NetBackup IT Analytics Report Reference Guide

Release 11.6

VERITAS
now part of Cohesity

NetBackup IT Analytics Report Reference Guide

Last updated: 2026-03-02

Legal Notice

Copyright © 2026 Cohesity, Inc. All rights reserved.

© 2026 Cohesity, Inc. All Rights Reserved. Cohesity, the Cohesity Logo and other Cohesity Marks are trademarks of Cohesity, Inc. in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. COHESITY SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Cohesity Support

21 September 2025

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Cohesity website.

Cohesity Services and Operations Readiness Tools (SORT)

Cohesity Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1 Introduction to NetBackup IT Analytics 26

Introduction to NetBackup IT Analytics 26
NetBackup IT Analytics Licensed Modules 27
NetBackup IT Analytics Components 28
Portal Server Software 29
Reporting Database 29
Data Collectors 29

Chapter 2 Alert Reports 31

Alert Reports 31
Alert Delivery Failure 32
Alert Detail 33
Alert Detail History 34
Alert Dashboard 35

Chapter 3 Ransomware reports 37

Ransomware reports overview 37
Ransomware Risk Assessment Dashboard 37
Ransomware Scorecard 40
Answer History 41
Item History 41
Override History 43

Chapter 4 Risk Mitigation Solution Reports 44

Introduction 45
Risk Mitigation Reports 45
    Risk Mitigation Analysis 45
    Risk Mitigation Trend 46
    Risk Mitigation Detail Reports 47
Backup Job Size Variance 47
Source Consecutive Failure 48
Source Overall Status Summary 49
Compliance RTO RPO 50
Host Multi-Pathing Exposure 52

Sources with No Recent Backups 52
Hot Array Ports 53
Hot LUNs by Read IO 54
Hot LUNs by Read Response 55
Hot LUNs by Write IO 56
Hot LUNs by Write Response 56
NetBackup Disk Pool Forecast 57
Thin Pool Forecast 58
VM Datastore Forecast 60
VM Guest Disk Forecast 61

Chapter 5 Storage Optimization Solution Reports 63

Introduction 64
Storage Optimization Overview 65
Storage Optimization Analysis 66
Storage Optimization by Tier 66
Storage Optimization Trend 67
Storage Optimization Detail Reports 68
Storage Optimization Dashboard 69
Unallocated LUNs by Array 71
Unallocated LUNs 71
Undiscovered LUNs by Array 72
Undiscovered LUNs 73
Unused LUNs by Array 74
Unused LUNs 75
Overprovisioned Hosts Summary 76
Overprovisioned Hosts 77
Inactive LUNS by Array 78
Inactive LUNS 79
VM Aged Snapshots by Datastore 80
VM Aged Snapshots 80
VMs Powered Off by VM Server 81
VMs Powered Off 82
VMs Low CPU by VM Server 83
VMs Low CPU 84
VMs Undiscovered Disks by VM Server 84
VMs Undiscovered Disks 85
VMs Not in VM Inventory by Datastore 86
VMs Not in Inventory 87
Non-VM Files by Datastore 87
Non-VM Files 88
Data Domain File Compression by Host 89

High Backup Retention Jobs by Sources 90
Sources Backed Up by Multiple Servers 90
Inactive Large Files by Host 91
Inactive Large Files 92
File Type Usage by Host 93
File Type Usage 93
AWS Orphan Snapshots 94
AWS Orphan Volumes by Account 95
AWS Orphan Volumes 96

Chapter 6 System Administration Reports 98

Audit Event Details 100
Audit Events Summary 101
Data Collection Message Summary (by group selection) 101
Data Collection Message Summary 102
Data Collection Message Detail 103
Data Collection Activity Detail 104
Data Collection Hourly Activity 105
Data Collection Performance Detail 106
Data Collection Performance Summary 107
Data Collection Performance Summary Trend 108
Data Collection Schedule Summary 109
Cases Not Represented in the Data Collection Schedule Summary 111
Example of a Data Collector Log Request 111
Data Collection Status Details 112
Data Collector Status Summary 114
Database Error Aggregation 115
Database Error Summary 116
Database Queries By Highest Execution Time 117
Database Queries By Total Execution Time 118
Database Queries Doing Full Table Scans 119
Duplicate Host Analysis 119
File Analytics Collection Status 123
NetApp Volume Status 125
License Summary 126
Oracle Job Overview 127
    Oracle Job Details 128
Capacity License Usage Details 129
File Analytics License Usage Details 130
Virtual Server License Usage Details 130
Fabric Manager License Usage Details 131

Backup License Usage Details 132
Replication License Usage Details 133
License Overage Detail 134
Report Activity Summary 135
Report Activity Detail 135
Scheduled Reports Summary 136
Disable or Delete a Scheduled Report 137
Disable a Scheduled Report 138
Delete a Scheduled Report 138
Modify a Schedule 138
System Health Check 138
Top Running Reports 141
Dashboard Considerations 142
Top Running Report Details 142
Version History 143
WMI Proxy Servers 144
Data Collector Version Updates Summary 145

Chapter 7 Capacity Manager Reports 146

Aggregate Mirror Summary 146
FlexClone Summary 147
HPE XP Efficiency Summary 148
HPE XP DP Pool Summary 149

Chapter 8 Application Capacity Reports 152

Application Storage Dashboard 152
Application Sub-Level of the Application Storage Dashboard 154
Applications at Risk 155
Over-Provisioned Applications 156

Chapter 9 Array Capacity and Utilization Reports 158

Array Capacity and Utilization 164
Array Capacity & Utilization (Generic Data) 165
Array/Cluster Summary 165
Physical Summary or Raw Capacity 166
Pool Summary 167
Volume/LUN Summary 167
Array Capacity & Utilization - (Dell Compellent View) 168
Array Capacity & Utilization - (EMC Data Domain Storage View) 170
Array Capacity & Utilization - (EMC VNX Celerra View) 173
Array Capacity & Utilization - (EMC VNX CLARiiON View) 176

Array Capacity & Utilization - (EMC VPLEX) 179
Array Capacity & Utilization - (EMC Isilon View) 181
Array Capacity & Utilization - (EMC Symmetrix View) 183
Array Capacity & Utilization (Hitachi Content Platform) 187
Array Capacity & Utilization (Hitachi Arrays & HDP View) 188
Array Capacity & Utilization (Hitachi NAS Arrays) 192
Array Capacity & Utilization - (Hitachi Vantara All-Flash and Hybrid Flash Storage) 219
HP 3PAR Overview 197
Array Capacity & Utilization (HP 3PAR View) 197
Array Summary 198
Physical Summary 199
Chunklets 200
Logical (LD) Summary 201
Common Provisioning Groups (CPG) 201
VV Summary 202
Array Capacity & Utilization (IBM Enterprise Array View) 202
Array Capacity & Utilization (IBM Modular Array View) 203
Array Capacity & Utilization (IBM SVC View) 203
How IBM SVC Capacity Values Are Derived 206
Array Capacity and Utilization (IBM XIV View) 207
XIV Array Summary 207
Physical Summary 208
Usable Volumes 208
Logical Volumes 209
Array Capacity and Utilization (NetApp View) 209
NetApp Summary 210
Physical Summary 211
Aggregate/Volume Capacity 212
LUN Capacity 213
Array Capacity and Utilization (NetApp Cluster) 214
NetApp Cluster Summary 214
Physical Summary 215
Aggregate Summary 216
Volume Capacity 217
LUN Summary 217
Performance 218
Array Capacity & Utilization - (Hitachi Vantara All-Flash and Hybrid Flash Storage) 219
NetApp Storage System Detail 221
Current Configuration 222
Cluster Status 223
Aggregate Utilization 223

Volume Utilization 223
Capacity Usage Bar Charts 224
Aggregate & Volume Capabilities 224
Array iSCSI Port Utilization 224
NetApp Cluster-Mode Array iSCSI Port Utilization 227
Array Capacity and Utilization (OpenStack Swift) 229
Cluster Summary 229
Physical Summary 230
Logical Summary 230
Performance 231
Disk Summary 231
RA Summary 235
Array Executive Summary 236
List Array Groups 239
HP EVA Physical Disks 240
Array FC Port Utilization 241
Data Domain Filesystem Compression 244
Data Domain System Summary - Storage 245
Data Domain System Details 246
Data Domain MTree Summary 249
Data Domain Replication History 250
Data Domain Snapshot History 251
EMC Isilon Disks 252
EMC Isilon HDFS Racks 254
EMC Isilon HDFS Roots 254
EMC Isilon License Detail 256
EMC Isilon NFS Exports 256
EMC Isilon Nodes 259
EMC Isilon Quota Summary 261
EMC Isilon SMB Share Summary 262
EMC Isilon Snapshot Summary 264
EMC Isilon Zones 265
EMC VNX (Celerra) Data Mover Summary 266
EMC VNX (Celerra) Virtual Data Mover Summary 267
EMC VNX (Celerra) File System Summary 268
EMC VNX (Celerra) SnapSure Summary 270
EMC VNX (Celerra) Storage Pools Summary 272
EMC VNX (Celerra) Storage Pools Detail 274
EMC VNX (Celerra) Tree Summary 276
EMC VNX (Celerra) Volume Summary 277
EMC VNX (Celerra) Volume Detail 278
EMC VNX (Celerra) Disk Summary 280
EMC VNX (Celerra) Storage System Detail 281

EMC VPLEX Backend Array Summary 282
EMC VPLEX Device Summary 284
EMC VPLEX Extent Summary 285
EMC VPLEX Port Summary 287
EMC VPLEX Storage View Summary 287
EMC VPLEX Storage Volume Summary 288
Object Store Namespace Summary 290
Node Summary 292
Object Store Tenant Summary 293
Hitachi NAS CIFS Share Summary 295
Hitachi NAS Disk Summary 296
Hitachi NAS EVS Summary 297
Hitachi NAS Ethernet Port Summary 298
Hitachi NAS File System Summary 298
Hitachi NAS NFS Export Summary 300
Hitachi NAS Snapshot Summary 302
Hitachi NAS Storage Pool Summary 303
Hitachi NAS Storage Server Summary 304
Hitachi NAS System Drive Summary 306
Hitachi NAS Tier Summary 308
Hitachi NAS Virtual Volume Summary 309
Overview of IBM Array Support 311
IBM Array Site Summary 312
Understanding IBM Enterprise Storage Capacity Calculations ...
3 1 2
IBM Array Summary 313
IBM Array Detail 314
IBM Storage System Summary 315
Raw Capacity 316
Array Site Capacity 317
Raw Array Capacity 317
Rank Capacity 318
Extent Pool Capacity 318
Logical Volumes 319
Top 3 of n Ranks 319
Top 3 of n Extent Pools 320
IBM Storage System Capacity Detail 320
IBM Storage System FB Volumes 320
Capacity and Object Usage Forecast 321
IBM Extent Pool Summary 321
IBM Rank Summary 322
IBM Disk Summary 323
IBM SAN Volume Controller Overview 324

Space-efficient VDisks 326
IBM SVC VDisk Summary 327
IBM SVC MDisk Groups 329
IBM SVC MDisk Group Details 330
IBM SVC MDisk Summary 332
How Capacity Values are Represented in XIV Reports 334
IBM XIV Consistency Groups Summary 334
IBM XIV Consistency Group Detail 335
XIV Physical Disks 337
XIV Snapshot Details 337
IBM XIV Storage Pools Summary 339
IBM XIV Storage Pool Details 341
LUN Utilization Summary 344
    Estimated LUN Usage/Availability 348
LUN Utilization Summary for EMC VPLEX Arrays 349
LUN Utilization Summary for IBM XIV Arrays 351
LUN Utilization Summary for Hitachi Vantara 353
LUN Utilization Summary for NetApp 7-Mode 355
LUN Utilization Summary for NetApp Cluster-Mode 356
LUN Utilization Summary for HP3PAR 358
Array Group Details 359
EMC CLARiiON Storage Array Group Specifics 361
Microsoft Azure Endpoint Summary 361
Microsoft Azure Storage Account Summary 362
NetApp Aggregate Summary 364
NetApp Cluster-Mode Aggregate Summary 365
NetApp Aggregate Detail 368
    Aggregate Capacity Usage 369
NetApp Cluster-Mode Aggregate Detail 370
    Aggregate Capacity Usage 371
NetApp Plex Details 373
    RAID Reconstruction Status 374
    Plex Disk Pool Details 374
NetApp CIFS Summary 374
NetApp Cluster-Mode CIFS Summary 375
NetApp Disk Summary 378
NetApp Cluster-Mode Disk Summary 379
NetApp Cluster-Mode FC Array Port Utilization 380
NetApp NFS Summary 383
NetApp Cluster-Mode NFS Summary 384
NetApp Cluster-Mode Node Summary 385
NetApp QTree Summary 389
NetApp Cluster-Mode QTree Summary 390

NetApp Snapshot Summary 392
NetApp Cluster-Mode Snapshot Summary 393
NetApp Volume Summary 395
NetApp Cluster-Mode Volume Summary 396
NetApp Volume Details 398
Volume Summary Status 400
Quota Status 400
Snapshot Schedule Details 400
Quota Details 401
NetApp Cluster-Mode Volume Detail 401
Quota Status 403
Quota Details 403
Snapshot Schedule Interval Details 404
Snapshot Schedule Cron Details 404
NetApp Cluster-Mode Summary 405
NetApp Cluster-Mode Thin Provisioning Summary 412
NetApp Cluster-Mode Vserver Summary 413
NetApp StorageGrid Active Alerts Summary 415
NetApp StorageGrid Node Summary 416
NetApp StorageGrid Site Summary 418
NetApp StorageGrid Tenant Summary 419
NetApp StorageGrid Users Summary 420
NetApp StorageGrid Buckets Summary 421
OpenStack Swift Container Summary 421
OpenStack Swift Policy Summary 423
OpenStack Swift Storage Utilization by Account 424
OpenStack Swift Storage Utilization by Zone 425
OpenStack Swift Storage Device Summary 426
OpenStack Swift Storage Utilization by Node 427
Hitachi Vantara Efficiency Summary 428

Chapter 10 Available/Reclaimable Capacity Reports 430

Allocated but Undiscovered LUNs 431
Allocated but Unused LUNs 432
Array Utilization Dashboard 435
Array Utilization Summary 437
Reclamation Summary 442
Reclaimable Space by Host Group 442
Reclaimable Space by Host Group Details 443
Reclaimable Space by Host 444
LUNS Added to Volume Group but Unused 444
Suspected De-provisioned Hosts 445

Unallocated LUNs 447
Unused Disks/Volume Groups 447
Unused Partitions/Logical Volumes 448

Chapter 11 Capacity at Risk Reports 450

Exception Summary 450
Exception Summary Report Details 451
Hosts at Risk 452
LUNs at Risk 453
NetApp Aggregates at Risk 455
NetApp Cluster-Mode Aggregates at Risk 456
NetApp Volumes at Risk 458
NetApp Cluster-Mode Volumes at Risk 459
Over-Provisioned Hosts 461
Over-Provisioned LUNs 462

Chapter 12 Capacity Chargeback Reports 464

Aggregated Chargeback 464
Aggregated Chargeback by Host Groups 465
Aggregated Chargeback by Attribute 466
Chargeback Detail 467
Chargeback Array Capacity 468
Chargeback By Host 469
Chargeback by LUN Detail 470
Chargeback Policy Capacity 471
Policy Detail 472
Object Store Chargeback 473
Object Store Chargeback Detail 475

Chapter 13 Host Capacity Utilization Reports 478

Host Capacity & Utilization 478
Host Utilization Detail 481
Host Details 482
Host Details by Source 483
Host Capacity Detail Report 484
Network Attached Storage Shares 485
Allocated and Unused LUNs 486
Disks Details 487
Disk Partitions/Volume Group Details 487
Host Groups 489
Host Utilization Detail (ASM Example) 490

Veritas Volume Details 492
Host Executive Summary 492
Host HBA Summary 494

Chapter 14 SnapMirror Reports 496

SnapMirror Summary 496
SnapMirror Details 498
SnapMirror Schedule 500
SnapMirror Status Summary 502
SnapMirror Data Transfer 503
SnapMirror Volumes/QTrees At Risk 503
SnapMirror Data Protection Dashboard 504

Chapter 15 SnapVault Reports 506

SnapVault Relationship Summary 506
SnapVault Secondary Configuration 508
SnapVault Summary 509
SnapVault Detail 510
Schedule Report (Primary and Secondary) 512
SnapVault Schedule 512
SnapVault Status 513
SnapVault Data Transfer 513
SnapVault Relationship Summary 514
SnapVault Data Protection Dashboard 515
SnapVault QTrees at Risk 516

Chapter 16 Capacity Forecasting Reports 517

Array Capacity Forecast 518
Data Domain Filesystem Capacity & Forecast 519
Data Domain Filesystem Status 520
EMC Symmetrix Thin Pool Capacity Forecast 521
EMC VNX (CLARiiON) Thin Pool Capacity Forecast 522
Object Store Usage Over Time 522
Hitachi NAS Storage Pool Capacity Over Time 523
Host Capacity Forecast 524
Host Volume Growth and Forecast 526
Host Storage Usage Over Time 527
Storage Capacity Over Time 528
IBM COS Storage Pool Growth and Forecast 529
IBM COS Vault Growth and Forecast 529
IBM COS Vault Summary Growth and Usage 530

IBM COS Site Summary 530
IBM COS Device Summary 530
IBM COS Mirror Summary 530

Chapter 17 Storage Performance Reports 532

Performance Profiles Overview 534
EMC Isilon Metrics 535
NetApp Cluster-Mode Metrics 535
EMC Symmetrix Enhanced Performance Metrics 535
Determining Performance States 535
Populating Performance States 536
Updating Performance Profiles 536
Aggregated Profiles 536
Array Performance by RAID Group 537
Array Performance by LUN 540
LUN Details 542
Data Domain Disk Performance 543
Data Domain System Performance 543
EMC Isilon CPU Performance by Cluster 544
EMC Isilon CPU Performance by Node 544
EMC Isilon Disk Performance by Cluster 545
EMC Isilon Disk Performance by Node 545
EMC Isilon File System Performance by Protocol 546
EMC Symmetrix Array Performance 546
EMC Symmetrix Device Group Performance 547
EMC Symmetrix Disk Performance 549
EMC Symmetrix FE Director Performance 550
EMC Symmetrix Storage Group Performance 551
EMC Symmetrix Thin Tier Performance 552
EMC VNX (Celerra) Data Mover Performance 553
Host Load Average 554
Host CPU Utilization 555
Host Memory Utilization 556
Host Memory Swap 557
Host Disk reads/writes 557
Host Network Packets sent/received 558
LUN Performance 559
Mission Control for Performance Analysis 560
Vendor-Specific Considerations for Array Performance Metrics 561
NetApp Cluster-Mode Aggregate Performance by Cluster 561
NetApp Cluster-Mode Aggregate Performance by Node 562

NetApp Cluster-Mode CIFS Performance by Cluster 562
NetApp Cluster-Mode CIFS Performance by Vserver 563
NetApp Cluster-Mode Disk Performance by Cluster 564
NetApp Cluster-Mode Disk Performance by Node 565
NetApp Cluster-Mode NFS Performance by Cluster 566
NetApp Cluster-Mode NFS Performance by Vserver 567
NetApp Cluster-Mode Volume Performance by Cluster 567
NetApp Cluster-Mode Volume Performance by Vserver 568
NetApp Cluster-Mode Performance by Cluster 569
NetApp Cluster-Mode Performance by Node 571
Performance Profiles 572
Port Throughput by Array 574

Chapter 18 Thin Provisioning Reports 576

Thin Provisioning Overview 576
EMC Symmetrix Compression Ratio Summary 577
EMC Symmetrix SRP Summary 579
EMC Symmetrix Thin Pool Summary 580
EMC Symmetrix Thin Pool Detail 581
EMC VNX (CLARiiON) Thin Pool Summary 583
EMC VNX (CLARiiON) Thin Pool Details 584
Hitachi Disk Summary 586
Hitachi Dynamic Provisioning Pool Summary 587
Hitachi Dynamic Provisioning Pool Utilization 590
    Reclaimable Storage 592
HP 3PAR Common Provisioning Group (CPG) Summary 592
HP 3PAR Common Provisioning Group (CPG) Detail 594
Storage Pool Summary 595
Hitachi Vantara DP Pool Summary 597
Hitachi Vantara Capacity by Storage 599

Chapter 19 File Analytics Reports 601

File Analytics Overview 602
Capacity Report 602
Capacity Forecast 603
Capacity Forecast by Device 604
Capacity Forecast by Volume/Share 605
Host Summary 605
Top Candidates for Duplicates 606
Volume Summary Report 607
Share Summary Report 608
New Files Added 609

Usage by Owner 609
Usage by Owner Forecast 610
Usage by Owner Forecast Details 611
Duplicate Files (top 50 in each device) 611
Deduplication Candidates 612
File Categories 613
Non-organizational Files 614
Ransomware File Types 615
File Sizes 616
File Types 617
Largest Files 617
File List Preview 618
Access Dates 622
Create Dates 623
Modified Dates 624

Chapter 20 Virtualization Manager Reports 626

Virtualization Manager Reports 627
Virtualization Terminology 628
Virtualization Manager Overview 629
Understanding the Datastore 630
    Datastore Characteristics 631
Using Virtualization Manager for Planning and Monitoring 631
VM Cluster Summary 632
VM Files Summary 634
VM Guest On-Prem Right-Sizing Detail by Host 636
VM Guest Disk R/W Profiles for Cloud Migration 637
Data Collection Options and VM Files Summary 639
VM Server Summary 639
VM Server Detail 642
    VM Server Detail Data 642
    VM Tags and Custom Attribute Details 644
    Performance & Forecast 644
    Largest VMs 644
    Datastores 645
    Disks 645
    Host Details by Source 645
    Data Collection Options and VM Server Detail 646
Microsoft Hyper-V Server Summary 646
VM Snapshot Summary 648
    Data Collection Options and VM Snapshot Summary 649
VM Storage Mapping Summary 649

VM Summary 650
VM Detail 652
VM Detail 653
VM Tags and Custom Attribute Details 654
Performance Over Time 654
VM Size Over Time 655
Volumes 655
Disks 655
VM Files 655
Snapshots 655
Host Details by Source 656
Microsoft Hyper-V VM Summary 657
VM Detail (Hyper-V) 658
VMs Over Time 660
Logical Disk Utilization 661
Datastore Usage Breakdown 662
VM Datastore Overview 664
Datastore Usage History and Performance 665
Datastore Usage History 665
Datastore Disk Read/Write Performance 665
Datastore Disk Latency Performance 666
Datastore Utilization Summary 666
Data Collection Options and Datastore Utilization 669
Datastore Detail 669
Datastore Detail: Datastore Usage Breakdown 671
Datastore Capacity & Forecast 672
Extents 672
Shared VM Servers 673
VM Utilization 673
Data Collection Options and Datastore Details 673
Physical Disk Utilization 673
Physical Disk Detail 675
Datastore Performance Summary 676
Physical Disk Performance Over Time 678
Physical Disk Performance Summary 679
VM Guest to LUN Performance 680
VM Performance Over Time 683
VM Performance Summary 683
Datastore Capacity Forecast 689
VM Size Forecast 690

Chapter 21 Fabric Manager Reports 691

Host to Storage Dashboard 692
Working with Topology 692
Paths Portlet 693
Paths in the Host to Array Connectivity Portlet 694
Zones in the Host to Array Connectivity Portlet 694
Port Summary 695
Port Details 696
SAN Fabric Summary 697
SAN Fabric Details 698
Switch Summary 700
Switch Details 701
Node Port Summary 704
Zone Summary 705
Zone Details 706
Zone Set Summary 708
Zone Members 708
Bottom Ranking 709
Port Distribution 710
SAN Resource Utilization 711
Top Ranking 712
SAN Fabric Change Status Summary 713
SAN Fabric Change Summary 714
Port Performance Dashboard 714
Ports Portlet 715
Port Details Portlet 716
Port Performance Summary 717
Port Performance Details 718
SAN Fabric Usage 719
Port Error Rate Forecast 720
Port Utilization Forecast 720

Chapter 22 Backup Manager Management Reports 721

Using Backup Manager Reports for Analysis 723
Anomalies Summary 724
Backup Executive Summary 727
Backup Server Performance Dashboard 728
Client Summary Report 731
Failed Backup Sources 732
Consecutive Errors 732
Data Domain Backup Summary 734
Data Domain File Compression Detail 734

Data Domain File Compression Summary 737
Deduplication Size Savings Report (Historical) 738
Error Log Summary 739
    Error Log Summary 739
    By Policy 739
    By Server 740
HP Data Protector Session Summary 741
HP Data Protector Session Details 742
Job Duration Report 744
    To diagnose excessive backup durations: 745
Job Duration by Source 746
Job Error Code 746
Job Status Summary 747
Job Status Summary by Client 749
Job Summary Report 750
Job Summary - Heterogeneous Backup Products 751
Job Throughput by Client 752
Avamar Job Summary Report (Homogeneous) 754
Backup Exec Job Summary Report (Homogeneous) 755
Cohesity DataProtect Job Summary Report (Homogeneous) 756
Commvault Simpana Job Summary Report (Homogeneous) 758
Commvault Simpana ESX Backup Summary 760
HP Data Protector Job Summary Report (Homogeneous) 761
IBM Spectrum Protect (TSM) Job Summary Report (Homogeneous) 762
NetBackup Job Summary Report (Homogeneous) 763
NetWorker Job Summary Report (Homogeneous) 766
Oracle Recovery Manager (RMAN) Job Summary Report (Homogeneous) 767
Veeam Backup & Replication Job Summary Report (Homogeneous) 768
    Job Summary Advanced Options 769
Job Details Report 770
Avamar Job Details 773
Commvault Simpana Job Details 776
IBM Spectrum Protect (TSM) Job Details Report 779
Veeam and RMAN Job Details Report 781
    Related Jobs 784
    Fragments 785
    Backup Items 786
    Job Messages 787
Adding a Note to a Job 787
    To add a note to a job 787

Job Summary by Server 788
Job Summary by Source 789
Job Type Count 790
Job Volume Summary Report 792
Use this procedure to analyze volume usage trends: 792
Largest Backup Volume 793
Primary Server Job Throughput 794
Monthly Backup Summary 795
NetBackup AIR-Replication Import Jobs 796
NetBackup deduplication to MSDP savings 798
NetBackup Deduplication to MSDP Savings Dashboard 798
NetBackup Deduplication to MSDP Savings - By Primary Servers 800
NetBackup Deduplication to MSDP Savings - By Clients 801
NetBackup Deduplication to MSDP Savings Trend Over Time 803
NetBackup Deduplication to MSDP Savings - By Policy Type 803
NetBackup Media Server Job Throughput 805
NetBackup SLP Status 806
NetBackup SLP Details 809
NetBackup SLP Status by Destination (AIR Activity) 811
NetBackup SLP Status by Client 813
NetBackup SLP Status by Image Copy 817
NetBackup SLP Status by Image 820
NetBackup SLP Status by SLP 823
Running and Queued Job Summary 827
NetBackup Running and Queued Jobs Summary 828
NetWorker Running and Queued Jobs Summary 828

Chapter 23 Backup Administration Reports 830

Backup RTO RPO Compliance 831
Avamar Grid Capacity Dashboard 833
Command Center Dashboard 834
Data Protection Dashboard 837
Host Details 838
Host Aliases 839
Edit Host Details 840
Host Details by Source 841
Host Group 842
IBM Spectrum Protect (TSM) Storage Pools Dashboard 843
Storage Pool Capacity 844

Storage Pool Attributes 844
Storage Pool Utilization 845
Tape Occupancy Summary 845
Tape Media Summary 845
Storage Pool Process History 845
Job Histogram 845
Job Throughput Histogram 847
Job Activity Histogram 848
Jobs Dashboard 848
Message of the Day Report 848
Mission Control - Backup 849
NetBackup Event Notification Summary - CRITICAL 853
NetBackup License Report 854
NetBackup Host Certificates 856
NetBackup Non-Host Certificates 858
Object Details 860
Operations Dashboard 862
Tailoring the Report Output by Backup Product 864
TSM Process Summary 865
NetBackup Audit Report 865
EEBs on Primary server, Media server and Clients reports 868
NetBackup Host Detail Installed EEBs and Packages 869
NetBackup Host Detail Installed EEBs 870
NetBackup Host Detail Installed Packages 870

Chapter 24 Backup Media Management Reports 871

Current Media Summary 871
Data Domain VTL Errors 872
Data Domain VTL Media Summary 873
Data Domain VTL Performance 874
Tape Media Summary Report 875
Tape Media Detail Report 877
Commvault Simpana Tape Media Detail 879
HP Data Protector Tape Media Table 881
NetBackup Tape Media Table 881
NetWorker Tape Media Table 881
TSM Tape Media Detail 882
TSM Tape Media Detail Table 883
Estimated Capacity Notes 883

Chapter 25 Backup Service Level Agreement (SLA) Reports 884

Available Service Level Agreement (SLA) Reports 884
Configure a Backup SLA Policy 885
Backup Start Time SLA 885
Determining and Improving Backup Start Time Performance 886
To determine and improve backup start time performance 887
Backup Status SLA 888
Changing the Success Percentage 888
Backup Primary Server Overview 888
Determining and Improving Backup Success Performance 890
To determine and improve backup success performance 890
Backup Duration SLA 890
Determining and Improving Backup Duration Performance 891
To determine and improve backup duration performance 892
SLA Job Summary 893

Chapter 26 Backup Storage Utilization Reports 895

Using Storage Utilization Reports 896
Data Domain System Summary - Backup 896
Data Domain System Details 898
Data Domain Disk Summary 900
Disk Usage and Performance Report 901
Drive Performance Summary Report 904
NetBackup Drive Performance Summary 906
Drive Performance Detail 907
Drive Utilization & Performance Report 908
NetBackup Deduplication Effect 910
NetBackup Deduplication Savings Trend 911
NetBackup Storage Unit Summary 911
NetBackup Storage Unit Usage 912
Real Time Library and Drive Status 913
Storage Unit Detail Report 914
Storage Unit Images 917
Storage Unit Event Details 918
Disk Volume Details 919
Tape Library and Drive Utilization 920
IBM Spectrum Protect (TSM) Database Utilization 921
IBM Spectrum Protect (TSM) Recovery Log Utilization 922

Chapter 27 Backup Manager Forecasting Reports 923

Forecasting Overview 923
NetBackup Library Capacity Forecast 924
Media Availability Forecast 925
Media Consumption Forecast 927
Media Usage Details 928
Scratch Pool Forecast Report 929
Tape Drive Usage and Forecast Report 930

Chapter 28 Backup Billing and Usage Reports 932

Available Billing and Usage Reports 932
Determining Media Chargebacks 933
To determine chargebacks for a business unit or customer 933
Determining Backup Chargebacks 934
To determine server consumption chargebacks 934
Billing and Chargeback Summary 934
NetBackup Front End Size by Duration and Workload 936
NetBackup Front End Size by Workload 937
Server Consumption Summary 938
NetBackup Front End Size Usage Summary 939
NetBackup Client Front End Size Report 940
NetBackup Policy Front End Size Report 941
NetBackup Images Report 941

Chapter 29 Backup Policies Reports 942

Commvault Simpana Policy Summary Report 942
Commvault Simpana Policy Backup Copy Details 944
EMC Avamar Groups 946
Avamar Server Detail 947
Avamar Group Detail 948
Clients for Group 949
Avamar Client Detail 950
Avamar Dataset Detail 953
Avamar Schedule Detail 954
Avamar Retention Policy Detail 955
EMC NetWorker Groups 957
Group Detail 958
Client Instance Detail 960
HP Data Protector Backup Specification Summary 961
HP Data Protector Backup Specification Detail 961
IBM Spectrum Protect (TSM) Policy Domains 964

Policy Domain Details 965
NetBackup Job Status Summary by Policy Type 966
Veritas NetBackup Policy 967
Backup Policy Details 968
Veritas Netbackup protection plan report 971

Chapter 30 Public Cloud Reports 974

About Public Cloud Reports 974
AWS Reports 974
AWS Account Billing Detail 975
AWS Account Listing 977
AWS Billing Summary Over Time 978
AWS Product Summary 979
AWS S3 Bucket Summary 980
AWS EC2 Snapshot 981
AWS EC2 Snapshot Detail 984
AWS EC2 Instance Summary 985
Microsoft Azure Reports 987
Azure Cloud Cost Spend Dashboard 987
Azure Daily Usage Future Forecast by Org Level 988
Azure Daily Usage Forecast by Subscription Details 990
Cost History By Region - Top N 991
Cost History by Service (Meter Name/Category) - Top N 991
Cost History by Subscription - Top N 992
Cost History by Tags 992
Microsoft Azure Virtual Machine Summary 992
Microsoft Azure Virtual Machine Detail 993
Microsoft Azure Billing Summary 994
Microsoft Azure Resource Group Summary 995
Microsoft Azure Subscription Summary 996
Microsoft Azure Unattached Disks Summary 997

Chapter 1

Introduction to NetBackup IT Analytics

This chapter includes the following topics:

- Introduction to NetBackup IT Analytics
- NetBackup IT Analytics Licensed Modules
- NetBackup IT Analytics Components
- Portal Server Software
- Reporting Database
- Data Collectors

Introduction to NetBackup IT Analytics

NetBackup IT Analytics provides storage reporting solutions for consolidated visibility into multi-vendor, complex storage environments, enabling you to increase data reliability and reduce costs with consolidated reporting and automated alerts.

NetBackup IT Analytics allows you to drill down and identify the root cause of performance issues, identify re-tiering opportunities, and effectively implement chargebacks so business units understand their overall storage environment.

Data collectors send relevant storage component information to the NetBackup IT Analytics platform, an integrated Oracle DB that maintains current and historical data on the entire storage environment. From this database, each NetBackup IT Analytics product can analyze and correlate relevant information. The result is consistent analysis and reporting, enabling you to manage all aspects of your IT environment.

NetBackup IT Analytics Licensed Modules

NetBackup IT Analytics offers several major licensed modules:

Backup Manager provides a centralized, real-time view of your backup environment. These reports increase the reliability of your data protection environment by providing the most in-depth backup reporting and management available. You can instantly view backup job status across thousands of clients in the enterprise. Benefits include:

- View successful, partial, and failed backups across the entire data protection infrastructure and discover the root cause of problems
- Institute best practices for data protection compliance and internal SLAs
- Increase reliability and performance of the backup/recovery environment

Capacity Manager increases storage utilization by providing a full view of the SAN storage environment from the applications, hosts, and storage arrays. Automated mapping shows the relationship between each host and LUN in your environment to convey how the storage resources are being used. Array utilization shows the amount of allocated storage from each storage array. Application utilization identifies the amount of storage used vs. allocated by application. Benefits include:

- Create chargeback reports by any logical grouping (for example, by department, geography, tiers)
- Optimize efficiency and reclaim hidden storage with allocated and used storage capacity views
- Improve planning with forecasts and trend analysis of future storage capacity

Fabric Manager provides a view of SAN Fabric components--zones, switches, ports, arrays, and hosts--enabling a full view of your resources. A variety of reports offer a quick view of capacity and utilization, along with change management and performance metrics. Benefits include:

- Visualize the SAN from the physical or virtual server to the storage system
- Determine if errors are creating issues
- View I/O response times, throughput, and read/write statistics

File Analytics collects and profiles unstructured data to enable you to identify storage that can be reclaimed. Three types of Data Collectors collect and categorize this data from: CIFS Shares, NetApp storage systems, and from host-attached storage. Benefits include:

- Consolidate and remove duplicate files for faster, optimized data backups
- Move data to specific tiers based on the value of the data
- Enforce corporate retention and unauthorized usage policies

Replication Manager can help you protect mission-critical data by measuring compliance, maximizing resources, controlling costs, and improving overall performance. Identify wasted storage, occupied by obsolete snapshot files, and reclaim space. Plan for growth as replication demands increase. Issue warnings when replication storage capacity becomes dangerously low. Benefits include:

- Measure compliance with business continuity goals and control costs associated with data replication environments

Virtualization Manager enables you to gain insight into how to plan, manage, and optimize storage allocated to virtual systems--improve efficiency and reduce costs. Directly map the storage used by each virtual machine, down to the actual storage array itself. Determine which virtual machines are over- or under-utilizing their allocated storage. Benefits include:

- Maximizing the storage efficiency in virtualized environments
- Improve performance of the virtualized infrastructure
- Monitor current usage and forecasting of virtual resource demands

NetBackup IT Analytics Components

When you install NetBackup IT Analytics, you are installing the following components:

- Portal Server. The physical server on which the NetBackup IT Analytics Portal Server software resides.
- Portal Server Software. The binaries, SQL scripts, configuration files, and open-source and third-party software products needed to retrieve and render reporting data from the Reporting Database.
- Reporting Database. The Oracle database stores all report data. The Reporting Database is usually installed on the Portal Server, but you can just as easily install it on a separate server, preferably a dedicated database server. These binaries are installed during the first step of the installation procedure.
- Data Collectors. The software that collects report data about your backup servers and storage arrays. The Data Collector is usually installed on a separate server, but you can just as easily install it on the Portal Server, with one exception: if you use Veritas NetBackup you must install it on each Primary Server.

Because of the relationship between the Portal Server, the Portal Server software, and the Reporting Database, you need to install these Portal Server components at the same time. However, you can install the Data Collector separately.

Portal Server Software

When you install Portal Server software, you install a compilation of software:

- Open-Source and Third-Party Software provide framework components.
- Binaries to provide NetBackup IT Analytics its intelligence and functionality. These binaries include the Portal engine, the Data Receiver, and stored procedures.
  - The Portal engine drives the Portal user interface.
  - The Data Receiver accepts data from the Data Collector, then saves the data to the Reporting Database.
  - The stored procedures retrieve data from the Reporting Database, input data into the Reporting Database, and perform calculations and analysis.
- Database SQL Scripts instruct the Oracle SQL scripts to create NetBackup IT Analytics's object module and database schema.

Reporting Database

The Reporting Database stores all report data (as metadata) that it receives from the Data Receiver. You can install the Reporting Database on a standalone server or on the same server as the Portal Server. If you use a Managed Services Provider, the Reporting Database and Portal Server are usually off site and hosted by that third-party provider.

NetBackup IT Analytics maintains most of the data in the Reporting Database indefinitely. Some reports are more valuable when they have access to historical data. Because the Reporting Database only stores metadata, the amount of data on the Reporting Database is relatively small (GBs). This report data is essential to retrieving reports, so the Portal Installation Guide walks you through a database backup process.

Data Collectors

The Data Collector is a centralized and remotely managed data collection mechanism. This Java application is responsible for gathering metadata from storage arrays, hosts, switches, and backup servers.

The Data Collector continuously collects data and sends this data, using an http or https connection, to another Java application, the Data Receiver. The Data Receiver runs on the Portal Server and stores the data that it receives in the Reporting Database. When you use the Portal to generate a report, the Portal requests this

information from the Reporting Database, then returns the results in one of the many available reports that support storage management and data protection.

The Data Collector obtains all of its monitoring rules from a Data Collector Configuration File. This file resides in the Reporting Database in XML format. When the Data Collector first starts, it downloads this file from the Reporting Database. The Data Collector uses this file to determine the list of enterprise objects--such as backup servers, hosts, or storage arrays--that are to be monitored and included in its data collection process.

A single installation of the Data Collector supports any number of servers. The only real limitation is the memory and CPU processing power of the server on which the Data Collector resides. See the Certified Configurations Guide for Data Collector system requirements.

Chapter 2

Alert Reports

This chapter includes the following topics:

- Alert Reports
- Alert Delivery Failure
- Alert Detail
- Alert Detail History
- Alert Dashboard

Alert Reports

NetBackup IT Analytics has identified many common trouble spots and built a set of pre-defined alerting rules to examine conditions within your enterprise. Rules are the first step in adding and editing Alert policies.

The Alert Policy defines the scope of what you'd like to monitor, such as a specific product or job types, the frequency of the symptom and alerting thresholds.

These rules are monitored in real-time as the data is collected or on a scheduled basis if the condition is time-based. After a Policy is created, you can choose to alert on the conditions you've specified or use reports to review the results on-demand and drill down to specifics. Because alerts can be suppressed for a variety of reasons, use reports to review the results of your monitoring criteria whenever required.

- Alert Delivery Failure - The Alert Delivery Failure lists details for each alert delivery failure by Alert Policy.
- Alert Detail - The Alert Detail allows you to view objects with alerts triggered by thresholds set in your Alert Policies. You can also use the report to suppress/unsuppress individual alerts. When running the report, you have the option to see all alerts or to filter them based on if they are suppressed or not.

- Alert Detail History - The Alert Detail History provides a summary of alerts based on definitions in the Alert Policies for a selected time period. This report includes the most recent alerts and past alerts. This comprehensive view by Alert Policy can be filtered by object types and Alert Rules.
- Alert Summary - The Alert Summary is an aggregated view of current alerts based on Alert Policy definitions. Filter this donut chart by object types and rules.
- Alert Summary History - The Alert Summary History is presented as a donut chart that displays all alerts by Alert Policy and your selected time period. Report can be filtered by object types and Alert Rules.
- Alert Trend - The Alert Trend is presented as a bar chart that displays all alerts by alert type (warning and critical) for a selected time period. If required, filter the report by object type and Alert Rules.

Alert Delivery Failure

Alert Delivery Failure lists details for each alert delivery failure by Alert Policy.

Table 2-1 Alert Delivery Failure report

Column name	Description
Policy Name	User defined Alert Policy name.
Alert Date	Date alert was triggered.
Status	Alert status: critical or warning.
Critical	
Email List	Lists email addresses defined for critical alert delivery.
Script	Location and name of user-created script for critical alert delivery mechanism.
SNMP Security Name	Username which is configured to send the SNMP trap.
SNMP Server Name	Name of the server used to receive the critical alert.
SNMP Version	Version of SNMP used to capture the critical alert.
Syslog Server	Displays the IP address of the syslog server.
Warning	
Email List	Lists email addresses defined for warning alert delivery.

Table 2-1 Alert Delivery Failure report *(continued)*

Column name	Description
Script	Location and name of user-created script for warning alert delivery mechanism.
SNMP Security Name	Username which is configured to send the SNMP trap.
SNMP Server Name	Name of the server used to receive the critical alert.
SNMP Version	Version of SNMP used to capture the critical alert.
Syslog Server	Displays the IP address of the syslog server.

Alert Detail

The Alert Detail allows you to view objects with alerts triggered by thresholds set in your Alert Policies. You can also use the report to suppress/unsuppress individual alerts. When running the report, you have the option to see all alerts or to filter them based on if they are suppressed or not.

To suppress or unsuppress an alert, click the **Row Selection** check box and click the relevant action button. When you suppress an alert, you select the duration for the suppression.

Table 2-2 Alert Detail

Column name	Description
Row1Column1	Row1Column2
Row Selection	Click the check box to select the alert. Use the **Suppress** and **Unsuppress** buttons to complete the action once the alert is selected.
Policy Name	User defined Alert Policy name. Use the link to access the policy directly
Alert Object Type	Defines the object type for which the alert was set. Object type dictates which drill down report is displayed.
Alert Object	Displays the object on which the alert was issued. Click the object to display the relevant report and the details for the object.
Parent Object Type	If a hierarchy structure exists, displays the object type of the parent for the alert object.

Table 2-2 Alert Detail *(continued)*

Column name	Description
Parent Object	If a hierarchy structure exists, displays the name of parent of the alert object.
Last Alert Message	Displays a brief description of the alert condition.
Status	Indicates the severity of the alert. Click the link to drilldown to view additional information about the root cause of the alert.
Alert Date	Displays the date the alert condition occurred.
Latest Alert Date	Displays the date of the latest alert that occurred for the alert policy and alert object combination. When this column has a value, the **Alert Date** column indicates the date when this alert occurred for the first time.
Occurrences	Displays the number of times the alert was processed on the same alert object in the defined number of **Look Back hours**.
Suppressed Until Date	Displays the date/time the alert is suppressed until. When you suppress an alert, you select a duration. This field displays that value.
Suppressed By	Displays the LDAP ID of the user responsible for the Alert suppression.
ServiceNow Ticket Number	Displays the value of **ServiceNow: Job Finalized** alert when ServiceNow integration is fully enabled.

Alert Detail History

The Alert Detail History provides a summary of alerts based on definitions in the Alert Policies for a selected time period. This report includes the most recent alerts and past alerts. This comprehensive view by Alert Policy can be filtered by object types and Alert Rules.

Table 2-3 Alert Detail History

Column name	Description
Policy Name	User defined Monitor Policy name. Use the link to access the policy directly
Alert Object Type	Defines the object type for which the alert was set. Object type dictates which drilldown report is displayed.
Alert Object	Displays the object on which the alert was issued. Click the object to display the relevant report and the details for the object.

Table 2-3 Alert Detail History *(continued)*

Column name	Description
Parent Object Type	If a hierarchy structure exists, displays the object type of the parent for the alert object.
Parent Object	If a hierarchy structure exists, displays the name of parent of the alert object.
Alert Date	Date/time the alert was logged into the system.
Alert Message	Displays the alert message as determined by the threshold settings.
Status	Indicates the severity of the alert. Click the link to drill down to view additional information about the root cause of the alert.

Alert Dashboard

This dashboard provides an overview of current and historical alert conditions. It provides a more detailed drilldown view that enables you to actively manage the alerts. Each chart on the dashboard is a report briefly explained in the sections below.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Alert Summary

This is a donut chart that displays all alerts by the Alert Policy, with drilldown to Alert Detail. You can filter the report by object types and alert rules.

Alert Trend

The bar chart displays all current alerts by alert type for a selected time period. You can filter the report by object type and alert rules or drill down to the Alert Summary History and view the alerts by policy.

Alert Summary History

Graphic representation of all historical alerts by Alert Policy. You can filter the report by object types and alert rules. Drill down to the Alert Detail History to view all alerts within the policy. Comparing this with the Alert Summary report lets you evaluate how the current alert breakdown compares to the historical.

Critical Alerts

Lists a summary of current alerts based on definitions in the Alert Policies and sorted by Policy. Report can be filtered by object types. You can **Suppress** or **Unsuppress** one or more alerts inline. Use this report to view objects with alerts triggered by thresholds set in the Alert Policy. Drilldown to view root causes for the alert.

Note: Alerts must be enabled and configured in order to populate the Alert Dashboard.

See “Alert Reports” on page 31.

Chapter 3

Ransomware reports

This chapter includes the following topics:

- Ransomware reports overview
- Ransomware Risk Assessment Dashboard
- Ransomware Scorecard
- Answer History
- Item History
- Override History

Ransomware reports overview

Ransomware reports consist of a user-friendly dashboard and scorecard that emphasize the vulnerability of assets within user environments. The Ransomware Risk Assessment Dashboard provides insights into risk mitigation trends by considering various factors, such as anomaly detection, inadequate or missing recent backups, and identification of potential ransomware files as a warning signal. On the other hand, the Ransomware Scorecard enables you to assess your overall resilience and recoverability against ransomware attacks through a series of evaluation questions and data points.

Ransomware Risk Assessment Dashboard

The integrated dashboard emphasizes the potential risks associated with inadequately protected assets. In the event of a ransomware threat, these assets may become inaccessible, leading to potential financial losses for your enterprise. Additionally, it identifies any files within your environment that bear well-known ransomware extensions. The dashboard offers a range of reports, effectively

showcasing and visualizing the risks present in your environment. This allows you to assess your level of preparedness in terms of data protection in case of a ransomware infection. To enhance clarity, the dashboard is divided into sections, with each section representing a specific report.

Risk Mitigation Analysis

This report in the dashboard describes the risk associated with several key aspects of your protection environment. The report lists servers which are divided across risk categories, such as:

- Sources with No Recent Backups: Lists all servers which have not been backed up in the recent past.
- Sources with No Recent Backups: Lists all servers which have not been backed up in the recent past.
- Suspect Backup by Job Size: Indicates servers whose backup job resulted in a very small size backup which indicates some misconfiguration.
- Sources Consecutive Failure: Describes the risk of backup jobs failing for certain servers and that means some servers may not be adequately protected.
- Compliance RTO RPO: Assists in computing RTO (Recovery Time Objectives) and RPO (Recovery Point Objectives) for backups using the time of the last full backup and then adding the time it takes to apply any incrementals.

See “Risk Mitigation Analysis ” on page 45.

You can click on the donut section of your choice in the chart to get a detailed view of the servers and sources under consideration for that selection. This drill down capability provides deeper insight and helps identify the sources of risk quickly and easily. To choose your categories to display, please hover your mouse over the top right part of the report to display icons that let you "edit scope". These rules and categories are also described in **Admin** > **Solutions** > **Risk Mitigation**, where you can double-click on any rule to customize it for your requirements. You can choose to view multiple categories at the same time or select the categories by clicking on the category legend without changing the scope of the report for quick view and analyze situations.

Risk Mitigation Trend

This report describes the same risk categories spread across time. This report gives you a sense of how your protection risk has been changing over time. This helps you create a baseline and work towards keeping your risk levels per category below this baseline.

See “Risk Mitigation Trend ” on page 46.

Alerts Summary Report

This report describes specific instances when things have not worked as expected with respect to data protection. A NetBackup Primary server connection failure, out of storage condition on the NetBackup Primary Server or an incomplete backup job may indicate risk in the environment that needs addressing. The alert summary chart can be used to drill down on any of the alert types to get deeper insight into the source of alert. This helps in mitigating any potential risk in the environment and consequently reducing a potential negative impact of a Ransomware infection.

Alert Trend Report

It describes the alert occurrences over time. This provides a good way of determining a baseline and staying below the baseline as a goal to mitigate any risk.

Ransomware File Types

This report describes the risk associated with the presence of files across the enterprise which have well known ransomware extension names. This helps to zoom in on servers which have files with potential ransomware infection. This visibility is crucial in identifying the spread of ransomware infection across the environment. This report provides the summary view of number of the such file extensions found. You can click on the export icon to download a CSV file which would contain a list of all file paths along with their server locations to help drill down into the source of the infection.

Note: This report is available with both Protection and Complete license suites for NetBackup IT Analytics.

You can edit the report scope by hovering over the top right hand corner of the report and choosing the **Edit Scope** option and choosing which Alert category to see in the report. The alerts can be configured in detail using the top level **Alerts** tab.

Job Anomalies Count

This chart displays the count of anomalies detected in the host group by it data protection tool. The chart drills down to the Anomalies Summary report that contains additional details about the hosts and the anomaly score.

See “Anomalies Summary” on page 724.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Ransomware Scorecard

Ransomware Scorecard is an interactive report that provides a view of your overall ransomware resilience and recoverability preparedness. The report rates your preparedness based on user responses and data points, and recommends improvements to get the best ransomware protection. See *Manage ransomware scorecard* section in *NetBackup IT Analytics User Guide* for detail on scorecard access privileges, configuration, and components.

The report is located here:

Reports > **Ransomware** > **Ransomware Scorecard**

The values displayed on the scorecard depend on the scope selection of the scorecard. You can use the Ransomware Scorecard Scope Selector to get the required scorecard view.

Table 3-1 Ransomware Scorecard components

Component	Description
Ransomware Score (%)	Ransomware score in percentage based on the user responses to the Question type queries and statistics reported by Data type queries. It is the sum of actual Score values of each query divided by the sum of maximum possible Score values of each query expressed in percentage.
Score Trend (%)	Ransomware recoverability trend over time specified in the report scope. Mouse-over each data point to view additional details on the trend line, such as the maximum score, percentage score, impacting query, and event.
Completed Queries (%)	Indicates the extent to which the ransomware preparedness evaluation is complete in percentage. A higher percentage reflects a better overall score.
Type	Indicates whether the query type is a Question or a Data input received from reports.
Query	The actual query text.
Result	Responses received from users or values derived from associated reports.
Best Practice Recommendation	Standard best practices recommended for the respective query.

Table 3-1 Ransomware Scorecard components *(continued)*

Component	Description
Score	Indicates the quality of the response submitted for the query. It is the sum of individual scores multiplied by the weight of the query.
Risk	Indicates the extent of risk associated with the query based on the submitted answer. Weak or Ineffective score for a query with high Weight results in high Risk and is represented by a longer bar. Significant risk is indicated with an added triangle next to the bar.
Edit Scope	Click to open the **Ransomware Scorecard Scope Selector** and edit the scorecard scope to get the required view.

Answer History

This report is invoked from the Ransomware Scorecard and is specific to the responses submitted for the query by the user. The details displayed in the report are described below.

Table 3-2 Answer history

Column name	Description
Query	The assessment query for which the user has responded.
Answer	User's response to the query
Deleted	Indicates whether the query has been deleted.
Event	Indicates whether the assessment event was an update, delete, or response.
Notes	Notes added by the user during the events.
Modified By	Username against which the event was attributed.
Modify Date	Timestamp of the affected change.

Item History

This report is invoked from Ransomware Scorecard and is specific to the change history of the query for which it is invoked. The report contains the trail of changes

to the overall query, such as Description or Weight as well as changes to each of the possible answers.

The report consists of two tables - first describing changes to the query and the second contains user responses over time.

Table 3-3 Item Metadata History

Column name	Description
Type	Indicates the type of assessment query. It can be either Data or Question.
Query	The assessment query that has undergone changes. The column also displays the text changes in the query.
Description	The column displays the text changes in the query description.
Best Practice Recommendation	Displays the text changes in recommendations over time.
Weight	Indicates the change in weight assigned to the query.
More Info Link	Displays changes to the hyperlink to additional information on the query.
Enabled	Indicates the enabled/disabled status of the query.
Deleted	Indicates whether the query was deleted.
Event	Records the type of change event.
Modified By	Displays the username that affected the changes.
Modify Date	Timestamp of the change event.

The table below describes the details displayed in the Answer History table.

Table 3-4 Item Answer History

Column name	Description
Query	The query for which the changes in responses are tracked.
Answer	Displays the change in responses over time.
Score	Indicates the change in score with change in responses.
Deleted	Indicates whether the response was deleted.
Event	Record of the event type.

Table 3-4 Item Answer History *(continued)*

Column name	Description
Modified By	Displays the username that affected the changes.
Modify Date	Timestamp of the affected change.

Override History

This report is invoked from Ransomware Scorecard and is specific to the override change history of the query for which it is invoked. This report can be invoked only for queries of type Data.

The table below describes the details displayed in the report.

Table 3-5 Item Answer History

Column name	Description
Query	The query for which the changes in responses are tracked.
Deleted	Indicates whether the response was deleted.
Event	Record of the event type.
Notes	Notes supplied by the user while performing the override
Override	Indicates where any actual override was performed.
Report Value	The report value when the override was set.
Modified By	Displays the username that affected the override.
Modify Date	Timestamp of the affected override.

Chapter 4

Risk Mitigation Solution Reports

This chapter includes the following topics:

- Introduction
- Risk Mitigation Reports
- Backup Job Size Variance
- Source Consecutive Failure
- Source Overall Status Summary
- Compliance RTO RPO
- Host Multi-Pathing Exposure
- Sources with No Recent Backups
- Hot Array Ports
- Hot LUNs by Read IO
- Hot LUNs by Read Response
- Hot LUNs by Write IO
- Hot LUNs by Write Response
- NetBackup Disk Pool Forecast
- Thin Pool Forecast
- VM Datastore Forecast

- VM Guest Disk Forecast

Introduction

Risk mitigation report templates provides actionable analytics, enabling administrators to identify and mitigate the risk profile of the datacenter. By providing visibility to the existing and potential risks associated with the objects, administrators can take appropriate action to mitigate potential dataloss, outages and performance issues.

The reports in this section employ Risk Mitigation Rules to assess areas within the enterprise that offer candidates for potential risk. These rules include parameters that can be configured to isolate specific conditions relevant for your environment.

Risk Mitigation Reports

The reports discussed in this section employ Risk Mitigation Rules to assess areas within your enterprise that may require attention to avoid performance issues and outages. These rules include parameters that can be configured to isolate specific conditions relevant for your environment.

Risk Mitigation Analysis

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

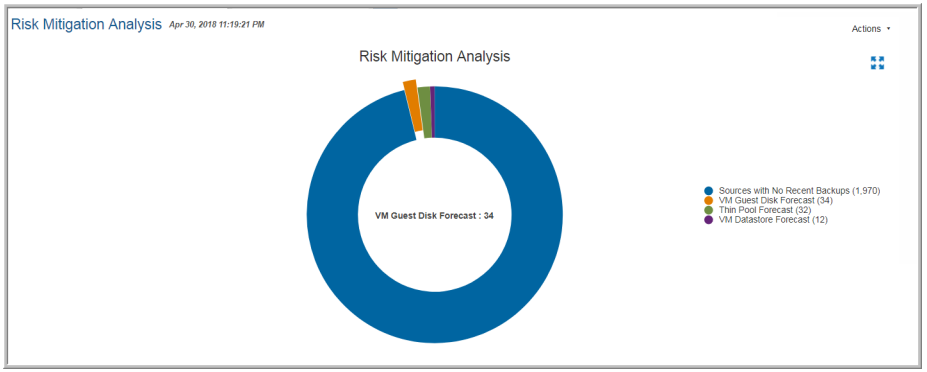
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

View a chart of infrastructure categories and parameters that may offer opportunities for storage optimization. Click a chart sector to view additional details for each of the relevant risk mitigation rules.

Each chart sector can drill down to increasing levels of detail, which can be used to isolate potential issues. For the types of risk mitigation details that can be accessed via this chart refer to the following sections.

See “ Storage Optimization Detail Reports ” on page 68.



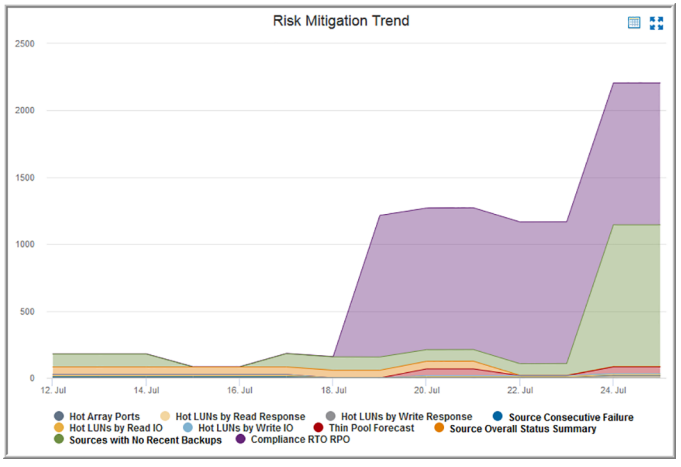
Risk Mitigation Trend

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Trend

View risk mitigation categories over time to identify risks that may impact performance and compliance goals.



Mouse over a data point to view a tooltip that lists the number of objects at risk, as illustrated in the following screen capture. If a risk mitigation rule value is blank, no

data was collected for that rule on the selected date. The total represents the number of objects at risk, for all the risk mitigation rules.

Risk Mitigation Detail Reports

Several risk mitigation reports include charts that enable drilldowns to details for each of the rules. This layered approach, with increasing levels of detail, enables different stages of analysis. These reports are driven by Risk Mitigation Rules.

- See "Backup Job Size Variance" on page 47.
- See "Source Consecutive Failure" on page 48.
- See "Source Overall Status Summary" on page 49.
- See "Sources with No Recent Backups" on page 52.
- See "Compliance RTO RPO" on page 50.
- See "Host Multi-Pathing Exposure" on page 52.
- See "Hot Array Ports" on page 53.
- See "Hot LUNs by Read IO" on page 54.
- See "Hot LUNs by Read Response" on page 55.
- See "Hot LUNs by Write IO" on page 56.
- See "Hot LUNs by Write Response" on page 56.
- See "NetBackup Disk Pool Forecast" on page 57.
- See "Thin Pool Forecast" on page 58.
- See "VM Datastore Forecast" on page 60.
- See "VM Guest Disk Forecast" on page 61.

Backup Job Size Variance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Backup Job Size Variance** sector.

Use this report to identify aberrations in backup job sizes, period over period, which may be an indicator of backup issues.

Table 4-1 Backup Job Size Variance

Column name	Description
Client	Client of the backup job.
Average Job Size for Prior Period	Average size for backup jobs in the previous period. The period for comparison is defined in the risk mitigation rule.
Prior Time Period	Prior time period of the selected period over period. The period for comparison is defined in the risk mitigation rule.
Average Job Size for Current Period	Average size for backup jobs in the current period. The period for comparison is defined in the risk mitigation rule.
Current Time Period	Current time period of the selected period over period. The period for comparison is defined in the risk mitigation rule.
Job Size Variance %	The difference in the average job size, expressed as a percentage, for the period over period. The period for comparison is defined in the risk mitigation rule.

Source Consecutive Failure

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Source Consecutive Failure**sector.

Use this report to identify consecutive backup failures for the last 14 days.

Table 4-2 Source Consecutive Failure

Column name	Description
Backup Server	Server from which the backup is executed.

Table 4-2 Source Consecutive Failure *(continued)*

Column name	Description
Source Name	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Message	Failure message for the source that has not been backed up.
Today	Today's backup status.
Yesterday	Yesterday's backup status.
Day 3 - 14 Status	Each column lists the backup status to aid in identifying when the consecutive failures occurred.

Source Overall Status Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Source Overall Status** Summary sector.

Use this report to assess the protection status of clients to ensure that compliance objectives are met. The data reflects historical data for the last 14 days.

Table 4-3 Source Overall Status Summary

Column name	Description
Source Name	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
Product	The backup product vendor name.

Table 4-3 Source Overall Status Summary *(continued)*

Column name	Description
Overall Status	Source backup status. Possible values: Failed, Failed and Incomplete, Successful, Successful and Failed, Successful and Failed and Incomplete, Successful and Warnings and Failed, Successful and Warnings and Failed and Incomplete.
# of Jobs	Number of backup jobs for the source.
Risk Level	Risk level, determined by the number of failed and incomplete jobs and also the job type.
# of Failed Jobs	The number of failed jobs for the source during the last 14 days.
# of Warning Jobs	The number of warnings for jobs for the source during the last 14 days.
# of Incomplete Jobs	The number of incomplete jobs for the source during the last 14 days.
# of Successful Jobs	The number of successful jobs for the source during the last 14 days.
First Attempt	Date and time of the first backup attempt.
Last Attempt	Date and time of the last backup attempt.
Job Types	Type of backup, such as scheduled, catalog, or full.
Exit Codes	Exit codes returned by the backup software.
Written	Amount of data written to the backed up copy.

Compliance RTO RPO

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Compliance RTO RPO** sector.

To remain in compliance for Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO), use this report to help determine how long it will take to restore

data if a recovery is required. Historical data for the last 14 days provides the details, such as the size and duration of the backups.

Table 4-4 Compliance RTO RPO

Column name	Description
Source Name	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
Restore Plan	Examples: No Backup, Full Backup.
Last Successful Full	Date and time of the last successful full backup of source data.
Full Size	The amount of data backed up during the last successful full backup.
Full Duration	The time it took to successfully complete the last full backup of source data.
Last Incremental After Full	Date and time of the last incremental backup after the successful full backup.
# of Increments After Full	The number of incremental backups that completed after the successful full backup.
Incr Size After Full	The amount of data backed up during the incremental backup.
Incr Duration After Full	The time it took to successfully complete the last incremental backup after the successful full backup of source data.
Last Incremental	Date and time of the last incremental backup of source data.
# Increments	Number of increments for the source.
Incr Size	The amount of data backed up during the last incremental backup for the source.
Incr Duration	The time it took to successfully complete the incremental backup of source data.
Recovery Date	The date and time of the expected completion of full data recovery, based on historical times and sizes.

Host Multi-Pathing Exposure

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Host Multi-Pathing Exposure**sector. This report identifies hosts that are at risk because they have less than the specified number of paths. It examines LUN mappings of hosts that do not have multiple HBA ports and Array ports configured between a host and a LUN.

Table 4-5 Host Multi-Pathing Exposure

Column name	Description
Host Name	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Array Name	Name of the array. Click to view the Array Capacity and Utilization report.
Logical Unit Name	Logical Unit Name (LUN). Click to view the LUN Details report.
Number of HBA Ports	Number of HBA ports.
List of HBA Ports	List of HBA ports.
Number of Array Ports	Number of array ports.
List of Array Ports	List of array ports.

Sources with No Recent Backups

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Sources No Recent Backups**sector.

Use this report to ensure that sources' data is sufficiently protected. Historical data for the last 14 days provides the details of the last backups.

Table 4-6 Sources with No Recent Backups

Column name	Description
Source Name	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
Product	Name of the backup software used to back up the source.
Days Since Last Backup	The number of days since the last backup of the source.
Primary IP Address	Primary IP address of the source.
Host Type	Indicates the type or role of the host, such as VM Guest.
Last Backup	Date and time of the last backup.
Job Status	The status of the last backup job for the source.
Job Type	The job type, such as Backup or Restore.

Hot Array Ports

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Hot Array Ports**sector.

Use this report to identify heavily used array ports. The high activity on these ports may impact application performance.

Table 4-7 Hot Array Ports

Column name	Description
Storage Array	The storage array identified with high port activity.
Vendor	Name of the storage array vendor.
Port	Name of the hot port.
Port Number	The number that identifies the hot port.
Port's Nickname	An alternative name for the port, in some cases the WWN is used.
WWN	The port's unique identifier, its worldwide name.
# of Hosts	The number of hosts accessing this port.
# of Logical Units	The number of array LUNs associated with this port.
Peak Data Transferred	The highest point of activity for this hot port.
Average Data Transferred	The average amount of data transferred during this reporting period.
Peak/Average Ratio	The transfer ratio is defined as peak divided by average.

Hot LUNs by Read IO

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Hot LUNs by Read IO**sector.

Table 4-8 Hot LUNs by Read IO

Column name	Description
Storage Array Name	Name of the storage array.
Logical Unit Name	Name of the LUN experiencing high volume of read activity.
Host Names	Comma-separated list of host names that are accessing the LUN.
Peak/Average Ratio	The ratio of the peak read I/O divided by the average read I/O.
Peak Read IO	The peak read I/O, which can be used to identify LUN overload.
Average Read IO	The average read I/O, which can be used to identify typical LUN activity.

Hot LUNs by Read Response

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Hot LUNs by Read Response**sector.

Table 4-9 Hot LUNs by Read Response

Column name	Description
Storage Array Name	Name of the storage array.
Logical Unit Name	Name of the LUN experiencing high volume of read response activity.
Host Names	Comma-separated list of host names that are accessing the LUN.
Peak/Average Ratio	The ratio of the peak read I/O response divided by the average read I/O response.
Peak Read Response	The peak read I/O response time, which can be used to identify LUN overload.

Table 4-9 Hot LUNs by Read Response *(continued)*

Column name	Description
Average Read Response Time	The average read I/O response time, which can be used to identify typical LUN activity.

Hot LUNs by Write IO

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Hot LUNs by Write IO**sector.

Table 4-10 Hot LUNs by Write IO

Column name	Description
Storage Array Name	Name of the storage array.
Logical Unit Name	Name of the LUN experiencing high volume of write activity.
Host Names	Comma-separated list of hosts that are accessing the LUN.
Peak/Average Ratio	The ratio of the peak write I/O divided by the average write I/O.
Peak Write IO	The peak write I/O, which can be used to identify LUN overload.
Average Write IO	The average write I/O, which can be used to identify typical LUN activity.

Hot LUNs by Write Response

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Hot LUNs by Write Response**sector.

Table 4-11 Hot LUNs by Write Response

Column name	Description
Storage Array Name	Name of the storage array.
Logical Unit Name	Name of the LUN experiencing high volume of activity.
Host Names	Comma-separated list of hosts that are accessing the LUN.
Peak/Average Ratio	The ratio of the peak write response time divided by the average write response time.
Port Name	The WWN of the relevant port.
Average Write Response Time	The average write response time, which can be used to identify typical LUN activity.

NetBackup Disk Pool Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **NetBackup Disk Pool Forecast**sector.

This data is used to specifically identify NetBackup storage that may be at risk

Table 4-12 NetBackup Disk Pool Forecast

Column name	Description
Server Name	Server name.
Disk Volume Name	Disk volume name.

Table 4-12 NetBackup Disk Pool Forecast *(continued)*

Column name	Description
Date Range: First Week	Date range of the first week's calculations.
Max Used	First week's maximum used capacity for the NetBackup Disk Pool within the selected number of statistics' days.
Date Range:Last Week	Date range of the last week's calculations.
Max Used	Last week's maximum used capacity for the NetBackup Disk Pool within the selected number of statistics' days.
Current Total Capacity	Current total capacity for the NetBackup Disk Pool.
Current Available Capacity	Current available capacity for the NetBackup Disk Pool.
Threshold %	Defined capacity threshold percentage. This applies to current or forecasted capacity.
Day Left (Linear)	Number of days remaining until the NetBackup Disk Pool reaches the capacity threshold percentage by using Linear Analysis.
Reach Threshold Date (Linear)	Week over week Linear Analysis of when the NetBackup Disk Pool will reach the specified total capacity threshold percentage.
Days Left (POP)	Number of days remaining until the NetBackup Disk Pool reaches the threshold by using POP Analysis.
Reach Threshold Date (POP)	Week over week POP Analysis of when the NetBackup Disk Pool will reach the specified total capacity threshold percentage.

Thin Pool Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **Thin Pool Forecast**sector.

This report uses multi-metric pool capacity for these vendors and products: EMC VMAX Thin Pools, EMC CLARiiON VNX Thin Pools, and HDS Journal Pools. This data is used to forecast storage that may be at risk

Table 4-13 Thin Pool Forecast

Column name	Description
Array Name	Name of the storage array.
Storage Pool Name	Name of the storage pool.
Date Range: First Week	Date range of the first week’s calculations.
Max Used	First week's maximum used capacity for the Thin Pool within the selected number of statistics' days.
Date Range:Last Week	Date range of the last week's calculations.
Max Used	Last week's maximum used capacity for the Thin Pool within the selected number of statistics' days.
Current Total Capacity	Current total capacity for the Thin Pool.
Current Available Capacity	Current available capacity for the Thin Pool.
Threshold %	Defined capacity threshold percentage. This applies to current or forecasted capacity.
Day Left (Linear)	Number of days remaining until the Thin Pool reaches the capacity threshold percentage by using Linear Analysis.
Reach Threshold Date (Linear)	Week over week Linear Analysis of when the Thin Pool will reach the specified total capacity threshold percentage.
Days Left (POP)	Number of days remaining until the Thin Pool reaches the threshold by using POP Analysis.
Reach Threshold Date (POP)	Week over week POP Analysis of when the Thin Pool will reach the specified total capacity threshold percentage.

VM Datastore Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **VM Datastore Forecast**sector.

This report examines VMWare datastore statistics for the number of weeks in the defined period to forecast the date when storage will run out within a three year period.

Table 4-14 VM Datastore Forecast

Column name	Description
Datastore	Datastore name.
Date Range: First Week	Date range of the first week’s calculations.
Max Used	First week's maximum used capacity for the datastore within the selected number of statistics' days.
Date Range: Last Week	Date range of the last week's calculations.
Max Used	Last week's maximum used capacity for the datastore within the selected number of statistics' days.
Current Total Capacity	Current total capacity for the datastore.
Current Available Capacity	Current available capacity for the datastore.
Threshold %	Defined capacity threshold percentage. This applies to current or forecasted capacity.
Days Left (Linear)	Number of days remaining until the datastore reaches the capacity threshold percentage by using Linear Analysis.
Reach Threshold Date (Linear)	Week over week Linear Analysis of when the datastore will reach the specified total capacity threshold percentage.

Table 4-14 VM Datastore Forecast *(continued)*

Column name	Description
Days Left (POP)	Number of days remaining until the datastore reaches the threshold by using POP Analysis.
Reach Threshold Date (POP)	Week over week POP Analysis of when the datastore will reach the specified total capacity threshold percentage.

VM Guest Disk Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Risk Mitigation > Risk Mitigation Analysis

Click the **VM Guest Disk Forecast**sector.

This report examines VMWare guest disk statistics for the number of weeks in the defined period to forecast the date when storage will run out within a three year period.

Table 4-15 VM Guest Disk Forecast

Column name	Description
VM Guest Name	VM Guest name.
Disk Path	Path for disk.
Date Range: First Week	Date range of the first week’s calculations.
Max Used	First week's maximum used capacity for the datastore within the selected number of statistics' days.
Date Range: Last Week	Date range of the last week’s calculations.
Max Used	Last week's maximum used capacity for the datastore within the selected number of statistics' days.

Table 4-15 VM Guest Disk Forecast *(continued)*

Column name	Description
Current Total Capacity	Current total capacity for the VM Guest.
Current Available Capacity	Current available capacity for the VM Guest.
Threshold %	Defined capacity threshold percentage. This applies to current or forecasted capacity.
Days Left (Linear)	Number of days remaining until the VM Guest reaches the capacity threshold percentage by using Linear Analysis.
Reach Threshold Date (Linear)	Week over week Linear Analysis of when the VM Guest will reach the specified total capacity threshold percentage.
Days Left (POP)	Number of days remaining until the VM Guest reaches the threshold by using POP Analysis.
Reach Threshold Date (POP)	Week over week POP Analysis of when the VM Guest will reach the specified total capacity threshold percentage.

Chapter 5

Storage Optimization Solution Reports

This chapter includes the following topics:

- Introduction
- Storage Optimization Overview
- Storage Optimization Analysis
- Storage Optimization by Tier
- Storage Optimization Trend
- Storage Optimization Detail Reports
- Storage Optimization Dashboard
- Unallocated LUNs by Array
- Unallocated LUNs
- Undiscovered LUNs by Array
- Undiscovered LUNs
- Unused LUNs by Array
- Unused LUNs
- Overprovisioned Hosts Summary
- Overprovisioned Hosts
- Inactive LUNS by Array

- Inactive LUNS
- VM Aged Snapshots by Datastore
- VM Aged Snapshots
- VMs Powered Off by VM Server
- VMs Powered Off
- VMs Low CPU by VM Server
- VMs Low CPU
- VMs Undiscovered Disks by VM Server
- VMs Undiscovered Disks
- VMs Not in VM Inventory by Datastore
- VMs Not in Inventory
- Non-VM Files by Datastore
- Non-VM Files
- Data Domain File Compression by Host
- High Backup Retention Jobs by Sources
- Sources Backed Up by Multiple Servers
- Inactive Large Files by Host
- Inactive Large Files
- File Type Usage by Host
- File Type Usage
- AWS Orphan Snapshots
- AWS Orphan Volumes by Account
- AWS Orphan Volumes

Introduction

Storage Optimization report templates provide actionable analytics, enabling administrators to monitor and manage storage consumption before making unnecessary purchases. Because NetBackup IT Analytics has end-to-end visibility,

from array storage to host applications, it supplies the data that is necessary to make informed decisions.

The reports discussed in this section employ Storage Optimization Rules to assess areas within your enterprise that offer candidates for optimization. These rules include parameters that can be configured to isolate specific conditions relevant for your environment.

Storage Optimization Overview

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

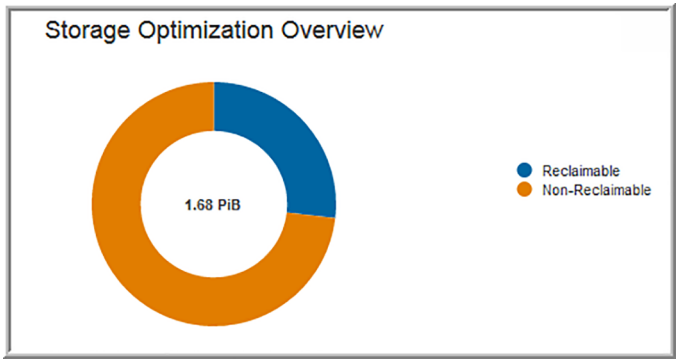
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Overview

View a chart of potentially reclaimable and non-reclaimable storage.

- Click the reclaimable chart sector to view the Storage Optimization Analysis chart which provides a more detailed breakdown of scenarios that may offer reclamation opportunities.
  See " Storage Optimization Analysis " on page 66.
- Click the non-reclaimable sector to view the Array Utilization Summary, which lists capacity utilization details.
- Click the reclaimable sector to view the Storage Optimization Analysis report.

See " Storage Optimization Detail Reports " on page 68.



Storage Optimization Analysis

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

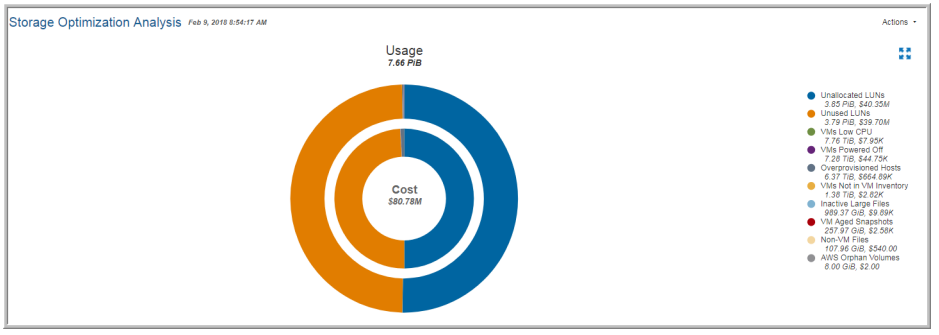
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

View a chart of infrastructure categories and parameters that may offer opportunities for storage optimization. You can also quickly examine the costs associated with optimization opportunities. By knowing costs, you can use the information to help prioritize the problems that have the greatest fiscal impact for your organization. Click a chart sector to view additional details for reclamation possibilities such as Undiscovered LUNs, Over-Provisioned Hosts, or Powered-Off VMs.

Each chart sector can drill down to increasing levels of detail, which can be used to isolate specific arrays, hosts, or VMs that may have reclaimable storage, plus view the associated costs at a granular level. Storage optimization details can be accessed via this chart.

See “ Storage Optimization Detail Reports ” on page 68.



Storage Optimization by Tier

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization by Tier

View a chart of potentially reclaimable and non-reclaimable storage, by tiers. You can also quickly examine the costs associated with optimization opportunities. By knowing costs, you can use the information to help prioritize the problems that have the greatest fiscal impact for your organization. Click a chart sector to view additional details in a similar chart, Storage Optimization Analysis, which renders data for a specific tier (chargeback policy), by infrastructure storage-optimization categories, such as Unallocated LUNs or Over-Provisioned Hosts. Capacity chargeback policies must be configured to support this report. Storage optimization details can be accessed via this chart.

See " Storage Optimization Detail Reports " on page 68.

Storage Optimization Trend

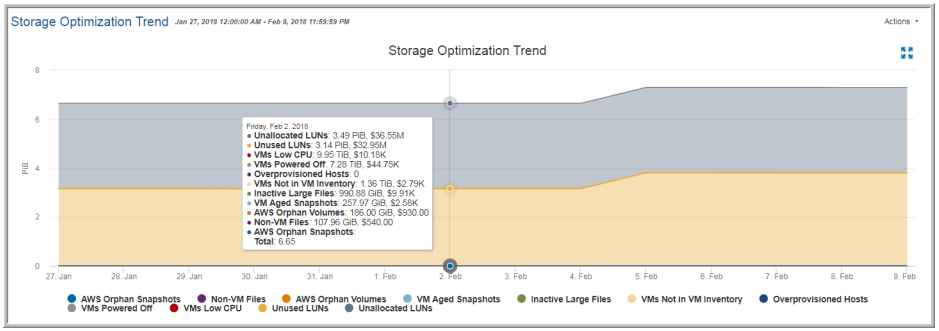
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Trend

View storage optimization categories over time to identify capacity and associated costs in areas that undermine storage optimization goals. By knowing costs, you can use the information to quickly assess the trends that have the greatest fiscal impact on your organization. Trends can be displayed with an emphasis on Cost or Usage, depending on choices made in the scope selector. Storage optimization details can be accessed via this chart.

See " Storage Optimization Detail Reports " on page 68.



Storage Optimization Detail Reports

Several storage optimization reports include charts that enable drilldowns to details for each of the optimization rules. This layered approach, with increasing levels of detail, enables different stages of analysis. These reports are driven by Storage Optimization Rule Categories:

- Cloud
  - See “AWS Orphan Volumes” on page 96.
  - See “AWS Orphan Volumes by Account” on page 95.
- Data Protection
  - See “Data Domain File Compression by Host” on page 89.
  - See “High Backup Retention Jobs by Sources” on page 90.
  - See “Sources Backed Up by Multiple Servers” on page 90.
- File Analytics
  - See “Inactive Large Files by Host” on page 91.
  - See “Inactive Large Files” on page 92.
  - See “File Type Usage by Host” on page 93.
  - See “File Type Usage” on page 93.
- Storage
  - See “Inactive LUNS by Array” on page 78.
  - See “Inactive LUNS” on page 79.
  - See “Overprovisioned Hosts Summary” on page 76.
  - See “Overprovisioned Hosts” on page 77.

- See “Source Overall Status Summary” on page 49.
- See “Unallocated LUNs” on page 71.
- See “Undiscovered LUNs by Array” on page 72.
- See “Undiscovered LUNs” on page 73.
- See “Unused LUNs by Array” on page 74.
- See “Unused LUNs” on page 75.

- Virtualization
  - See “VM Aged Snapshots by Datastore” on page 80.
  - See “VM Aged Snapshots” on page 80.
  - See “VMs Powered Off by VM Server” on page 81.
  - See “VMs Low CPU by VM Server” on page 83.
  - See “VMs Powered Off” on page 82.
  - See “VMs Undiscovered Disks by VM Server” on page 84.
  - See “VMs Undiscovered Disks” on page 85.
  - See “VMs Not in VM Inventory by Datastore” on page 86.
  - See “VMs Not in Inventory” on page 87.
  - See “Non-VM Files by Datastore” on page 87.
  - See “Non-VM Files” on page 88.

Storage Optimization Dashboard

This dashboard can be reviewed regularly to maintain efficient storage use.

Each section on the chart is a report and you can click the charts to get a drilldown view for additional details of particular infrastructure categories, such as:

- Undiscovered LUNs
- Powered-off VMs
- Over-provisioned hosts

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

A trend report shows usage over time by storage category. The reports displayed on the dashboard are briefly described below:

Storage Optimization Overview

This is a chart of potentially reclaimable and non-reclaimable storage. Click the reclaimable chart sector to view the **Storage Optimization Analysis** chart, which provides a more detailed breakdown of scenarios that may offer reclamation opportunities. The non-reclaimable sector provides a drilldown view of the Array Utilization Summary, which lists capacity utilization details.

See " Storage Optimization Overview " on page 65.

Storage Optimization Analysis

Displays a chart of infrastructure categories and use cases that may offer opportunities for storage optimization. Click a chart sector to view additional details for reclamation possibilities such as undiscovered LUNs, over-provisioned hosts, or powered-off VMs.

See " Storage Optimization Analysis " on page 66.

Storage Optimization By Tier

This is a chart of potentially reclaimable and non-reclaimable storage, by tiers. Click a chart sector to view additional details in a similar chart, Storage Optimization Analysis by Tier, which renders data for a specific tier, by infrastructure storage-optimization categories, such as undiscovered LUNs or over-provisioned hosts.

See " Storage Optimization by Tier " on page 66.

Storage Optimization Trend

Provides a view of storage optimization categories over time to identify storage consumption in areas that undermine storage optimization goals. Storage optimization cannot be treated as a once-and-done event. Data growth, combined with dynamic IT configurations, easily undermine the best attempts to minimize waste. Processes that regularly monitor historical data can forestall inefficient storage usage.

See " Storage Optimization Trend " on page 67.

Note: As a best practice, use a monthly schedule to capture point-in-time data that can be compared month-to-month in this storage optimization trending report.

Unallocated LUNs by Array

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Unallocated LUNs** sector.

Table 5-1 Unallocated LUNs by Array

Column name	Description
Storage Array	Storage array name.
Product	The storage family product name, such as Dell Compellent or EMC Symmetrix.
Array Family	Array family.
Unallocated Capacity	Total capacity for the unallocated LUNs in this storage array.
# of Unallocated LUNs	Total number of associated logical units.
Total Cost	Total cost of unallocated logical units.
Row Selection	Select the rows for drilldown.

Unallocated LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Unallocated LUNs** sector to view the Unallocated LUNs by Array report.

Then, click a **Unallocated Capacity** link or check several boxes and **Show Selected**.

Table 5-2 Unallocated LUNs

Column name	Description
Storage Array	Storage array name.
LUN Name	LUN name.
Capacity	Size of the LUN.
RAID Type	RAID type, such as RAID-10, RAID-6, or 2-Way Mirror.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Tier	Name of the chargeback policy (Tier).
Total Cost	Total cost of the Unallocated LUN.
Array Group	Array group name.
Product	The storage family product name, such as Dell Compellent or EMC Symmetrix.
Array Family	Array family.
Is Thin Provisioned	Indicates if thin provisioning is enabled: Yes or No.
Thin Provision Source	Thin-provisioning source, such as EMC Symmetrix Thin Storage Pool.

Undiscovered LUNs by Array

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Undiscovered LUNs** sector.

Table 5-3 Undiscovered LUNs by Array

Column name	Description
Storage Array	Storage array name.
Product	The storage family product name, such as Dell Compellent or EMC Symmetrix.
Array Family	Array family.
Undiscovered Capacity	Total capacity for undiscovered LUNs in this array.
# of Undiscovered LUNs	Total number of associated undiscovered Logical Units.
Total Cost	Total cost of undiscovered Logical Units.
Row Selection	Select the rows for drilldown.

Undiscovered LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Undiscovered LUNs** sector to view the Undiscovered LUNs by Array report.

Then, click a **Undiscovered Capacity** link or check several boxes and **Show Selected**.

Table 5-4 Undiscovered LUNs

Column name	Description
Storage Array	Storage array name.
LUN Name	LUN name.
Capacity	Size of the LUN.
RAID Type	RAID type, such as RAID-10, RAID-6, or 2-Way Mirror.

Table 5-4 Undiscovered LUNs *(continued)*

Column name	Description
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Tier	Name of the chargeback policy (Tier).
Total Cost	Total cost of Undiscovered LUN.
Array Group	Array group name.
Product	The storage family product name, such as Dell Compellent or EMC Symmetrix.
Array Family	Array family.
Is Thin Provisioned	Indicates if thin provisioning is enabled: Yes or No.
Thin Provision Source	Thin-provisioning source, such as EMC Symmetrix Thin Storage Pool.
Host List	List of host ID and name associated with the LUN.
Last I/O Date	Date and time the LUN had I/O operations.

Unused LUNs by Array

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Unused LUNs** sector.

Table 5-5 Unused LUNs by Array

Column name	Description
Storage Array	Storage array name.
Product	The storage family product name, such as Dell Compellent or EMC Symmetrix.
Array Family	Array family.

Table 5-5 Unused LUNs by Array *(continued)*

Column name	Description
Unused Capacity	Total capacity for unused LUNs in this array.
# of Unused LUNs	Total number of associated logical units.
Total Cost	Total cost of unused logical units.
Row Selection	Select the rows for drilldown.

Unused LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Unused LUNs** sector to view the Unused LUNs by Array report.

Then, click a **Unused Capacity** link or check several boxes and **Show Selected**.

Table 5-6 Unused LUNs

Column name	Description
Storage Array	Storage array name.
LUN Name	LUN name.
Capacity	Size of the LUN.
RAID Type	RAID type, such as RAID-10, RAID-6, or 2-Way Mirror.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Tier	Name of the chargeback policy (Tier).
Total Cost	Total cost of the unused LUN.
Array Group	Array group name.
Product	The storage family product name, such as Dell Compellent or EMC Symmetrix.

Table 5-6 Unused LUNs *(continued)*

Column name	Description
Array Family	Array family.
Is Thin Provisioned	Indicates if thin provisioning is enabled: Yes or No.
Thin Provision Source	Thin-provisioning source, such as EMC Symmetrix Thin Storage Pool.
Host List	List of host ID and name associated with the LUN.
Last I/O Date	Date and time the LUN had I/O operations.

Overprovisioned Hosts Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Overprovisioned Hosts** sector.

Table 5-7 Overprovisioned Hosts Summary

Column name	Description
Host	Host name.
File System Available	Total available capacity of all file systems associated with overprovisioning.
File System Capacity	Total capacity of all file systems associated with overprovisioning.
# of File Systems	Total number of associated file systems.
Total Cost	Total cost of available capacity of overprovisioned hosts.
Row Selection	Select the rows for drilldown.

Overprovisioned Hosts

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Overprovisioned Hosts** sector.

Then, click a **File System Capacity** link or check several boxes and **Show Selected**

Table 5-8 Overprovisioned Hosts

Column name	Description
Host	Host name.
File System	File system name.
Storage Array	Storage array name.
LUN Name	Logical unit name.
Tier	Name of the chargeback policy (Tier).
Volume Group	Name of the chargeback policy (Tier).
Available %	Percentage of the file system capacity that is available.
Used Growth %	Growth percentage of file system used.
Avail 90-Day Avg	90-day average of available file system capacity.
Avail 90-Day Max	90-day maximum available file system capacity.
Avail 90-Day Min	90-day minimum available file system capacity.
Avail 90-Day Delta Avg	90-day delta average of available file system capacity.
Avail 90-Day Delta Max	90-day delta maximum available file system capacity.
Avail 90-Day Delta Min	90-day delta minimum available file system capacity.

Table 5-8 Overprovisioned Hosts *(continued)*

Column name	Description
File System Available	Amount of available space in the file system.
File System Capacity	File system capacity.
File System Used	Amount of available space in the file system.
Is Thin Provisioned	Indicates if thin provisioning is enabled: Yes or No.
Mount Point	Mount point of the share.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Storage Type	Specifies the storage server properties, such as Network Attached.
Used 90-Day Avg	Last 90-days average used capacity for the file system.
Total Cost	Total cost of available capacity of Over-Provisioned hosts.

Inactive LUNS by Array

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Inactive LUNs** sector.

Table 5-9 Inactive LUNS by Array

Column name	Description
Storage Array	Storage array name.
Array Family	Storage array family, such as Symmetrix VMAX.
Inactive Capacity	Total capacity for inactive LUNs in this array.
Product	The storage product name, such as Dell Compellent or EMC Symmetrix.

Table 5-9 Inactive LUNS by Array *(continued)*

Column name	Description
# of LUNs	Total number of associated Logical Units.
Total Cost	Total cost of inactive Logical Units.
Row Selection	Select the rows for drilldown.

Inactive LUNS

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Inactive LUNs** sector.

Then, either click the **Inactive Capacity** link or select checkboxes and click **Show Selected**.

Table 5-10 Inactive LUNS

Column name	Description
Storage Array	Storage array name.
LUN Name	Logical unit name.
Capacity	Size of the LUN.
RAID Type	RAID type, such as RAID-10, RAID-6, or 2-Way Mirror.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Tier	Name of the chargeback policy (Tier).
Total Cost	Total cost of the inactive LUN.
Array Group	Array group name.
Product	The storage product name, such as Dell Compellent or EMC Symmetrix.
Status	Status of collection.

Table 5-10 Inactive LUNS *(continued)*

Column name	Description
Array Family	Storage array family, such as Symmetrix VMAX.
Is Thin Provisioned	Indicates if thin provisioning is enabled: Y or N.
Host List	List of host IDs and names associated with the LUN.
Last I/O Date	Date and time the LUN had I/O operations.

VM Aged Snapshots by Datastore

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VM Aged Snapshots** sector.

Table 5-11 VM Aged Snapshots by Datastore

Column name	Description
Datastore	Name of the Datastore.
# of Orphan Snapshots	Number of Orphan Snapshots for this datastore.
File Size	Size of the snapshot file.
Total Cost	Total cost of VMWare aged snapshots.
Row Selection	Select the rows for drilldown.

VM Aged Snapshots

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VM Aged Snapshots** sector.

Then, click a **File Size** link or check several boxes and **Show Selected**.

Table 5-12 VM Aged Snapshots

Column name	Description
VM Name	Name of the VM guest.
VM Server	Name of the virtual system.
Host	Host name.
Datastore	Name of the datastore.
Array List	List of storage arrays associated with datastore.
Snapshot Description	Description of the snapshot.
Snapshot Name	Name of the snapshot.
Creation Date	Date and time the snapshot was created.
File Name	Name of the snapshot file.
File Size	Size of the snapshot file.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of VMWare snapshots.
File Path	Path of the snapshot file.
File Type	Type of the snapshot file.

VMs Powered Off by VM Server

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Powered Off** sector to view the VMs Powered Off by VM Server report.

Table 5-13 VMs Powered Off by VM Server

Column name	Description
VM Server	Name of the virtual system.
Total Size	Capacity of the datastores for this VM.
# of Powered Off VMs	Total number of associated VMs.
Total Cost	Total cost of VM powered off.
Row Selection	Select the rows for drilldown.

VMs Powered Off

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Powered Off** sector.

Then, click a **Total Size** link or check several boxes and **Show Selected**.

Table 5-14 VMs Powered Off

Column name	Description
VM Name	VM name.
VM Server	Name of the virtual system.
Host	Host name.

Table 5-14 VMs Powered Off *(continued)*

Column name	Description
Datastore	Name of the datastore.
Array List	List of storage arrays associated with datastore.
Guest State	Guest state, such as notRunning.
Total Size	Capacity of the datastores for this VM.
Powered Off Date	Date and time when the VM guest power state changed to poweredOff, suspended, or unknown.
State	The power state of the VM.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of VM powered off.

VMs Low CPU by VM Server

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Low CPU** sector.

Table 5-15 VMs Low CPU by VM Server

Column name	Description
VM Server	Name of the virtual system.
Total Size	Capacity of the datastores for this VM.
# of Low CPU VMs	Total number of associated Low CPU VMs.
Total Cost	Total cost of the VM with low CPU.
Row Selection	Select the rows for drilldown.

VMs Low CPU

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Low CPU** sector.

Then, click a **Total Size** link or check several boxes and **Show Selected**.

Table 5-16 VMs Low CPU

Column name	Description
VM Name	VM name.
VM Server	Name of the virtual system.
Host	Host name.
Datastore	Name of the datastore.
Array List	List of storage arrays associated with datastore.
Power	The power state of the VM.
Guest State	Guest state, such as notRunning.
Total Size	Capacity of the datastores for this VM.
CPU Usage	CPU usage.
CPU Total	CPU total.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of the VM with low CPU.

VMs Undiscovered Disks by VM Server

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Undiscovered Disks** sector.

Table 5-17 VMs Undiscovered Disks by VM Server

Column name	Description
VM Server	Virtual system name.
Total Size	Total size of the undiscovered disks for the VM server.
# of Undiscovered Disk VMs	Total number of associated Undiscovered Disk VMs.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of the VM with undiscovered disks.
Row Selection	Select the rows for drilldown.

VMs Undiscovered Disks

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Undiscovered Disks** sector.

Then, click a **Total Size** link or check several boxes and **Show Selected**.

Table 5-18 VMs Undiscovered Disks

Column name	Description
VM Name	Name of the VM Guest.
VM Server	Virtual server name.
Host	Host name.

Table 5-18 VMs Undiscovered Disks *(continued)*

Column name	Description
Datastore	Name of the datastore associated with the VM.
Array List	List of storage arrays associated with datastore.
Power	Power state.
Guest State	Guest state, such as running.
Guest Disk Size	VM Guest disk size.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of the VM with undiscovered disks.
Total Size	Total size for the undiscovered disks.

VMs Not in VM Inventory by Datastore

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Not in VM Inventory** sector.

Table 5-19 VMs Not in VM Inventory by Datastore

Column name	Description
Datastore	Name of the datastore associated with the VM.
File Size	Total size of VM files identified for this datastore.
# of Files	Total number of associated files.
Total Cost	Total cost of the VMs not in the VM Inventory.
Row Selection	Select the rows for drill down.

VMs Not in Inventory

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **VMs Not in VM Inventory** sector.

Then, click a **File Size** link or check several boxes and **Show Selected**.

Table 5-20 VMs Not in Inventory

Column name	Description
Datastore	Name of the datastore associated with the VM.
File Name	Name of the VM file that is not in the inventory.
Top Folder	Top folder for the file location.
File Path	Path of the VM file that is not in the inventory.
Virtual Disk Capacity	Path of the VM file that is not in the inventory.
File Size	Size of the file not in the inventory.
Array List	List of storage arrays associated with datastore.
File Type	Type of the VM file that is not in the inventory.
Last Modified Date	Date and time the VM data was last modified.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of the VM files not in inventory.

Non-VM Files by Datastore

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Non-VM Files** sector.

Table 5-21 Non-VM Files by Datastore

Column name	Description
Datastore	Name of the datastore associated with the VM.
File Size	Total size of non-VM files identified for this datastore.
# of Files	Total number of associated files.
Total Cost	Total cost of non-VM files in inventory.
Row Selection	Select the rows for drilldown.

Non-VM Files

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Non-VM Files** sector.

Then, click a **File Size** link or check several boxes and **Show Selected**.

Table 5-22 Non-VM Files

Column name	Description
Datastore	Name of the datastore associated with the VM.
File Name	Name of the non-VM file in the inventory.
Top Folder	Top folder for the file location.
File Path	Path of the non-VM file in the inventory.

Table 5-22 Non-VM Files *(continued)*

Column name	Description
File Type	Type of the non-VM file in the inventory.
Virtual Disk Capacity	Capacity of the VM disks,
File Size	Size of the non-VM file in the inventory.
Array List	List of storage arrays associated with datastore.
Last Modified Date	Date and time the VM data was last modified.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of the non-VM file in the inventory.

Data Domain File Compression by Host

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Data Domain Compression** sector.

Table 5-23 Data Domain File Compression by Host

Column name	Description
Data Domain Host Name	Name of the Data Domain appliance.
Server Name	The name is derived from the host's DNS name.
Original	Amount of logical capacity associated with backups from this host to the appliance.
# of Clients	Total number of associated clients.
Total Cost	Total cost of Data Domain capacity jobs (original size).
Row Selection	Select the rows for drilldown.

High Backup Retention Jobs by Sources

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **High Backup Retention Jobs**sector.

Table 5-24 High Backup Retention Jobs by Sources

Column name	Description
Source Name	Name of the source (host or job object) that was backed up.
Source Type	Type of the job object.
Total Job Size	Total size of jobs associated with high retention.
# of Jobs	Total size of jobs associated with high retention.
Max Retention Date	Maximum expiration date for a job for this source.
Total Cost	Total cost of High Retention Jobs.
Max Retention Days	Maximum number of days to left for a job to expire for the source.
Row Selection	Select the rows for drilldown.

Sources Backed Up by Multiple Servers

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Sources Backed Up by Multiple Servers** sector.

Table 5-25 Sources Backed Up by Multiple Servers

Column name	Description
Source Name	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
# of Servers	Number of servers the source is being backed up by.
Backup Product Names	Backup product vendor name.
Backup Server Names	Comma separated list of backup server names.
Extra Backup Capacity	Total capacity of extra backups.
Cost Per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of the extra backups.

Inactive Large Files by Host

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Inactive Large Files** sector.

Table 5-26 Inactive Large Files by Host

Column name	Description
Host Name	Name of the host.
Host Type	Host type, such as CIFS.

Table 5-26 Inactive Large Files by Host *(continued)*

Column name	Description
File Size	File size.
# of Inactive Large Files	Total number of associated inactive large files.
Total Cost	Total cost of inactive large files.
Row Selection	Select the rows for drilldown.

Inactive Large Files

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **Inactive Large Files** sector.

Then, click a **File Size** link or check several boxes and **Show Selected**.

Table 5-27 Inactive Large Files

Column name	Description
Host Name	Name of the host.
Last Accessed	File last accessed.
Host Type	Host type, such as CIFS.
File Name	Name of the inactive large file.
File Type	File type, derived from the filename extension.
Owner	Owner of the file.
Is Active	Is the context active? Y or N.
File Size	File size.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.

Table 5-27 Inactive Large Files *(continued)*

Column name	Description
Total Cost	Total cost of inactive large files.

File Type Usage by Host

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

File Analytics data collection is required for this rule.

Click the **File Type Usage**sector.

Table 5-28 File Type Usage by Host

Column name	Description
Host Name	Name of the host.
Host Type	Host type.
Total # of Files	Total number of files found for the File Type Usage rule's criteria.
Domain Name	Domain name.
Total Size	Total number of files found for the File Type Usage rule's criteria.
Total Cost	Total cost of FA file types.
Row Selection	Select the rows for drilldown.

File Type Usage

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **File Type Usage** sector.

Then, click a **Total Size** link or check several boxes and **Show Selected**.

Table 5-29 File Type Usage

Column name	Description
Host Name	Name of the host.
Context Name	Name of the context.
Host Type	Host type.
Context Type	Indicates if this file type is a volume (V) or a share (S).
Is Active	Indicates if the context is active: Y or N.
File Type	File type, such as tif.
# of Files	Number of files for this type.
Total Size	Total amount of storage consumed by this file type.
Total %	Percentage of a file type's size compared to the total file capacity for a host.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of FA file types.

AWS Orphan Snapshots

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **AWS Orphan Snapshots** sector.

Table 5-30 AWS Orphan Snapshots

Column name	Description
Account #	The AWS account ID, a 12-digit number.
Availability Zone Name	Name of the Availability Zone.
Resource ID	Unique identifier for the resource.
Volume Size	Size of the EBS volume when the snapshot was taken.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of AWS orphan volumes.
Region Name	Name of the region.
Volume Name	Size of the EBS volume when the snapshot was taken.
Volume Type	The type of Amazon Elastic Block Store (EBS) volume, such as standard or gp2.
Status	Snapshot status, such as Pending, Completed, or Error.
Volume Creation Date	Date and time the AWS volume was created.
Volume Status	The status of the volume.
Encrypted	Indicates if the volume is encrypted, Y or N.
Marked for Delete	Indicates if the volume has been marked for deletion (soft deleted).

AWS Orphan Volumes by Account

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **AWS Orphan Volumes** sector.

Table 5-31 AWS Orphan Volumes by Account

Column name	Description
Account #	The AWS account ID, a 12-digit number.
# of Orphan Volumes	Total number of volumes associated with EC2 instances that no longer exist.
Volume Size	Size of the Amazon Elastic Block Store (EBS) volume when the snapshot was taken.
Total Cost	Total cost of AWS orphan volumes.
Row Selection	Select the rows for drilldown.

AWS Orphan Volumes

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Solutions > Storage Optimization > Storage Optimization Analysis

Click the **AWS Orphan Volumes** sector.

Then, click a **Volume Size** link or check several boxes and **Show Selected**.

Table 5-32 AWS Orphan Volumes

Column name	Description
Account #	The AWS account ID, a 12-digit number.
Availability Zone Name	Name of the Availability Zone.
Resource ID	Unique identifier for the resource.
Volume Size	Size of the EBS volume when the snapshot was taken.
Cost per GiB	Cost per GiB for storage associated with this storage optimization rule.
Total Cost	Total cost of AWS orphan volumes.
Region Name	Name of the region.

Table 5-32 AWS Orphan Volumes *(continued)*

Column name	Description
Volume Name	Size of the EBS volume when the snapshot was taken.
Volume Type	The type of Amazon Elastic Block Store (EBS) volume, such as standard or gp2.
Status	Snapshot status, such as Pending, Completed, or Error.
Volume Creation Date	Date and time the AWS volume was created.
Volume Status	The status of the volume.
Encrypted	Indicates if the volume is encrypted, Y or N.
Marked for Delete	Indicates if the volume has been marked for deletion (soft deleted).

Chapter 6

System Administration Reports

This chapter includes the following topics:

- Audit Event Details
- Audit Events Summary
- Data Collection Message Summary (by group selection)
- Data Collection Message Summary
- Data Collection Message Detail
- Data Collection Activity Detail
- Data Collection Hourly Activity
- Data Collection Performance Detail
- Data Collection Performance Summary
- Data Collection Performance Summary Trend
- Data Collection Schedule Summary
- Cases Not Represented in the Data Collection Schedule Summary
- Example of a Data Collector Log Request
- Data Collection Status Details
- Data Collector Status Summary
- Database Error Aggregation

- Database Error Summary
- Database Queries By Highest Execution Time
- Database Queries By Total Execution Time
- Database Queries Doing Full Table Scans
- Duplicate Host Analysis
- File Analytics Collection Status
- NetApp Volume Status
- License Summary
- Oracle Job Overview
- Capacity License Usage Details
- File Analytics License Usage Details
- Virtual Server License Usage Details
- Fabric Manager License Usage Details
- Backup License Usage Details
- Replication License Usage Details
- License Overage Detail
- Report Activity Summary
- Report Activity Detail
- Scheduled Reports Summary
- Disable or Delete a Scheduled Report
- Disable a Scheduled Report
- Delete a Scheduled Report
- Modify a Schedule
- System Health Check
- Top Running Reports
- Dashboard Considerations
- Top Running Report Details

- Version History
- WMI Proxy Servers
- Data Collector Version Updates Summary

Audit Event Details

The report displays the details of each event captured within the time period of the report scope in addition to the details present in the Audit Events Summary Report.

You can use this report to view the old and new values of the attributes captured in the audit logs. This report also appears as a drill down view of the Audit Events Summary report when you click the Attributes Modified count.

This report is located here:

System Administration Reports > **Audit Event Details**

Table 6-1 Audit Event Details

Column name	Description
User Name	User name of the user who performed the audited action.
Log Date	Timestamp of event logged.
Event Group	Event group of the logged event.
Event Category	Event category of the logged event.
Event Action	Event action of the logged event.
Audited Object ID	ID of the updated object.
Audited Object Name	Name of the updated object.
Attribute Name	Name of the attribute updated during the event.
New Value	New value of the updated attribute. The new value is blank when the object is deleted.
Old Value	Old value of the attribute before it was updated. The value is blank when the attribute is updated for the first time.

For more information on event groups and categories, see *Event audit logging* in the *NetBackup IT Analytics System Administrator Guide*.

Audit Events Summary

The report displays the audit summary of the events captured within the time period specified in the report scope. You can use this report to view a summary of events audited and the number of attribute changes within a specific time period. The report also provides a detailed drill down view of each attribute change during the event.

This report is located here:

System Administration Reports > **Audit Events Summary**

Table 6-2 Audit Events Summary

Column name	Description
User Name	User name of the user who performed the audited action.
Log Date	Timestamp of event logged.
Event Group	Event group of the logged event.
Event Category	Event category of the logged event.
Event Action	Event action of the logged event.
Audited Object ID	ID of the updated object.
Audited Object Name	Name of the updated object.
Attributes Modified	Number of attributes updated during the event. The link provides a detailed drill down view of each attribute change.

For more information on event groups and categories, see *Event audit logging* in the *NetBackup IT Analytics System Administrator Guide*.

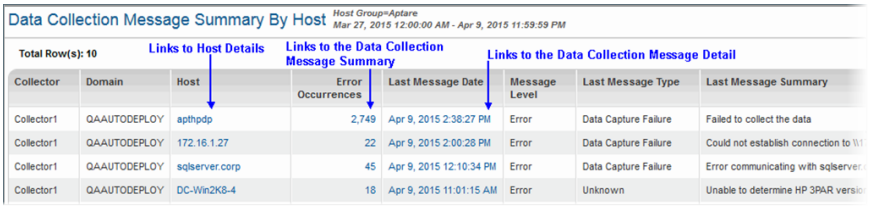
Data Collection Message Summary (by group selection)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collection Message Summary

The Data Collection Message Summary report lists issues that occurred during the data collection process. Use this report to identify the reason for missing backup and storage capacity status.



Data Collection Messages can be filtered (in the scope selector) by the following criteria:

- by Host (see previous example)

Table 6-3 Data Collection Message Summary (by group selection)

Column name	Description
Collector	Data Collector name.
Domain	Portal domain of the host.
Host	Host name. Links to the Host Details.
Error Occurrences	Number of error occurrences. Links to the Data Collection Message Summary report.
Last Message Date	The last time an error message was logged during data collection. Links to the Data Collection Message Detail report.
Message Level	Indicates the severity of the error.
Last Message Type	Indicates the type of error.
Last Message Summary	A brief description of the error. Click on the Last Message Date to see details.

Data Collection Message Summary

This report is accessed via the Error Occurrences link in the Data Collector Message Summary.

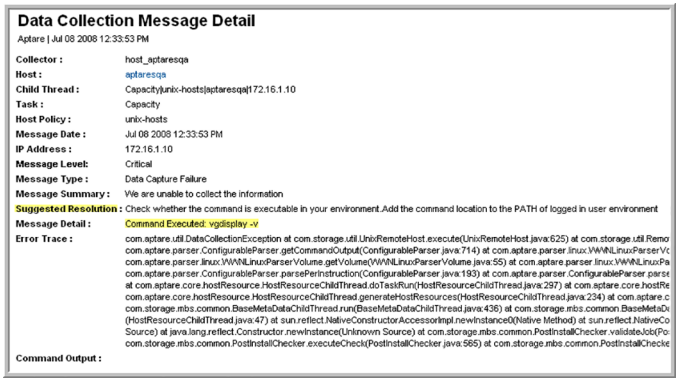
Table 6-4 Data Collection Message Summary

Column name	Description
Collector	Data Collector name.
Host	Host name. Links to the Host Details.
Child Thread	Lists child threads and their configurations encapsulated within a data collector configuration.
Task	Defines the data collection task.
Host Policy	The host policy associated with this collector.
Message Date	The last time an error message was logged during data collection. Links to the Data Collection Message Detail report.
Message Level	Indicates the severity of the error.
Message Type	Indicates the type of error.
Message Summary	A brief description of the error. Click on the Last Message Date to see details.

Data Collection Message Detail

This report is accessed via the Last Message Date link in the Data Collection Message Summary.

See " Data Collection Message Summary (by group selection) " on page 101.



Data Collection Message Detail

Aptare | Jul 08 2008 12:33:53 PM

Collector :	host_aptaresqa
Host :	aptaresqa
Child Thread :	Capacity[unix-hosts]aptaresqa[172.16.1.10
Task :	Capacity
Host Policy :	unix-hosts
Message Date :	Jul 08 2008 12:33:53 PM
IP Address :	172.16.1.10
Message Level:	Critical
Message Type :	Data Capture Failure
Message Summary :	We are unable to collect the information
Suggested Resolution :	Check whether the command is executable in your environment.Add the command location to the PATH of logged in user environment
Message Detail :	Command Executed: vgdisplay -v
Error Trace :	com.aptare.util.DataCollectionException at com.storage.util.UnixRemoteHost.execute(UnixRemoteHost.java:625) at com.storage.util.Remo com.aptare.parser.ConfigurableParser.getCommandOutput(ConfigurableParser.java:714) at com.aptare.parser.linux.VWNLinuxParserVc com.aptare.parser.linux.VWNLinuxParserVolume.getVolume(VWNLinuxParserVolume.java:55) at com.aptare.parser.linux.VWNLinuxPa com.aptare.parser.ConfigurableParser.parsePerInstruction(ConfigurableParser.java:193) at com.aptare.parser.ConfigurableParser.parse at com.aptare.core.hostResource.HostResourceChildThread.doTaskRun(HostResourceChildThread.java:297) at com.aptare.core.hostRe com.aptare.core.hostResource.HostResourceChildThread.generateHostResources(HostResourceChildThread.java:234) at com.aptare.c com.storage.mbs.common.BaseMetaDataChildThread.run(BaseMetaDataChildThread.java:436) at com.storage.mbs.common.BaseMetaDa (HostResourceChildThread.java:47) at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeCo Source) at java.lang.reflect.Constructor.newInstance(Unknown Source) at com.storage.mbs.common.PostInstallChecker.validateJob(Po com.storage.mbs.common.PostInstallChecker.executeCheck(PostInstallChecker.java:565) at com.storage.mbs.common.PostInstallChecke
Command Output :	

Data Collection Activity Detail

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collection Performance Summary

Click a link within the Data Collection Performance Summary

This report provides the elapsed time taken by a particular database component and process, when persisting data during a data collection cycle. This list correlates activity with the logged hour to aid in isolating the process that may have impacted performance.

Primarily used to troubleshoot performance issues, this report can aid in determining if a task is taking an unreasonable amount of time or if load-balancing could improve overall performance. When evaluating the data, this report narrows the results to pinpoint specific activity that impacted performance.

Refer to the *NetBackup IT Analytics User Guide* for information about *Advanced Filtering for Tabular Reports* and *Alerting*.

Note: Data shown in this report may not match the date range selected in the parent report. Only the most recent set of 100 database transactions is retained and reflected in this report.

Table 6-5 Data Collection Activity Detail

Column name	Description
Component Name	The specific data collection link will render the Data Collection Activity Detail report, which lists details specific to the selected component.
Process Group	The Process Group aids in isolating the portion of the code that may be causing performance degradation. This field links to the Data Collection Activity Detail report, which lists details specific to the process group.
Process Name	The Process Name links to the Data Collection Activity Detail report, with an even more granular view of the time period and section of code that may have impacted performance.

Table 6-5 Data Collection Activity Detail *(continued)*

Column name	Description
Total Elapsed Time in Secs	The total time the component required to process the associated number of records.
Number of Records Processed	The number of records processed by the database process.
Average Time Per Record	The average time taken to process each record. The value is displayed in seconds.
Last Updated	Date and time the Portal database was last updated with data from a collection cycle.

Data Collection Hourly Activity

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collection Performance Summary

Click a link within the Data Collection Performance Summary

This report provides the elapsed time taken by a particular database component and process, when persisting data during a data collection cycle. This list correlates activity with the logged hour to aid in isolating the process that may have impacted performance.

Primarily used to troubleshoot performance issues, this report can aid in determining if a task is taking an unreasonable amount of time or if load-balancing could improve overall performance. When evaluating the data, this report narrows the results to pinpoint specific activity that impacted performance.

Refer to the *NetBackup IT Analytics User Guide* for information about *Advanced Filtering for Tabular Reports* and Alerting.

Table 6-6 Data Collection Hourly Activity

Column name	Description
Component Name	The specific data collection link will render the Data Collection Activity Detail report, which lists details specific to the selected component.
Process Group	The Process Group aids in isolating the portion of the code that may be causing performance degradation. This field links to the Data Collection Activity Detail report, which lists details specific to the process group.
Process Name	The Process Name links to the Data Collection Activity Detail report, with an even more granular view of the time period and section of code that may have impacted performance.
Total Elapsed Time in Secs	The total time the component required to process the associated number of records.
Log Hour	The time stamp of the log activity.
Number of Records Processed	The number of records processed by the database process.
Average Time Per Record	The average time taken to process each record. The value is displayed in seconds.

Data Collection Performance Detail

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collection Performance Detail

This report provides the elapsed time taken by a particular database component and process, when persisting data during a data collection cycle. This list correlates activity with the logged hour to aid in isolating the process that may have impacted performance.

Primarily used to troubleshoot performance issues, this report can aid in determining if a task is taking an unreasonable amount of time or if load-balancing could improve

overall performance. When evaluating the data, this report narrows the results to pinpoint specific activity that impacted performance.

Refer to the *NetBackup IT Analytics User Guide* for information about *Advanced Filtering for Tabular Reports* and *Alerting*.

Table 6-7 Data Collection Performance Detail

Column name	Description
Component Name	The specific data collection system link will render the Data Collection Hourly Activity report.
Process Group	The Process Group aids in isolating the portion of the code that may be causing performance degradation. This field links to the Data Collection Hourly Activity report.
Process Name	The Process Name links to the Data Collection Hourly Activity report, with an even more granular view of the time period and section of code that may have impacted performance.
Total Elapsed Time in Secs	The total time the component required to process the associated number of records.
Number of Records Processed	The number of records processed by the database process.
Average Time Per Record	The average time taken to process each record. The value is displayed in seconds.

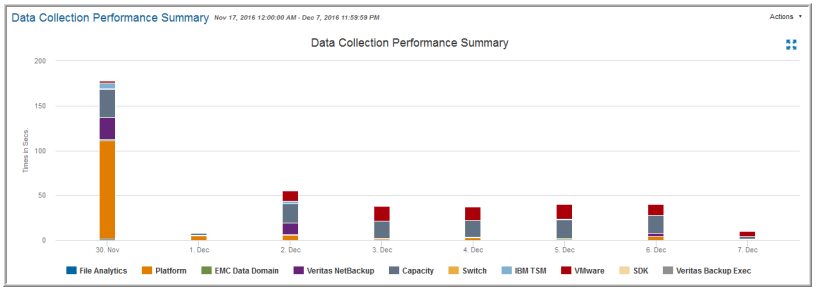
Data Collection Performance Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collection Performance Summary

This report provides a high-level view of days where data collection may have had performance issues. Click a section of a bar to view the Data Collection Performance Detail, which lists specific database components with the elapsed time and number of records processed.



In the stacked bar chart:

- Disable a category by clicking the legend label.
- Click a bar segment go view the Data Collection Performance Detail. See “ Data Collection Performance Detail ” on page 106. report.
- Hover the mouse over a bar segment to view the Elapsed Time tooltip.

Data Collection Performance Summary Trend

Use this report to create a report based alert with filter set on the column that displays percentage linear change. This report is located here:

System Administration Reports > **Data Collection Performance Summary Trend**

This report displays the Data Collection Performance Trend over a period of time. You can view the data Collector components, daily linear change, and total elapsed time in this report. The performance history and linear trend are represented as sparklines in this report.

Table 6-8 Data Collection Performance Summary Trend

Column name	Description
Component	System Component / Subsystem represented by the data/line of report.
Performance History	Trend line representing historical data.
Linear Trend	Calculated Trend line / slope using linear regression.
Daily Change Linear	Daily Change.

Table 6-8 Data Collection Performance Summary Trend *(continued)*

Column name	Description
Total Elapsed Time (Secs)	Total Elapsed time in seconds.
Largest Elapsed Time (Secs)	Largest Elapsed time in seconds.
Avg Elapsed time (Secs)	Average Elapsed time in seconds.
#Days	Number of days of historical data selected in scope.
First Date	First date of historical data selected in scope.
First Value	First value of data selected in scope.
Last Date	Last date of historical data selected in scope.
Last Value	Last value of data selected in scope.
%Change Linear	Percent of Change of data selected in scope (Linear).
Elapsed Time Change (Secs)	Summary of data selected in scope representing if data Changed, Increased, Decreased or Did not Change.
Trend	Icon representing the direction of change.

Data Collection Schedule Summary

Data collection performance may be impacted when multiple policies initiate at the same time. Use this single view of all successful active collection policy schedules to determine scheduling conflicts. Compare the schedule frequency with the start/end times to understand which data collectors may be competing for resources. Note that often long-running collections could be due to database persistence processing. This report can be useful for troubleshooting and implementing database efficiencies.

See "Cases Not Represented in the Data Collection Schedule Summary" on page 111.

Table 6-9 Data Collection Schedule Summary

Column name	Description
Policy Name	The type of system from which data is being collected, such as EMC VPLEX, Data Domain, or NetApp Cluster. The host name or IP address of the system that is the source of the data is appended to the product name, except for Host Resources collection.
Collector Name	Name of the data collector.
Collector Host Name	Name of host where data collector is deployed. Drill down to the Host Details.
Probe Type Name	The type of data collection probe, such as Capacity, Performance, or Inventory.
Schedule	The frequency of the data collection, such as Execute at 12:00 PM Daily. All times are relative to the Data Collector server.
Last Start Time	The date and time the last successful data collection started.
Last End Time	The date and time the last successful data collection ended.
Last Duration	Length of the last successful collection cycle. A color indicator highlights durations that are 50% above the average duration. Note that the duration is calculated from the time the first set of data is persisted in the database until the last data from the collection thread is persisted.
Average Duration	The average duration of all successful executions of this probe, including database persistence processes.
Duration Delta	The % difference between the last successful duration and the average successful duration, which is an indicator of collection and database persistence performance. If the last duration is unknown, the delta will be unknown.
# of Collections	Number of times the probe successfully collected data within the report scope.
Total Duration	Total Duration = (# of Collections) * (Average Duration). This column is useful for sorting purposes, to identify long-running collection and database persistence processes.
Collector Identifier	This metadata collector ID is useful for correlating collection events with the relevant log files. See "Example of a Data Collector Log Request" on page 111. For a VM, instead of a metadata ID, the name of the ESX server is listed in this table column.

Cases Not Represented in the Data Collection Schedule Summary

Due to the nature of certain types of collected data, the Data Collection Schedule Summary report cannot always display schedules for all collected products. Typically, these exception cases cannot be included in the report because there are no explicit start and finish transactions to be captured.

The following table lists the exceptions to the Data Collection Schedule Summary report, where the collected subsystems cannot be represented in the report.

Table 6-10 Cases Not Represented in the Data Collection Schedule Summary

Collected Subsystem	NetBackup IT Analytics Licensed Module	Why the data is not listed in the report
File Analytics - CIFS	File Analytics	No explicit start and finish transactions are collected.
Host Resources	Capacity Manager	May be supported in a future NetBackup IT Analytics version.

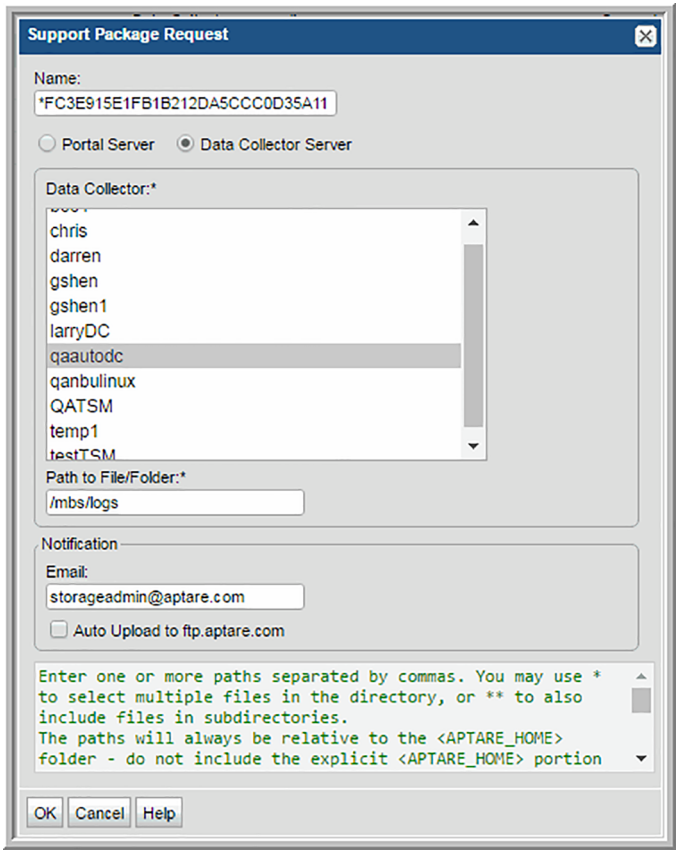
Example of a Data Collector Log Request

When troubleshooting data collection issues, you can use the Support Tools feature to create a Support Package Request to download relevant files. When you add a Support Package Request, the Collector Identifier, listed in the Data Collection Schedule Summary, can be entered for the name to retrieve a specific data collection log.

See " Data Collection Schedule Summary " on page 109.

1. Generate the Data Collection Schedule Summary report.
2. In the Data Collection Schedule Summary, select the Collector Identifier for the data collection thread that needs further investigation.
3. Ctrl-C copy the identifier.
4. Go to the Support Tools: **Admin > Advanced > Support Tools**
5. Click **Add** to create a Support Package Request, supply the following details, and click **OK**.
   - Name: Paste the Collector Identifier into this field and also place a wildcard asterisk (*) before and after the identifier.

- Data Collector Server: Select this radio button.
- Data Collector: Select the data collector associated with the collection thread listed in the Data Collection Schedule Summary.
- Path to File/Folder: Enter the path to the log files. For a Linux data collector server, this path typically is/mbs/logs.
- Email: Enter your email address so that you can be notified when the package becomes available.



Data Collection Status Details

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collection Status Details

The Data Collection Status Details report allows you to monitor the health and progress of data collection. Collection status, available at the data collector and policy level, provides a value for the last time collection was attempted. This report can provide essential details that enable you to diagnose collection issues. You can schedule/email/export this report plus set alerts for failures.

Table 6-11 Data Collection Status Details

Column name	Description
Probe Name	The name of a specific probe on the associated Data Collector policy.
Device	IP address or name of the host/server from which the data was collected.
Collector Name	The name of the Data Collector.
Policy Name	The type of system from which data is being collected, such as EMC VPLEX, Data Domain, or NetApp Cluster. The host name or IP address of the system that is the source of the data is appended to the product name, except for Host Resources collection.
Run Type	The type of collection run: **Scheduled** or **On-Demand**.
Probe State	Defines the state of the probe: queued, in progress, interrupted or canceled.
Status	The status of data collection: Success, Warning, Failure or No Status Available. Status is only displayed after the collection run has completed. Filter the report by all status, no status, a single status or multiple statuses with the scope selector.
Start Date	The date data collection started for the specific probe/collector combination.
Finish Date	The date data collection finished for the specific probe/collector combination.
Duration	The time span of the collection run.
Schedule	Indicates if the policy schedule is enabled or if it has been manually disabled.
Policy Changed	Indicates if the Data Collector policy has been modified since its creation date.

Table 6-11 Data Collection Status Details *(continued)*

Column name	Description
Details	Click the link to access the Data Collection Detail report to help with troubleshooting issues. From this report, you can view specifics about status issues as well as possible resolutions.

Data Collector Status Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Data Collector Status Summary

This report provides an at-a-glance view of any issues that may be occurring with your data collectors.

Refer to the *NetBackup IT Analytics User Guide* for information about Advanced Filtering for Tabular Reports and Alerting.

Note: The report scope selector includes a Show tasks checkbox. When this is checked, the following columns are displayed. When it is unchecked, only the Name, Product Name, and Status columns are included in the report's table.

Table 6-12 Data Collector Status Summary

Column name	Description
Name	The name of the data collector policy.
Product Name	Defines the product subsystem that is being collected.
Task	Lists the specific data collection task. These are defined in the policy. You can choose to omit this column by un-checking the Show Tasks option in the scope selector.
Status	Indicates the status of the collection task, such as warning, failed, or ok.

Table 6-12 Data Collector Status Summary *(continued)*

Column name	Description
Last Update	The date and time the last collection task was run. You can choose to omit this column by un-checking the Show Tasks option in the scope selector.
Comments	The comments provide details to help you determine how often the data was accessed. You can choose to omit this column by un-checking the Show Tasks option in the scope selector.

Database Error Aggregation

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Database Error Aggregation

This report provides a diagnostic tool for System Administrators to monitor typical operational functions. More often than not, you will be requested by Veritas Support to use this diagnostic tool and email the report to them for further investigation.

Entries in this report do not necessarily signal problems that require corrective action. For example, if a user inadvertently attempts to create a Domain Name that already exists, an entry will be written to the database error log file and also displayed in this report.

Ideally, this report will rarely list errors that require additional actions, but periodic monitoring is recommended.

Use this report to get a high-level view of all the occurrences of each database error for the report's interval. Rather than a row representing each error occurrence, this report lists one row for all the occurrences of a particular error. Errors are sorted by error counts, highest to lowest. This aggregation may be helpful in assessing the severity of the issue.

Note: Given that this report extracts customer-specific database messages from a log file, System Administrators should ensure that only authorized personnel are given access to this report.

Table 6-13 Database Error Aggregation

Column name	Description
Error Number	Unique error number associated with specific database errors, to aid in troubleshooting issues. Drill down to the Database Error Summary to list the date and time of each error occurrence.
Number of Errors	The number of occurrences of this error provides insight into the severity of an issue.
First Error Date	Date and time the first error occurred.
Last Error Date	Date and time the last error occurred in this interval.
Module	The software module that was executing when the error occurred.
Portal Error Message	Lists the message displayed on the Portal, if that is the origin of the error.
Oracle Error Message	Lists the message displayed in Oracle if that is the origin of the error.
Oracle Stack Trace	Lists the SQL stack trace error, which pinpoints the line of code that encountered the error.

Database Error Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Database Error Summary

This report provides a diagnostic tool for System Administrators to monitor typical operational functions. More often than not, you will be requested by Veritas Support to use this diagnostic tool and email the report to them for further investigation.

Entries in this report do not necessarily signal problems that require corrective action. For example, if a user inadvertently attempts to create a Domain Name that already exists, an entry will be written to the database error log file and also displayed in this report.

Ideally, this report will rarely list errors that require additional actions, but periodic monitoring is recommended.

Note: Given that this report extracts customer-specific database messages from a log file, System Administrators should ensure that only authorized personnel are given access to this report.

Table 6-14 Database Error Summary

Column name	Description
Error Number	Unique number associated with specific database errors to aid in troubleshooting issues.
Date	Date and time the error was logged.
Module	Defines area of the code in which the error occurred.
Portal Error Message	Lists the message displayed on the Portal, if that is the origin of the error.
Oracle Error Message	Lists the message displayed in Oracle, if that is the origin of the error.
Oracle Stack Trace	Lists the SQL stack trace error, which pinpoints the line of code that encountered the error.

Database Queries By Highest Execution Time

This report is accessible only to the NetBackup IT Analytics Portal super user.

This report is located at:

System Administration Reports > **Database Queries By Highest Execution Time**

Use this report to diagnose Oracle database performance issues. This report lists the database queries ordered by highest query execution time (in seconds) for a given time frame.

Table 6-15 Database Queries By Highest Execution Time

Column name	Description
DB User	Name of the database user account used to run the SQL query.
Module	Name of the Program executing the SQL Query (e.g. SQLPlus.exe, JDBC Thin Client).
Run Time	Date and time when the SQL query was run.

Table 6-15 Database Queries By Highest Execution Time *(continued)*

Column name	Description
Rows Processed	Number of database rows processed by the query.
Total Time (Seconds)	Time required to run the SQL query on the database in seconds.
SQL Query	The actual SQL query that was run on the NetBackup IT Analytics database.

Database Queries By Total Execution Time

This report is accessible only to the NetBackup IT Analytics Portal super user.

This report is located at:

System Administration Reports > **Database Queries By Highest Execution Time**

Use this report to diagnose Oracle database performance issues. This report lists the database queries ordered by total execution Time (in seconds) for a given time frame.

Table 6-16 Database Queries By Total Execution Time

Column name	Description
DB User	Name of the database user account used to run the SQL query.
Module	Name of the Program executing the SQL Query (e.g. SQLPlus.exe, JDBC Thin Client).
Avg. Execution Time (Seconds)	Average time duration taken by the SQL query to run over the time specified while generating the report.
# of Executions	Number of times the SQL query was run during the specified report duration.
Avg. Executions Per Day	Average count of the SQL Query executions each day.
Total Time (Seconds)	Time required to run the SQL query on the database in seconds.
Earliest Run Time	Date and time of the first query execution found in Oracle library cache.
Last Run Time	Date and time of the last query execution found in Oracle library cache.

Table 6-16 Database Queries By Total Execution Time *(continued)*

Column name	Description
SQL Query	The actual SQL query that was run on the NetBackup IT Analytics database.

Database Queries Doing Full Table Scans

This report is accessible only to the NetBackup IT Analytics Portal super user.

This report is located at:

System Administration Reports > **Database Queries Doing Full Table Scans**

Use this report to diagnose Oracle database performance issues. This report lists the database queries ordered by highest query execution time (in seconds) for a given time frame.

Table 6-17 Database Queries Doing Full Table Scans

Column name	Description
DB User	Name of the database user account used to run the SQL query.
Disk Reads	Disk I/O accounted to the SQL query.
# of Executions	Number of times the SQL query was run during the specified report duration.
SQL ID	The unique identifier Oracle assigns to each unique SQL query.
Last Load Time	Date and time when the SQL query was last parsed in Oracle library cache.
SQL Query	The actual SQL query that was run on the NetBackup IT Analytics database.

Duplicate Host Analysis

This report is available to the Super User or System Administrator by default and can be made available to other admins and users by enabling special access permissions. You can use this report to detect duplicate hosts in your environment. You can use the CSV export of this report as an input to the CLI-based utility to merge the duplicate hosts in your environment.

You can find this report at: **Reports** > **System Administration** > **Duplicate Host Analysis**

Report scope selector

Since the report processes information from already collected jobs while processing, running it for specific host groups can help you to generate the report more quickly. This recommendation is applicable for an environment with large - a few tens of thousands of hosts. The scope selector fields are described in the table below.

Table 6-18 Duplicate Host Analysis Scope Selector

Field name	Description
Host Group	Allows you to select the host groups, backup servers, or hosts for the report scope. Your selection narrows down the scope and helps find duplicates more efficiently by targeting specific host groups or host name.
Find Duplicates Using: ■ **Host Name** ■ **Display Name**	You can choose between **Host Name** (default) and **Display Name**. Both searches are case-sensitive. ■ For **Host Name**, the system compares the internal host names to find duplicates. This is the default criterion in the legacy host merge option. ■ For **Display Name**, the system uses the display or external names of the hosts to find duplicates.
Host Type for the Duplicate Host ■ **Clients Only** ■ **All**	■ **Clients Only** allows you to find duplicates only for hosts that are identified as Clients (hosts backed up by any backup system). ■ **All** detects duplicates for all types of hosts.
Surviving host Selection Criteria ■ **Highest Job Count** ■ **Most Recently Updated**	Allows you to specify the criteria to select the surviving host among the duplicates when performing a host merge. ■ **Highest Job Count**: Selects the host with most associated jobs as the surviving host. This is the default criterion of the legacy host merge option, as a higher job count suggests that the host has more data associated with it. ■ **Most Recently Updated**: Selects the most recently updated host as the surviving host. Use this option when the duplicate hosts found are no longer actively collecting new data, as it helps to retain the most current host.

Table 6-18 Duplicate Host Analysis Scope Selector *(continued)*

Field name	Description
Cascade into sub-groups	The scope selector default is to cascade to all child sub-groups when generating the report. If you prefer to report ONLY on the host group you selected, then uncheck Cascade into sub-groups.
Filter by Common Attributes	Select this checkbox to have the report scope display attributes using "AND" logic. By selecting this feature, the report will display those results with the intersection of the selected criteria. If this checkbox is not selected, the report will display attributes using "OR" logic. For example, if you select attribute values, Campbell, Engineering, Cost Center 1 and select Filter by Common Attributes, the report will display only the results that contain all 3 attribute values. If you do not select Filter by Common Attributes, the report will display all results with attributes Campbell, Engineering, or Cost Center 1.
Apply Attributes to Backup Servers	Select this checkbox to apply the attributes only to the backup servers, instead of hosts.

This report displays one row each for each suspected duplicate pair. If multiple duplicate hosts are detected, the report displays one row for each duplicate pair. For example, if host A has three potential duplicates, the report displays three rows - one for each duplicate. The report columns are described below.

Table 6-19 Duplicate Host Analysis

Column name	Description
Sequence Number	Consecutive number for each row. This is a system-defined number. Do not change this number as it can be utilized while manipulating the exported CSV for Duplicate Host Merge process. Also, it allows an easy return to the original sort order.
Surviving Host	The host which will be retained after merge is successful between the row of the duplicate hosts.
Main Host Internal ID	The ID internal to NetBackup IT Analytics for the surviving host as per report scope selection.

Table 6-19 Duplicate Host Analysis *(continued)*

Column name	Description
Duplicate Host Count	Total number of duplicate hosts found for the Main host as per report scope selection. For example, number 3 indicates that this row plus next two rows have duplicate hosts with Main host as the surviving host.
Duplicate Host Internal ID	Internal ID for the duplicate host as per report scope selection.
Main Host Display Name	Main host's display/external name. This is the name which is displayed on various reports of the NetBackup IT Analytics Portal. You can modify it during collection, unlike internal/host name which is used in host matching and is not updated once a record is created in NetBackup IT Analytics database.
Duplicate Host Display Name	Duplicate host's display/external name.
Main Host Name	Main host's internal name. This is the name with which and entry was created for the first time NetBackup IT Analytics database. This name cannot be updated as it is used in host name matching.
Duplicate Host Name	Duplicate host's internal name.
Is Duplicate Host's Merge supported:	Values could be Yes/No. This indicates whether host merging is possible for this host. The hosts can be merged provided: ■ The duplicate host has backup server type 'Client' AND ■ The duplicate host is only collected from NetBackup, Commvault, and legacy Legato subsystems.
Main Host IP Address	IP address of the main host.
Main Host Jobs Count	Total number of jobs associated with the main host.
Duplicate Host Jobs Count	Total number of jobs associated with the duplicate host.

Table 6-19 Duplicate Host Analysis *(continued)*

Column name	Description
Main Host Last Modified Date	Last modified date of the main host which is updated with collections or with an action on the host from the UI.
Duplicate Host Last Modified Date	Last modified date of the duplicate host.
Main Host Attribute Count	Number of NetBackup IT Analytics-shipped or user-defined attributes associated with the main host.
Duplicate Host Attribute Count	Number of attributes associated with the duplicate host.
Main Host Server Type	Backup server type of the main host.
Duplicate Host Server Type	Backup server type of the duplicate host.
Main Host Matching Identifier	Host Matching Identifier for the main host. Applicable only for environments which support 'Host Matching Identifier' advanced option. It is empty for other environments.

File Analytics Collection Status

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > File Analytics Collection Status

This product-specific data collection status report provides essential details that enable you to diagnose collection issues. File Analytics collection is different from other data collection in that it has multiple approaches to interrogating systems, depending on the location of the shares.

Table 6-20 File Analytics Collection Status

Column name	Description
Host	Host name for which volume/share data has been collected; NetApp collection links to the NetApp Volume Status report: CIFS display the host and the share See “ NetApp Volume Status ” on page 125.
Type	NetApp, CIFS, or Windows
Collector	Name of the Data Collector
Status	■ Timed Out ■ Collecting - File meta data is being collected. ■ Collection Completed - File meta data has been collected. ■ Data Aggregation Started - High level report are being created from the meta data and loaded into the Oracle database. ■ Completed - Aggregation has completed. Updated reports can now be viewed in the Portal. ■ Error - Error occurred during collection or aggregation. ■ Unknown
Current Collection	
Start	Time stamp for the data aggregation start time
Finish	Time stamp for data aggregation end time; blank until successful collection and aggregation complete
Elapsed Time	Time, in minutes, that it took for the last data collection to complete; refresh the report if collection is in progress. A total is listed at the bottom of the table.
File Count	Number of files from which metadata was collected; a total is listed at the bottom of the table. This value initially is estimated based on the size of the collection file and then is updated once aggregation completes.
Files/Min	Number of files per minute; updated during active collection and upon completion, populated with the averages for the total duration of the collection period
% Completed	This progress bar reflects the elapsed time relative to the elapsed time of the last data collection/aggregation (current/previous). When there is no previous collection, it will be 0%. In some cases, for example, prolonged collection periods, the % can exceed 100%. Excessively large % values (greater than 200%) may require investigation.

Table 6-20 File Analytics Collection Status *(continued)*

Column name	Description
Previous Collection	
File Count	Number of files from which metadata was collected during the last collection cycle; blank if no previous collection. A total is listed at the bottom of the table.

NetApp Volume Status

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > File Analytics Collection Status

Click a Host link for NetApp collection.

Table 6-21 NetApp Volume Status

Column name	Description
Volume	Volume name
Status	Partial, Collecting, Completed, Error; an Error status links to the Data Collection Message Summary for additional details
Current Collection	
Start	Time stamp for Data Collection start time
Finish	Time stamp for Data Collection end time; blank if collection is still in progress
Elapsed Time	Time, in minutes, that it took for the last data collection to complete; refresh the report if collection is in progress

Table 6-21 NetApp Volume Status *(continued)*

Column name	Description
File Count	Number of files from which metadata was collected
Previous Collection	
File Count	Number of files from which metadata was collected during the last collection cycle; blank if no previous collection

License Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > License Summary

For Super Users only

Use this report to identify what NetBackup IT Analytics products have been licensed, how much of the licensed capacity has been used, and when the terms of the product license expire. In addition, you can access the usage details to determine how close you may be to licensing limits.

Once the licensing limits have been reached, data for any new devices will not be captured. For example, if a new array is added, data for that array will not be captured and persisted. However, if there is an existing array where the capacity increases, the data for the existing array is captured and persisted.

Drilldowns enable access to details.

- See “ Capacity License Usage Details ” on page 129.
- See “ File Analytics License Usage Details ” on page 130.
- See “ Virtual Server License Usage Details ” on page 130.
- See “ Fabric Manager License Usage Details ” on page 131.
- See “ Backup License Usage Details ” on page 132.
- See “ Replication License Usage Details ” on page 133.
- See “ License Overage Detail ” on page 134.

Table 6-22 License Summary

Column name	Description
Licensed Module	Identifies the modules licensed for your organization.
Licensed Unit	The product-specific units that are licensed: Backup Units, Raw TB of capacity, Fabric ports, Utilized TB of File Analytics data collected, Arrays, or VM servers.
Licensed	Number of units that have been licensed for this product.
Used	Amount of licensed units that have been used; links to the License Usage Details report, relevant to the licensed module.
Used %	Used percentage of product license.
Remain	Unused portion of the product license.
Rejected	Identifies the number of objects not collected because licensed values have been exceeded. See " License Overage Detail " on page 134.
Portal Version	Identifies the version of the portal.
Oracle Version	Identifies the Oracle version.
License Expiration	Expiration date and time zone of the license module; N/A for perpetual licenses.

Oracle Job Overview

Monitor the health and progress of your scheduled Oracle jobs. The report lists the overview of the Oracle jobs configured to perform a defined set of actions at regular intervals. This report also provides a drill-down view for each job. The drill-down view gives a detailed view of each instance the job ran within the time interval considered for the report.

System Administrators must ensure that only authorized personnel have access to this report.

Table 6-23 Oracle Job Overview

Column name	Description
Job Name	The name of the scheduled job.
Job Action	The stored database procedure it calls.

Table 6-23 Oracle Job Overview *(continued)*

Column name	Description
Previous Job Status	Status of Previously executed job. (SUCCEEDED/FAILED)
State	Current status of the job, such as Scheduled, Disabled, Broken, or Running.
Last Start Date	Last date when the job was run.
Last Run Duration	The time taken to complete the last job.
Next Run Date	The next scheduled date for the job to run.
Job Run Count	■ The number of times the job ran within the time interval considered for the report. ■ You can click the count to drill down and view the Oracle Job details report for each job.

Oracle Job Details

This report appears as a drill-down report for Oracle Jobs Overview report. This report includes details of every run of the selected job within the time interval considered for the report.

Following information is available from this drill-down report:

Table 6-24 Oracle Job Details

Column name	Description
Job Name	Name of the Oracle job.
Job Action	The stored database procedure that it calls.
Status	Status of the job run.
Oracle Error Number	Oracle Error Number in case the job encountered an error.
Requested Start Date	Requested start date of the job.
Actual Start Date	Actual start date of the job.
Run Duration	Time duration for which the job ran.
CPU Run Duration	The time spent by CPU to process the job.

Table 6-24 Oracle Job Details *(continued)*

Column name	Description
Log Date	Date and time of the job log.

Capacity License Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the Used column for the Capacity Manager module.

For Super Users only

Table 6-25 Capacity License Usage Details

Column name	Description
Array	Array name; links to the Array Capacity & Utilization report.
Model	Array model.
Vendor	Array vendor.
Raw	Total raw capacity that has been collected.
Virtual	For Hitachi USP-V arrays only: Total capacity including attached open systems capacity (back-end arrays) collected by an alternate collector.
Known Virtual	For Hitachi USP-V arrays only: Known virtual capacity.
Licensed Used	Total licensed capacity.
Last Access	Date and time of last successful data collection.

File Analytics License Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the Used column for the File Analytics module.

For Super Users only

Table 6-26 File Analytics License Usage Details

Column name	Description
Device	IP address or name of the host/server from which the data was collected.
Type	CIFS, NetApp, Windows, or Linux.
Hidden	Indicates if the device is not being counted toward licensing.
Utilized	Utilization associated with the data collected.
Last Access	Date and time of last successful data collection.

Virtual Server License Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the Used column for the Virtualization Manager module.

For Super Users only

Table 6-27 Virtual Server License Usage Details

Column name	Description
Name	Name of the VM server; links to the VM Server Detail report
IP Address	IP address of the virtual server
Active VMs	Number of active VMs
Total VMs	Total number of VMs
Last Access	Date and time of last successful data collection

Fabric Manager License Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the Used column for the Fabric Manager module.

For Super Users only

Table 6-28 Fabric Manager License Usage Details

Column name	Description
SAN Fabric Name	The name of the SAN Fabric. Click the link to view the SAN Fabric Details.
Physical Switch Name	The name of the physical switch.
Switch Name	Switch that has been collected
Model	Switch model name
Vendor	Vendor: Cisco, Brocade, or Unknown
Ports	Port number for the switch
Last Access	Date and time of last successful data collection

Backup License Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the Used column for the Backup Manager module.

For Super Users only

This report lists the details that are used to determine license compliance, thereby assisting users in understanding the data that is represented in backup reports.

Table 6-29 Backup License Usage Details

Column name	Description
Source	Name of the host/client that has been backed up.

Table 6-29 Backup License Usage Details *(continued)*

Column name	Description
Source Internal Name	Internal name of the host/client that has been backed up.
Source ID	Internal numeric ID of the host/client.
Source Type	Type of backup collection, such as Server, Client, HP Data Protector Cell Manager, Unknown.
Vendor	Backup software vendor.
Last Updated	Date and time of the last successful data collection.
Last Job	The last backup job completion date and time. The client is counted for the license count if it was backed up (success or failure) in the last 48 hours.
Media Server	Indicates if this host is or ever has been a media server: Y or N.
Primary Server	Indicates if this host is a Primary backup server: Y or N.
Policy	Backup policy used for this backup. When the policy name is not relevant for a vendor, the following data is displayed for this field. TSM: last access date, NetWorker: last updated, Avamar: group name, HP Data Protector: backup specification name, Commvault Simpana: Y, Backup Exec: Y.

Replication License Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the Used column for the Replication Manager module.

For Super Users only

Table 6-30 Replication License Usage Details

Column name	Description
Primary Array	Name of the primary array/system.
Partner Array	Name of the secondary array/system
Model	NetApp model
State	Connected or Disconnected
Enabled	Indicates if replication is enabled
Last Access	Date and time of last successful data collection

License Overage Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

System Administration Reports > License Summary

Click the value in the **# Objects not Collected** column for the **License Overage Details**. This identifies the number of objects not collected for that Licensed Module because licensed values have been exceeded.

For Super Users only

Table 6-31 License Overage Detail

Column name	Description
Object Name	Names of objects not collected for the licensed module.
Object Type	Types associated with objects not collected for the licensed module.

Table 6-31 License Overage Detail *(continued)*

Column name	Description
First Collected Date	The first date the object was collected.
Last Collected Date	The last date the object was collected.

Report Activity Summary

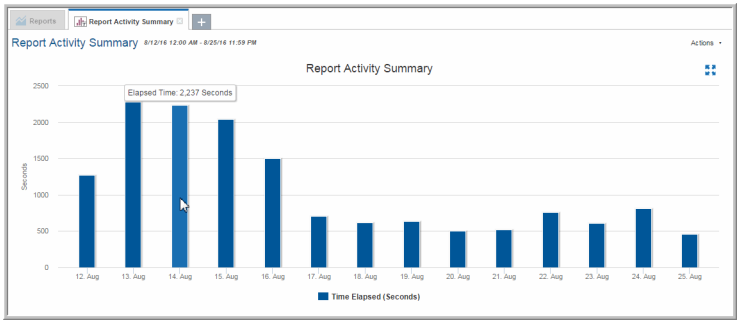
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Report Activity Summary

This report presents a graphical representation of the aggregate of the time for Portal report generation.

- Mouse over a bar to view the total elapsed time.
- Click a bar to view the Report Activity Detail.
  See “ Report Activity Detail ” on page 135.



Report Activity Detail

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Report Activity Detail

This report lists a line item for each report execution with associated user and execution time for the specified time interval. The report can be scheduled, to provide greater visibility about what reports are being run on the Portal.

Table 6-32 Report Activity Detail

Column name	Description
Date	Date and time the report or template was run.
Report Template Name	Name of the report template that was run. In rare circumstances, this field may be blank if the report template was deleted after the report activity was captured in the database.
Report Instance Name	If the report was rendered from a saved instance of the report, a name of the instance will be listed. If the report was generated from the report template itself, this field will be blank. Also, this field may be blank if the report instance was deleted after the report activity was captured in the database.
Dashboard Name	If the report instance was a portlet in a user-defined dashboard, the relevant dashboard will be listed.
User	The User ID of the person who ran the report.
Custom	Y or N indicates if the report was derived from a report template developed by an administrator or user in one of the report template designers, SQL or Dynamic Template Designer. This is not a report that is provided in the Portal software.
Time Elapsed (ms)	This time represents the time required to render the report.

Scheduled Reports Summary

User type defines the access to scheduled reports:

- Administrators can view, reschedule, delete and modify the export format for all users' scheduled reports. This can be done using the **Scheduled Reports Summary**. For example, if a user leaves the company, but her scheduled reports continue to be emailed, you can delete the report's schedule, as described below.

- Non-Administrators can manage their own report schedules using the **My Scheduled Reports** feature.

System Administration Reports > Scheduled Reports Summary

List a summary of reports that are scheduled as alerts, exports, or emails. Use this report to view and manage scheduled reports.

The following options enable you to manage scheduled reports:

- See "Disable or Delete a Scheduled Report" on page 137.
- See "Modify a Schedule" on page 138.

Table 6-33 Scheduled Reports Summary

Column name	Description
Alert Type	Displays how the alert is delivered.
Report	The name of the report instance.
Template	The report template name.
User's Email	The email of the user generating the scheduled report.
Recipient's Email	The email of the user receiving the scheduled report.
Function	The scheduled action: Alert, Export, or Email.
Frequency	The schedule frequency. Click the frequency information to edit and save the new schedule. You can also delete the schedule.
Owner	Displays the owner of the report instance.
Format	The report format that will be delivered, such as html, txt, or image.
Last Executed Date	The date and time the report was last run.

Disable or Delete a Scheduled Report

When a report no longer needs to be regularly scheduled, two options provide a means for removing the report from the schedule:

See "Disable a Scheduled Report" on page 138.

See "Delete a Scheduled Report" on page 138.

These are the recommended ways to remove a report from the schedule.

See "Modify a Schedule" on page 138.

Note: An Administrator can disable or delete another user's scheduled report.

Disable a Scheduled Report

Typically, this option is used to temporarily stop the schedule. The schedule is retained so that it can easily be enabled in the future.

- To disable a scheduled report, click the **Frequency** link of the report and select a **Disabled** run status.

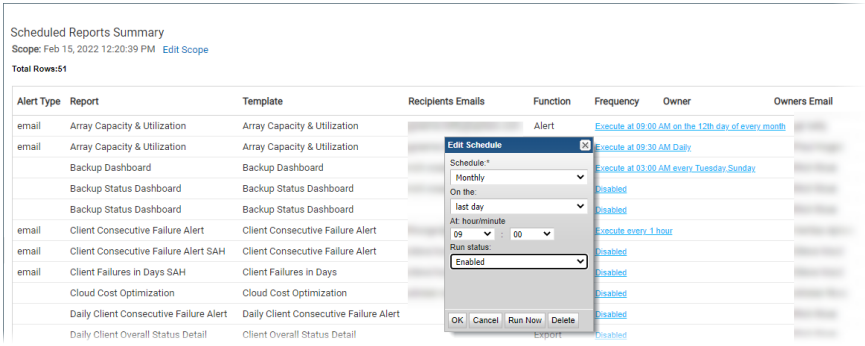
Delete a Scheduled Report

Typically, this option is used to permanently remove a report from the schedule. This option simply removes the report from the schedule, but the report remains intact in the database and can be run manually.

- To delete a scheduled report, click the **Frequency** link of the report and then click **Delete** at the bottom of the **Edit Schedule** window.

Modify a Schedule

1. Click the **Frequency** link to modify the schedule.



2. Modify or delete the report's schedule. You can also run the report on-demand.

System Health Check

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > System Health Check

For Super Users Only

This environmental summary provides a comprehensive overview of operational wellness. Numerous system-level reports are combined and displayed with drilldowns to specifics to identify operational issues. This report is a vital Support tool helping to efficiently troubleshoot issues across your environment.

The System Health Check is available in two modes: Detailed and Summary. Detailed mode displays all user data, whereas Summary mode masks user names and email addresses. Set this preference when you define the scope of the report.

The following system information is included within the report:

- Operating system
- Total Memory
- Portal Memory
- Data Receiver Memory
- Oracle Memory
- Oracle Max Memory
- Application URL
- Data Receiver URI
- Split Architecture
- Processors
- Tomcat Version
- Java Version
- Apache Version
- Portal Version
- Oracle Version
- # Active Users
- Filesystem

The following reports are displayed:

- License Summary. See “ License Summary ” on page 126.

- Version History See “ Version History ” on page 143.
- Advanced Parameter Details.
- Oracle Patch History: Includes the Oracle patch name and the date when it was applied or rolled back.
- System Notifications: Displays notifications with their configuration details.
- Database Queries By Total Execution Time
- Database Queries By Highest Execution Time
- Database Queries doing Full Table Scans
- See “ Data Collection Performance Summary ” on page 107.
- Report Activity Summary
- Backup Volume by NetBackup Policy Type.
- Top 25 Reports by Resource Consumption.
  See “ Top Running Report Details ” on page 142.
- Top 50 Largest Tables
- Inventory Object Summary
- Top 25 Frequently Run Reports.
  See “ Top Running Reports ” on page 141.
- Last 25 Unique Logins
- Oracle Jobs Overview - Last 60 Days
- See “ Database Error Aggregation ” on page 115.
- Portal Error Aggregation
- Alert Policy Summary: Includes alter details, such as alert status, policy name, rule, category, type (warning or critical), creation and modified dates, and last modified by. See “Alert Reports” on page 31.
- Data Collection Message Summary by Child Thread. See “Data Collection Message Summary” on page 102.
- See “ Data Collection Status Details ” on page 112.
- See “ File Analytics Collection Status ” on page 123.
- See “ Data Collection Schedule Summary ” on page 109.

Top Running Reports

Using this report, you can list the reports that have been running for long periods of time or that have run an excessive number of times. A report will be listed as Top Running, based on its overall cost to the system, determined by the average execution time * run count. By default, the most costly reports are listed first. Note that the average execution time includes not only the database query, but also the time associated with building the report on the Portal server.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Top Running Reports

Select the time range for the list in the scope selector.

Click the **Run Count** link to view the Top Running Reports.

See " Top Running Reports " on page 141.

Table 6-34 Top Running Reports

Column name	Description
Template Name	The report template name.
Report Name	The name of the saved report instance.
Is Drilldown	Indicates if the report is a drilldown from a parent report.
Last Run By	The login name of the user who last ran the report.
Run Count	Total number of times the report was run.
Max Execution Time	The maximum amount of time the report has taken to run.
Avg Execution Time	The average amount of time the report has taken to run.
Max DB Time	The maximum amount of time used to access the Portal database.
Avg DB Time	The average amount of time used to access the Portal database.
Last Run Time	The date and time the report was last run during this time period.

Table 6-34 Top Running Reports *(continued)*

Column name	Description
First Run Time	The date and time the report was last run during this time period.

Dashboard Considerations

A dashboard that contains four different reports (portlets) will appear in this Top Running Reports list as four separate records. This summary does not display a row for the high-level dashboard, just the reports contained within the dashboard.

- If the portlet within the dashboard is a saved report, the name of the saved report will be displayed in the Report Name column.
- If the portlet is not a saved report--that is, it was simply generated and dragged into the dashboard--then Dashboard Report will be displayed in the Report Name column.

Top Running Report Details

The Run Count in the Top Running Reports Summary provides access to the list of reports that have consumed higher than average resources. Use this list to determine if any actions need to be taken. For example:

- Adjust the scope of the report to narrow the content.
- Schedule certain top running reports to be run at off hours.
- Verify that automatically schedule reports are scheduled appropriately and are not running more frequently than is required.
- Consider scheduling frequently run reports to run once, with filesystem exports or emails to users.

Table 6-35 Top Running Report Details

Column name	Description
Template Name	The report template name.
Report Name	The name of the saved report instance.
Is Drilldown	Indicates if the report is a drilldown from a parent report.
Dashboard Name	Name of the dashboard in which this report is located.
Schedule Name	The name attached to the schedule for the report.

Table 6-35 Top Running Report Details *(continued)*

Column name	Description
Schedule Frequency	The schedule frequency. Click the frequency information to edit and save the new schedule. You can also delete the schedule.
LDAP ID	The LDAP ID of the user running the report.
Total Execution Time	The total amount of time the report took to render.
DB Execution Time	The total amount of time used for Portal database access.
Run Time	The transaction timestamp for the report.

Version History

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > Version History

The Version Historyreport lists the installation and upgrade events so that you can determine the current version and history of NetBackup IT Analytics software upgrades. This report enables you to verify that the latest patches have been installed.

The Status indicates the success or failure of the installation/upgrade process. The Comments, if available, are derived from the process as well. In many cases, there won’t be any specific comments related to the upgrade.

Table 6-36 Version History

Column name	Description
Start Time	Portal upgrade/installation start date and time.
End Time	Portal upgrade/installation end date and time.
Version	Version number of upgrade.
Status	Upgrade/installation end status.

Table 6-36 Version History *(continued)*

Column name	Description
Comments	Additional information related to the version.

WMI Proxy Servers

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

System Administration Reports > WMI Proxy Servers

For Super Users only

Each Data Collector can communicate with one or more WMI Proxies, and hence the Data Collector captures the server name/IP and the Proxy version from each of the servers running the WMI Proxy.

Note: This report is populated with data only after Windows host collection has occurred, at least once.

Use this report to determine if the WMI Proxy Servers that are used by data collectors are up-to-date.

Table 6-37 WMI Proxy Servers

Column name	Description
Domain Name	The domain name under which the proxy server resides.
WMI Proxy Name	WMI Proxy Name.
Host	WMI Proxy Server IP address or host name.
Version	WMI Proxy version.
Build Number	Build number of the WMI Proxy version.
Collector Name	The data collector that is using the WMI Proxy.

Data Collector Version Updates Summary

This report helps you to identify the current versions of Data Collector and Upgrade Manager. Information such as Version details, Date/Time when the last version update was requested, Date/Time when the last version update happened, if Auto-Update is enabled are also displayed in this report.

Navigate to *Reports* > *System Administration Reports* > *Data Collector Version Updates Summary report.*

Table 6-38 Report columns

Field	Description
Domain	NetBackup IT Analytics domain name
Collector	Data Collector name
Enabled	Displays if the Data Collector is enabled. Y/N
Collector online	Displays if the Data Collector is online. Y/N
Collector version	Displays the current version of Data Collector.
Auto update	Displays if the Data Collector is configured for auto updates. Y/N
Upgraded on	Displays the date/time of the last upgrade of the Data Collector.
Upgrade manager version	Displays the current version of Update Manager.
Auto Upgrade	Displays if the Upgrade Manager is configured for auto updates. Y/N
Update Requested	Displays the date/time of the last Upgrade Manager update request.
Upgraded on	Displays the data/time of the last upgrade of the Upgrade Manager.
Hosts	Data Collector Host name.

Chapter 7

Capacity Manager Reports

This chapter includes the following topics:

- Aggregate Mirror Summary
- FlexClone Summary
- HPE XP Efficiency Summary
- HPE XP DP Pool Summary

Aggregate Mirror Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Replication Manager > Aggregate Mirror Reports > Aggregate Mirror Summary

This report lists each NetApp Aggregate that was mirrored and its status. Use this report to get an overview of the state of all aggregate mirrors.

Table 7-1 Aggregate Mirror Summary

Column name	Description
Aggregate Name	The NetApp Aggregate that was mirrored; links to the details. See " NetApp Aggregate Detail " on page 368.
Storage Array	Links to the Array Detail report.

Table 7-1 Aggregate Mirror Summary *(continued)*

Column name	Description
Mirror Status	The Aggregate Mirror status: invalid, uninitialized, needs CP count check, CP count check in progress, unmirrored, mirrored, mirror degraded, mirror resynchronizing, failed, or limbo.
Detected Date	Date and timestamp of the mirror detection.
Split Detected Date	Date and timestamp a split was detected.

FlexClone Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Replication Manager > FlexClone Reports > FlexClone Summary

Table 7-2 FlexClone Summary

Column name	Description
Parent Volume	The parent NetApp volume links to the details. See " NetApp Volume Details " on page 398.
Child Volume	The child volume of the FlexClone transaction links to its details. See " NetApp Volume Details " on page 398.
Storage Array	The name of the storage array links to the NetApp Storage System Detail.
Parent Base Snapshot	The name of the initial, base snapshot.
Clone Creation Date	The date and timestamp when the clone was created.
Clone Split Date	The date and timestamp when the clone was split.
Split Estimate	The amount of storage estimated to be consumed by the split.
Aggr Available	The amount of storage on the aggregate that is not in use.

HPE XP Efficiency Summary

Explore your datacenter using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates, and dashboards across the portal. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

Navigate to *Reports* tab > *Capacity Manager* > *Array Capacity & Utilization* > *HPE XP Efficiency Summary*.

This report displays the summarized view of the HPE XP operational efficiency. You can define the scope of the array and configure the report parameters using filters.

Table 7-3 HPE XP Efficiency Summary displayed columns

Column name	Description
Storage Array Id	Storage Array Id of APS_STORAGE_ARRAY table.
Storage Array Name	Name of the storage system.
Total Efficiency Ratio	The total efficiency of the entire storage system.
Data Reduction Ratio	The efficiency of capacity saving performed by using the capacity saving function (dedupe and compression) or accelerated compression.
Snapshot Efficiency Ratio	The efficiency of capacity saving performed by using snapshots to back up the data.
Provisioning Efficiency Percentage	The percentage of saving efficiency for consumption of capacity achieved by using Dynamic Provisioning to virtualize capacity.
Calculation Start Time	The date and time when the calculation of the total efficiency began (UTC), format(YYYY-MM-DD hh:mm:ss AM/PM).
Calculation End Time	The date and time when the calculation of the total efficiency ended (UTC), format(YYYY-MM-DD hh:mm:ss AM/PM).
Software Saving Ratio	Total efficiency of capacity saving performed by using the capacity saving function.
Software Compression Ratio	Efficiency of capacity saving performed by using compression.
Software Deduplication Ratio	Efficiency of capacity saving performed by using deduplication.
Software Pattern Matching Ratio	Efficiency of capacity saving performed by reclaiming the specified data pattern.

Table 7-3 HPE XP Efficiency Summary displayed columns *(continued)*

Column name	Description
FMD Pattern Matching Ratio	The efficiency of capacity saving performed by using accelerated compression are displayed.
FMD Saving Ratio	
FMD Compression Ratio	
Creation Date	Date of creation.
Last Updated	Last updated date.

HPE XP DP Pool Summary

Explore your datacenter using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates, and dashboards across the portal. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

Navigate to*Reports* tab > *Capacity Manager* > *Thin Provisioning Capacity & Utilization* > *HPE XP DP Pool Summary*.

This report displays the virtual volumes for data storage with actual capacity available. Use this report to ensure that the pool has sufficient free capacity, with percentage and ratio values.

Table 7-4 HPE XP DP Pool Summary report columns

Column name	Description
Model	Storage array type.
Serial Number	Serial number of the storage array.
Pool Type	Type of the pool e.g. Dynamic Provisioning, Dynamic Provisioning (Tiering),Data Direct Mapping, Thin Image
Collection Time	Time (in GMT) when the data was collected from the storage system
Logical Capacity	Capacity of the pool
Used Logical Capacity	Amount of pool capacity that is actually being used (KB).
Used Logical Capacity in Percentage	Usage rate of the pool capacity

Table 7-4 HPE XP DP Pool Summary report columns *(continued)*

Column name	Description
Available Physical Volume Capacity	Total amount of the physical capacity of the parity groups that make up the pool (KB).
Available Physical FMC Pool Volume Capacity	Total amount of used capacity in the physical capacity of the parity groups that make up the pool (KB).
Subscribed Capacity	Virtual capacity of all Dynamic Provisioning V- VOLs mapped to the pool (KB)
Subscription Percentage	Virtual capacity of all Dynamic Provisioning V- VOLs mapped to the pool (KB)
Threshold Rate	System threshold or depletion threshold set for the percentage of used pool space
Auto Expansion	Indicates whether to automatically add pool volumes according to the compression ratio of a parity group for which accelerated compression is enabled
Data Volumes Used Capacity	Capacity used by a virtual volume (KB)
Data Reduction Capacity	The capacity saved by capacity saving or accelerate compression of the parity group (KB)
Data Reduction Ratio	Saving effect achieved by capacity saving or accelerate compression of the parity group. This is the value of N when the ratio of capacity before and after the reduction is N:1
Deduplication Capacity	Capacity saved by deduplication of capacity saving (KB)
Deduplication Ratio	Ratio of disk space that is saved through deduplication when using capacity saving. This is the value of N when the ratio of capacity before and after the reduction is N:1
Compression Capacity	Capacity compressed by capacity saving (KB)
Compression Ratio	Ratio of disk space that is saved through compression when using capacity saving. This is the value of N when the ratio of capacity before and after the reduction is N:1
Deduplication and Compression Capacity	Capacity reduced when deduplication and compression are used in a pool
Deduplication and Compression Ratio	Saving effect achieved by capacity saving. This is the value of N when the ratio of capacity before and after the reduction is N:1
Pattern Matching	Capacity saved by reclaiming the specified data pattern using capacity saving (KB)
System Data	The amount of consumed system data (KB)

Table 7-4 HPE XP DP Pool Summary report columns *(continued)*

Column name	Description
Preprocessed Data	Capacity of data (before reduction) to be reduced by capacity saving (KB)
FMC Logical Total	The total logical capacity of the pool volume for FMC or the pool volume that supports accelerated compression used in the pool (KB)
FMC Logical Used	The logically used capacity of the pool volume for FMC or the pool volume that supports accelerated compression used in the pool (KB)
FMC Physical Total	The total physical capacity of the pool volume for FMC or the pool volume that supports accelerated compression used in the pool (KB)
FMC Physical Used	The actually used capacity of the parity group for FMC or the parity group that supports accelerated compression (KB)
FMC Compression Capacity	The compressed capacity with the accelerated compression function (KB)
FMC Compression Ratio	Saving effect achieved by accelerated compression. This is the value of N when the ratio of capacity before and after the reduction is N:1
FMC Pattern Matching	The reclaimed capacity with the accelerated compression function (KB)
Creation Date	Date of creation
Last Updated	Last updated date

Chapter 8

Application Capacity Reports

This chapter includes the following topics:

- Application Storage Dashboard
- Application Sub-Level of the Application Storage Dashboard
- Applications at Risk
- Over-Provisioned Applications

Application Storage Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Application Capacity & Utilization > Application Storage Dashboard

This report essentially interrogates the application to identify what storage is free. In addition, this report provides data that can be used to resolve performance issues. The table has three main sections of information: Application Storage, File System Storage, and Array Storage.

Table 8-1 Application Storage Dashboard

Column name	Description
	Click the **arrow** to drill down to a hierarchical view. Each subsequent drilldown takes you to the next level of the application, such as the table space and filesystem level. See " Application Sub-Level of the Application Storage Dashboard " on page 154.
Application	The name of the level. For the application level, the name will be specific to the application, such as Exchange or Oracle. If it is the tablespace level, the name of the respective tablespace is displayed.
Host	The host(s) associated with the application.
Application Storage Total	Total storage allocated to the application. App. Used + App. Free = App. Total The total may not be precise to the decimal, as 1 Kbyte = 1024.
Application Storage Free	Amount of storage available for the application. In some cases, storage does not need to be pre-allocated to the application, so this field will display as N/A.
Application Storage Used	Amount of storage consumed by the application. This is the amount of consumed storage as reported by the application. For example, in Oracle within a data file it would display how much space is really being consumed by Oracle.
Filesystem Storage Total	This total indicates the amount of storage used by the filesystem associated with the application.
Filesystem Storage Free	Amount of unused and available filesystem storage. Refer to this value to determine if there is space to expand a files system.
Filesystem Storage Used	Amount of used filesystem storage.
Array Storage Total	This total indicates the amount of storage used by the array on which this application is located.
Array Storage Free	Amount of available array storage.
Array Storage Used	Amount of used array storage.
Arrays	Storage array; if this field is empty, it means it is a local disk.

Table 8-1 Application Storage Dashboard *(continued)*

Column name	Description
LUNs	A concatenated list of LUNs from which the space is being taken by this application.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Application Sub-Level of the Application Storage Dashboard

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Application Capacity & Utilization > Application Storage Dashboard

Click the arrow to view the Application Level detail.

The data displayed in this window is simply the next level within the application--that is, instance-level reporting, as described in the Application Storage Dashboard.

See “ Application Storage Dashboard ” on page 152.

At the application sub-level, a host link launches the Host Utilization Detail report.

See “ Host Utilization Detail ” on page 481.

For example, an Oracle Database Administrator may be interested in managing storage at the tablespace level. Using the Application Storage Dashboard, the administrator can expand the Oracle application to view table space and also filesystem usage. In essence, this view provides the physical implementation of the filesystem, including the LUNs allocated to it.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Applications at Risk

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Application Capacity & Utilization > Applications At Risk

Using this report, you can identify applications that are at risk of running out of storage space. Identify which applications are expected to run out of storage and when.

Note: The **Applications at Risk** report supports the following applications: SQL Server, MS Exchange, and Oracle Database.

Table 8-2 Applications at Risk

Column name	Description
Name	The application at risk.
Risk Date	The first date of the period for which the usage is projected.The risk date is based on the historical growth of the application's usage.
Usage	The projected usage based on the average usage history (or historical average usage for periods in history).
Projected High-Water Mark	The projected maximum usage based on historical average usage.
Capacity	The database capacity.
Current Usage	Database storage already in use.

Table 8-2 Applications at Risk *(continued)*

Column name	Description
Threshold	If a threshold has been set for a particular application--that is, a database threshold--then that threshold is used to determine the risk. If there is no threshold associated with the application's database, then other thresholds are checked, in the following sequence: Host for the associated LUN/Filesystem/Database, then Host group to which the host belongs, and then the Top-level, global host group. If no threshold settings are found at the above levels, the system defaults are used: 90% for Critical, 70% for Warning, and 30% for Low. Thresholds can be set by an administrator using **Admin > Reports > Threshold Policies**.

Over-Provisioned Applications

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Application Capacity & Utilization > Over Provisioned Applications

Using this report, you can identify applications that are at risk of having over-provisioned storage. Often, storage is allocated to a project, but never used. This report helps you to identify excess, unused storage so that it can be reclaimed for another purpose.

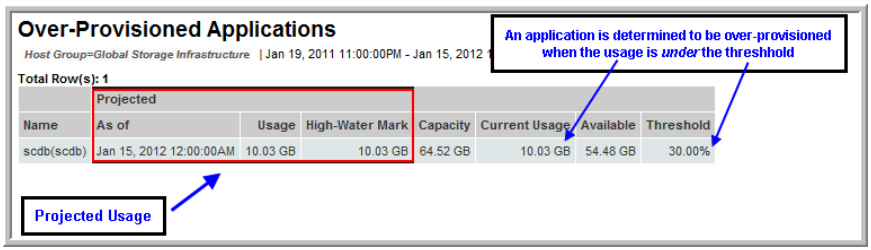


Table 8-3 Over-Provisioned Applications

Column name	Description
Name	The application that has been given storage that remains unused.
Projected	
As of	The first date of the period for which the usage is projected.
Usage	The projected usage based on the average usage history (or historical average usage for periods in history).
High-Water Mark	The projected usage based on the maximum usage history.
Current	
Capacity	The database capacity.
Current Usage	Database storage already in use.
Available	Amount of storage still available.
Threshold	An application is determined to be over-provisioned when the usage is under the threshold.

Chapter 9

Array Capacity and Utilization Reports

This chapter includes the following topics:

- Array Capacity and Utilization
- Array Capacity & Utilization (Generic Data)
- Array Capacity & Utilization - (Dell Compellent View)
- Array Capacity & Utilization - (EMC Data Domain Storage View)
- Array Capacity & Utilization - (EMC VNX Celerra View)
- Array Capacity & Utilization - (EMC VNX CLARiiON View)
- Array Capacity & Utilization - (EMC VPLEX)
- Array Capacity & Utilization - (EMC Isilon View)
- Array Capacity & Utilization - (EMC Symmetrix View)
- Array Capacity & Utilization (Hitachi Content Platform)
- Array Capacity & Utilization (Hitachi Arrays & HDP View)
- Array Capacity & Utilization (Hitachi NAS Arrays)
- Array Capacity & Utilization - (Hitachi Vantara All-Flash and Hybrid Flash Storage)
- HP 3PAR Overview
- Array Capacity & Utilization (HP 3PAR View)
- Array Summary

- Physical Summary
- Chunklets
- Logical (LD) Summary
- Common Provisioning Groups (CPG)
- VV Summary
- Array Capacity & Utilization (IBM Enterprise Array View)
- Array Capacity & Utilization (IBM Modular Array View)
- Array Capacity & Utilization (IBM SVC View)
- Array Capacity and Utilization (IBM XIV View)
- Array Capacity and Utilization (NetApp View)
- Array Capacity and Utilization (NetApp Cluster)
- Array Capacity & Utilization - (Hitachi Vantara All-Flash and Hybrid Flash Storage)
- NetApp Storage System Detail
- Array iSCSI Port Utilization
- NetApp Cluster-Mode Array iSCSI Port Utilization
- Array Capacity and Utilization (OpenStack Swift)
- Disk Summary
- RA Summary
- Array Executive Summary
- List Array Groups
- HP EVA Physical Disks
- Array FC Port Utilization
- Data Domain Filesystem Compression
- Data Domain System Summary - Storage
- Data Domain System Details
- Data Domain MTree Summary

- Data Domain Replication History
- Data Domain Snapshot History
- EMC Isilon Disks
- EMC Isilon HDFS Racks
- EMC Isilon HDFS Roots
- EMC Isilon License Detail
- EMC Isilon NFS Exports
- EMC Isilon Nodes
- EMC Isilon Quota Summary
- EMC Isilon SMB Share Summary
- EMC Isilon Snapshot Summary
- EMC Isilon Zones
- EMC VNX (Celerra) Data Mover Summary
- EMC VNX (Celerra) Virtual Data Mover Summary
- EMC VNX (Celerra) File System Summary
- EMC VNX (Celerra) SnapSure Summary
- EMC VNX (Celerra) Storage Pools Summary
- EMC VNX (Celerra) Storage Pools Detail
- EMC VNX (Celerra) Tree Summary
- EMC VNX (Celerra) Volume Summary
- EMC VNX (Celerra) Volume Detail
- EMC VNX (Celerra) Disk Summary
- EMC VNX (Celerra) Storage System Detail
- EMC VPLEX Backend Array Summary
- EMC VPLEX Device Summary
- EMC VPLEX Extent Summary
- EMC VPLEX Port Summary

- EMC VPLEX Storage View Summary
- EMC VPLEX Storage Volume Summary
- Object Store Namespace Summary
- Node Summary
- Object Store Tenant Summary
- Hitachi NAS CIFS Share Summary
- Hitachi NAS Disk Summary
- Hitachi NAS EVS Summary
- Hitachi NAS Ethernet Port Summary
- Hitachi NAS File System Summary
- Hitachi NAS NFS Export Summary
- Hitachi NAS Snapshot Summary
- Hitachi NAS Storage Pool Summary
- Hitachi NAS Storage Server Summary
- Hitachi NAS System Drive Summary
- Hitachi NAS Tier Summary
- Hitachi NAS Virtual Volume Summary
- Overview of IBM Array Support
- IBM Array Site Summary
- IBM Array Summary
- IBM Array Detail
- IBM Extent Pool Summary
- IBM Rank Summary
- IBM Disk Summary
- IBM SAN Volume Controller Overview
- Space-efficient VDisks
- IBM SVC VDisk Summary

- IBM SVC MDisk Groups
- IBM SVC MDisk Group Details
- IBM SVC MDisk Summary
- How Capacity Values are Represented in XIV Reports
- IBM XIV Consistency Groups Summary
- IBM XIV Consistency Group Detail
- XIV Physical Disks
- XIV Snapshot Details
- IBM XIV Storage Pools Summary
- IBM XIV Storage Pool Details
- LUN Utilization Summary
- LUN Utilization Summary for EMC VPLEX Arrays
- LUN Utilization Summary for IBM XIV Arrays
- LUN Utilization Summary for Hitachi Vantara
- LUN Utilization Summary for NetApp 7-Mode
- LUN Utilization Summary for NetApp Cluster-Mode
- LUN Utilization Summary for HP3PAR
- Array Group Details
- EMC CLARiiON Storage Array Group Specifics
- Microsoft Azure Endpoint Summary
- Microsoft Azure Storage Account Summary
- NetApp Aggregate Summary
- NetApp Cluster-Mode Aggregate Summary
- NetApp Aggregate Detail
- NetApp Cluster-Mode Aggregate Detail
- NetApp Plex Details
- NetApp CIFS Summary

- NetApp Cluster-Mode CIFS Summary
- NetApp Disk Summary
- NetApp Cluster-Mode Disk Summary
- NetApp Cluster-Mode FC Array Port Utilization
- NetApp NFS Summary
- NetApp Cluster-Mode NFS Summary
- NetApp Cluster-Mode Node Summary
- NetApp QTree Summary
- NetApp Cluster-Mode QTree Summary
- NetApp Snapshot Summary
- NetApp Cluster-Mode Snapshot Summary
- NetApp Volume Summary
- NetApp Cluster-Mode Volume Summary
- NetApp Volume Details
- NetApp Cluster-Mode Volume Detail
- NetApp Cluster-Mode Summary
- NetApp Cluster-Mode Thin Provisioning Summary
- NetApp Cluster-Mode Vserver Summary
- NetApp StorageGrid Active Alerts Summary
- NetApp StorageGrid Node Summary
- NetApp StorageGrid Site Summary
- NetApp StorageGrid Tenant Summary
- OpenStack Swift Container Summary
- OpenStack Swift Policy Summary
- OpenStack Swift Storage Utilization by Account
- OpenStack Swift Storage Utilization by Zone
- OpenStack Swift Storage Device Summary

- OpenStack Swift Storage Utilization by Node
- Hitachi Vantara Efficiency Summary

Array Capacity and Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Use this report to identify storage arrays that are under-utilized or highly utilized.

Several versions of this report display data relevant to the type of storage:

- See “Array Capacity & Utilization - (Dell Compellent View)” on page 168.
- See “Array Capacity & Utilization - (EMC Data Domain Storage View)” on page 170.
- See “Array Capacity & Utilization - (EMC VNX Celerra View)” on page 173.
- See “Array Capacity & Utilization - (EMC VNX CLARiiON View)” on page 176.
- See “Array Capacity & Utilization - (EMC VPLEX)” on page 179.
- See “Array Capacity & Utilization - (EMC Isilon View)” on page 181.
- See “Array Capacity & Utilization - (EMC Symmetrix View)” on page 183.
- See “Array Capacity & Utilization (Hitachi Content Platform)” on page 187.
- See “Array Capacity & Utilization (Hitachi Arrays & HDP View)” on page 188.
- See “Array Capacity & Utilization (Hitachi NAS Arrays)” on page 192.
- See “Array Capacity & Utilization (HP 3PAR View)” on page 197.
- See “ HP EVA Physical Disks ” on page 240.
- See “Array Capacity & Utilization (IBM Enterprise Array View)” on page 202.
  See “Array Capacity & Utilization (IBM Modular Array View)” on page 203.
- See “Array Capacity & Utilization (IBM SVC View)” on page 203.
- See “Array Capacity and Utilization (IBM XIV View)” on page 207.
  See “How Capacity Values are Represented in XIV Reports” on page 334.
- See “Array Capacity and Utilization (NetApp View)” on page 209.

- See “Array Capacity and Utilization (NetApp Cluster)” on page 214.
- See “ Array Capacity and Utilization (OpenStack Swift) ” on page 229.
- See “Array Capacity & Utilization (Generic Data)” on page 165.

Array Capacity & Utilization (Generic Data)

Each storage array vendor provides subsystem-specific details. Vendor-specific versions of the Array Capacity & Utilization report are available. For certain arrays, a basic set of data is provided in this report. Specifically, when a Data Collector connector has been written with the NetBackup IT Analytics SDK, the Array Capacity & Utilization report will display capacity metrics that are commonly available from array vendors.

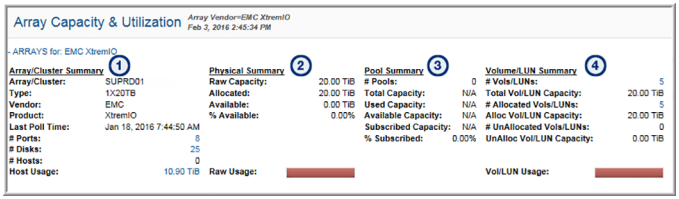


Figure 6.1 Example of an Array Capacity & Utilization Report

This report is organized into four columns, representing common categories of capacity metrics. When this report is generated for a specific storage vendor’s data, these four columns display categories relevant to the vendor.

1. See “Array/Cluster Summary” on page 165.
2. See “Physical Summary or Raw Capacity” on page 166.
3. See “Pool Summary” on page 167.
4. See “Volume/LUN Summary” on page 167.

Array/Cluster Summary

The array/cluster information can be used as a high-level summary of the elements collected from the array or cluster.

Table 9-1 Array/Cluster Summary

Column name	Description
Array/Cluster	The name that you assigned to the array. Available values include all supported storage arrays/clusters.

Table 9-1 Array/Cluster Summary *(continued)*

Column name	Description
Type	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
Last Poll Time	The date and time of the last successful data collection cycle.
# FC Ports	Number of storage director Fibre Channel ports.
# iSCSI Ports	Number of iSCSI ports. Click the link to drill down to the Array iSCSI Port Utilization report.
# Disks	Number of disks for this storage unit (available for certain array vendors). Click the value to view the Disk Summary report. See "Disk Summary" on page 231.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. Click the value to go to the Host Capacity & Utilization report. See " Host Capacity & Utilization " on page 478.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.

Physical Summary or Raw Capacity

This column of information summarizes capacity and usage. The percentages, in particular, capture the essence of your current utilization status.

Table 9-2 Physial Summary or Raw Capacity

Column name	Description
Raw Capacity	Total raw or physical capacity.
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.

Table 9-2 Physial Summary or Raw Capacity *(continued)*

Column name	Description
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the array/cluster that has not been assigned to a host and that is available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
Raw Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.

Pool Summary

View this data for details specific to storage pools.

Volume/LUN Summary

Several fields enable drilldown access to additional details so that you can determine the capacity and usage of each configured volume or LUN.

Note: Make sure the user has the access rights to the objects. For more information, see *NetBackup IT Analytics User Guide >> Manage attributes and objects >> Enable granular user access using authorization attributes* section.

Table 9-3 Volume/LUN Summary

Column name	Description
# Vols/LUNs	The number of LUNs that have been created on the array. Click the link to view the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Total Vol/LUN Capacity	The total capacity for the volumes/LUNs.
# Allocated Vols/LUNs	The number of volumes/LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. Click the link to view the LUN Summary. See " LUN Utilization Summary " on page 344.

Table 9-3 Volume/LUN Summary *(continued)*

Column name	Description
Allocated Vol/LUN Capacity	The total storage capacity for the allocated LUNs.
# UnAllocated Vols/LUNs	The number of LUNs that have yet to be allocated.
UnAllocated Vol/LUN Capacity	The available LUN capacity.
Vol/LUN Usage	Thermometer with a tooltip that displays the percentage allocated, calculated as Allocated divided by (Allocated + Unallocated).

Array Capacity & Utilization - (Dell Compellent View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents storage array details.

Table 9-4 Array Capacity & Utilization - (Dell Compellent View)

Column name	Description
Storage Array	The name that you assigned to the array. Available values include all supported storage arrays/clusters.
Family	The family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Type	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.

Table 9-4 Array Capacity & Utilization - (Dell Compellent View) *(continued)*

Column name	Description
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See " Host Capacity & Utilization " on page 478.
# Ports	The number of ports in the array. Click the value to view the Array Port Utilization.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Physical Summary	
Capacity	Total raw or physical capacity.
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the LUNs in the array that have not been assigned to a host and that are available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
Raw Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.
Raw RAID Group Summary	
# RAID Groups	The number of array groups in the storage array. Links to the List Array Groups. See " List Array Groups " on page 239.
RAID Groups Capacity	Total capacity of the array groups in the storage array.

Table 9-4 Array Capacity & Utilization - (Dell Compellent View) *(continued)*

Column name	Description
RAID Groups Used	Total used capacity of the array groups (by LUNs) in the storage array.
RAID Groups Available	Available capacity of the array groups in the storage array.
RAID Group Usage	Thermometer with a tooltip that displays the percentage used.
LUN Summary	
# LUNs	The number of LUNs that have been created on the array. See " LUN Utilization Summary " on page 344.
Total LUN Capacity	The total capacity for the volumes/LUNs.
# Allocated LUNs	The number of volumes/LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. See " LUN Utilization Summary " on page 344.
Allocated LUN Capacity	The total storage capacity for the allocated LUNs.
# UnAllocated LUNs	The number of LUNs that have yet to be allocated.
UnAllocated LUN Capacity	The available LUN capacity.
LUN Usage	Thermometer with a tooltip that displays the percentage allocated, calculated as Allocated divided by (Allocated + Unallocated).

Array Capacity & Utilization - (EMC Data Domain Storage View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under-utilized or highly utilized.

The Array Capacity & Utilization report represents storage array details.

Table 9-5 Array Capacity & Utilization - (EMC Data Domain Storage View)

Column name	Description
Array Name	Data Domain deduplication storage system name.
Model	Model of the Data Domain system.
OS Version	Version of the Data Domain system's operating system.
Serial #	The serial number that uniquely identifies the Data Domain system.
Filesys Encrypt	Yes (Y) indicates the file system is encrypted.
# of Disks	Number of disk drives in the Data Domain system.
Last Updated	Date and time of the data was last collected.
Physical Summary	
Raw Capacity	Total of all of disk capacities (physical storage).
Active Capacity	Amount of capacity in active storage tiers.
Used Capacity	Amount of Data Domain system capacity that is in use.
Usage	The bar represents the percentage of capacity that is in use.
Archive Capacity	Total amount of usable physical space in the archive.
Used Archive	Archived physical space used for compressed data.
Usage	The bar represents the percentage of archived capacity that is in use.
Logical Summary	
# Storage Units	Number of logical storage units associated with the Data Domain system.
# File Systems	Number of file systems associated with the Data Domain system.
# MTrees	Number of MTrees on the Data Domain system.
# CIFS Shares	Number of CIFS shares on the Data Domain system.

Table 9-5 Array Capacity & Utilization - (EMC Data Domain Storage View) *(continued)*

Column name	Description
# NFS Mounts	Number of NFS mounts on the Data Domain system.
# Files	Number of files associated with the Data Domain system.
Compression Factor	De-duplication ratio (original bytes divided by locally compressed).
Total Original Bytes	The original size before compression.
Globally Compressed	Indicates de-duplication efficiency.
Global Compress Factor	De-duplication ratio (original bytes divided by globally compressed).
Locally Compressed	File compression.
Local Compressed Factor	De-duplication ratio (original bytes divided by locally compressed).
Metadata	Amount of LSU storage consumed by metadata; this is only collected if SNMP is configured as your data collection method.
Performance Summary	
CPU/Avg Busy%	System statistics for the CPU over the 5-second sample interval.
CPU/Max Busy%	System statistics for the CPU over the 5-second sample interval.
Disk/Read(MB/s)	System statistics for disk reads over the 5-second sample interval.
Disk/Write(MB/s)	System statistics for disk writes over the 5-second sample interval.
Disk/Busy	The average percentage (%) of time that each disk has at least one command queued.
NFS/Busy	System statistics for NFS over the 5-second sample interval.
Net/In (MB/s)	System statistics for all Ethernet ports over the 5-second sample interval.
Net/Out (MB/s)	System statistics for all Ethernet ports over the 5-second sample interval.
Replication	
Last Updated	Date and time of the last data collection.

Array Capacity & Utilization - (EMC VNX Celerra View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents EMC VNX File (Celerra) storage array details.

Table 9-6 Array Capacity & Utilization - (EMC VNX Celerra View)

Column name	Description
Storage Array	The name that you assigned to the array. Available values include all supported storage arrays.
Family	The family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor’s product specifications.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
Version	Version of the storage array software.
Status	Status of the array.
# DMs	Number of Data Movers links to the Data Mover report where you can chart CPU and Memory utilization.
# VDMs	Number of Virtual Data Movers links to the Virtual Data Mover report.
# FC Ports	Number of storage director Fibre Channel ports.
# iSCSI Ports	Number of iSCSI ports. Click the link to drill down to the Array iSCSI Port Utilization report.

Table 9-6 Array Capacity & Utilization - (EMC VNX Celerra View) *(continued)*

Column name	Description
# 10/100/1000 Base-T	Number of Ethernet ports.
# Disks	Number of disks links to the Disk Summary report for additional information.
# Control Stations	Number of Control Stations for the Data Movers.
Physical Summary	
Capacity	Sum of the capacity of all physical disks from all array groups, including spares (drives that are not allocated to an array group).
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	% of storage that is available.
Raw Usage	% indicates raw usage.
Raw RAID Group Summary	
# RAID Groups	Number of raw RAID groups for this storage array.
RAID Groups Capacity	Amount of capacity associated with the array's RAID groups.
RAID Groups Used	Amount of raw RAID group capacity that is currently in use.
RAID Groups Available	Amount of raw RAID group capacity that is available.
RAID Group Usage	% of the RAID group capacity that is in use.
Storage Pools	
# Storage Pools	Number of storage pools for this array; links to the Storage Pools Summary report.

Table 9-6 Array Capacity & Utilization - (EMC VNX Celerra View) *(continued)*

Column name	Description
Raw Capacity	The storage pool's raw capacity. Raw capacity is the summation of the physical disks.
Used Capacity	The amount of storage pool capacity that is in use.
Available Capacity	The amount of storage pool capacity that is available.
Potential Capacity	Storage pool capacity, including available space that can be allocated (depending on the associated storage profile).
Thin Pool Total	Total of thin pool capacity for the storage pools.
Thin Pool Used	Amount of thin pool capacity in the storage pools that is in use, not including LUNs that have not yet been physically allocated.
Thin Pool Available	Amount of thin pool capacity in the storage pools that is available.
% Subscribed	% of the thin pool capacity that is subscribed.
Volume Summary	
# Disk Vols	Number of Celerra disk volumes (storage system volumes).
# Slice Vols	Number of slices within the volumes.
# Stripe Vols	Number of strip volumes configured to improve performance.
# Meta Vols	Number of metavolumes; Metavolumes offer a mechanism for creating a logical volume that is larger than a single disk.
# Pool Vols	Number of pool volumes, used to store Celerra SnapSure checkpoint data (SavVol) or Celerra Replicator copies (ConfigVol).
Volume Capacity	The amount of capacity associated with all the volumes.
# Allocated Vols	Number of volumes allocated to Celerra file systems.
Allocated Vol Capacity	Amount of capacity associated with the allocated volumes.
# Unallocated Vols	Number of volumes that have not been allocated to Celerra file systems.
UnAlloc Vol Capacity	Amount of capacity associated with the unallocated volumes.
Vol Usage	% of the volumes that is in use.

Array Capacity & Utilization - (EMC VNX CLARiiON View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized. The Array Capacity & Utilization report represents EMC VNX (CLARiiON) storage array details.

Table 9-7 Array Capacity & Utilization - (EMC VNX CLARiiON View)

Column name	Description
Storage Array	The name that you assigned to the array. Available values include all supported storage arrays.
Family	The family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Type	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
Last Poll Time	Indicates the last time the data was collected from the storage array and persisted to the Portal database.
# Ports	The number of ports in the array. Click the value to view the Array Port Utilization.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See " Host Capacity & Utilization " on page 478.

Table 9-7 Array Capacity & Utilization - (EMC VNX CLARiiON View) *(continued)*

Column name	Description
# Disks	The number of physical disks for this storage array.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Physical Summary	
Capacity	For EMC CLARiiON arrays: Sum of the capacity of all physical disks from all array groups--including spares (drives that are not allocated to an array group).
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the LUNs in the array that have not been assigned to a host and that are available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
Raw Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.
Raw RAID Group Summary	
# RAID Groups	The number of array groups in the storage array. Links to the List Array Groups report. See " List Array Groups " on page 239.
RAID Groups Capacity	Total capacity of the array groups in the storage array.
RAID Groups Used	Total used capacity of the array groups (by LUNs) in the storage array.
RAID Groups Available	Available capacity of the array groups in the storage array.
RAID Group Usage	Thermometer with a tooltip that displays the percentage used.
LUN Summary	

Table 9-7 Array Capacity & Utilization - (EMC VNX CLARiiON View) *(continued)*

Column name	Description
# LUNs	The number of LUNs that have been created on the array. See " LUN Utilization Summary " on page 344.
Total LUN Capacity	The total capacity for the LUNs.
# Allocated LUNs	The number of LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. storage array --> array port --> host storage domain --> HBA port --> host. See " LUN Utilization Summary " on page 344.
Allocated LUN Capacity	The total storage capacity for the allocated LUNs.
# UnAllocated LUNs	The number of LUNs that have yet to be allocated.
UnAllocated LUN Capacity	The available LUN capacity.
LUN Usage	Thermometer with a tooltip that displays the percentage allocated, calculated as Allocated divided by (Allocated + Unallocated).
Thin Provisioning	
# of Pools	Number of thin pools. Drills down to the relevant Thin Pool Summary report.
Raw Capacity	The raw capacity of the pool. Raw capacity is the summation of the physical disks.
User Capacity	The physical capacity that's available to all LUNs in the storage pool; also known as usable capacity or host-visible capacity.
Consumed Capacity	The physical capacity that is assigned to all LUNs in the storage pool. From the array's perspective, consumed capacity is the same as allocated capacity.
% Used	Consumed Capacity divided by User Capacity
Available Capacity	The physical pool storage that is not allocated to pool LUNs
Subscribed Capacity	The physical capacity that has been given to hosts/applications; this is from the host's perspective

Table 9-7 Array Capacity & Utilization - (EMC VNX CLARiiON View) *(continued)*

Column name	Description
% Oversubscribed	Oversubscribed capacity refers to Subscribed Capacity that exceeds User Capacity.

Array Capacity & Utilization - (EMC VPLEX)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Use this report to identify VPLEX storage arrays that are under-utilized or highly utilized. This array information can be used to reference the array vendor's capacity limits.

Table 9-8 Array Capacity & Utilization - (EMC VPLEX)

Column name	Description
Cluster Summary	
Cluster Name	Name of the cluster.
Operational	Operational status: degraded, isolated, ok, or blank.
Health State	Indicates the health of the cluster.
Last Poll Time	The date and time of the last successful data collection cycle.
# of Connected	Number of clusters to which this cluster is connected.
# of Engines	Number of engines: 1, 2, or 4. Each engine includes 2 directors, each having x4 8G FC front-end ports and x4 8G FC back-end ports.
# of FC Ports	Number of fibre-channel ports in the VPLEX instance.
# of Arrays	Number of back-end connected arrays.
# of Storage Volumes	Number of back-end storage volumes.

Table 9-8 Array Capacity & Utilization - (EMC VPLEX) *(continued)*

Column name	Description
# of Extents	Number of extents.
# of Logical Devices	Number of local devices with local visibility, available to the VPLEX cluster.
# of Global Devices	Number of local devices with global visibility, available to resident or remote VPLEX clusters.
# Hosts	Number of hosts associated with this cluster. Drill down to the Host Capacity & Utilization report.
Logical Summary	
Capacity	Total capacity of all Storage Volumes (Unclaimed, Claimed, and Used).
Claimed	Total capacity of all Storage Volumes Claimed.
Unclaimed	Total capacity of all Storage Volumes Unclaimed.
Used	Total capacity of all Storage Volumes Used
Usage	A progress bar (red, yellow, green) indicates usage, based on Used/Capacity.
Virtual Summary	
# Virtual Volumes	Number of all Virtual Volumes managed by this instance.
Virt Vol Capacity	Total capacity of all Virtual Volumes managed by this instance.
# Exp Virt Vol	Number of all Virtual Volumes managed by this instance, exported to hosts.
Ex Virt Vol Capacity	Total capacity of all Virtual Volumes managed by this instance, exported to hosts.
# UnExp Virt Vol	Number of all Virtual Volumes managed by this instance, un-exported to hosts.
UnEx Virt Vol Capacity	Total capacity of all Virtual Volumes managed by this instance, un-exported to hosts.
Virt Vol Usage	A progress bar (green/yellow/red) indicates usage, based on UnEx Virt Vol Capacity/Virt Vol Capacity.
# Storage Views	Total number of Storage Views on this instance of VPLEX.

Array Capacity & Utilization - (EMC Isilon View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Use this report to identify storage arrays that are under-utilized or highly utilized. This array information can be used to reference the array vendor's capacity limits.

Table 9-9 Array Capacity & Utilization - (EMC Isilon View)

Column name	Description
Cluster Summary	
Name	Name of the Isilon cluster.
OneFS Version	Isilon OS version.
Last Poll Time	Date and time of the last poll.
Health	Status of the Isilon cluster.
# Nodes	Number of nodes configured to the Isilon cluster. This link drills down to the Isilon Node report.
# Online Nodes	Number of online nodes in the cluster.
# Offline Nodes	Number of offline nodes in the cluster.
# Chassis	Number of chassis in the cluster.
# Drive Bays	Number of disk bays in all the chassis in the cluster.
# Disks	Number of disks in all the bays in all the chassis in the cluster. This link drills down to the Isilon Disk report.
# Zones	Number of zones associated with the cluster. This link drills down to the EMC Isilon Zones report.
# NFS Exports	Number of NFS exports in the cluster. This link drills down to the Isilon NFS report.
# SMB Shares	Number of Windows server message block (SMB) shares in the cluster. This link drills down to the Isilon SMB Shares report.

Table 9-9 Array Capacity & Utilization - (EMC Isilon View) *(continued)*

Column name	Description
# HDFS Roots	Number of Hadoop Distributed File Systems (HDFS) roots in the cluster. This link drills down to the EMC Isilon HDFS Roots report.
# HDFS Racks	Number of Hadoop Distributed File Systems (HDFS) in the cluster. This link drills down to the Isilon HDFS report.
# FTP Sessions	Number of FTP sessions for the cluster.
Capacity Summary	
Disk Capacity	Sum of capacity of all physical drives in associated block storage arrays.
Raw OneFS Capacity	Total capacity for filessystems on the cluster, including any space in the virtual hot spare reserve.
Virtual Hot Spare Capacity	Total amount of space in the virtual hot spare reserve.
Raw OneFS Usage	(Filesystem Used + Virtual Hot Spare Capacity) / Raw CapacityHover your mouse over the bar to determine the percentage of raw capacity that is in use.
Filesystem Capacity	Total capacity for all filesystems on the cluster. This does not include any space in the virtual hot spare reserve.
Filesystem Used	Total used capacity for all filesystems on the cluster.
Filesystem Usage	Percentage of filesystem used divided by the filesystem capacity. This does not include any space in the virtual hot spare reserve.
Snapshot Summary	
Snapshot Create	Indicates if Snapshot Create is enabled. Possible values: enabled or disabled.
Snapshot Disable	Indicates if Snapshot Disable is enabled. Possible values: enabled or disabled.
Snapshot Reserve	Percentage of storage space reserved for snapshots; a percentage of the total cluster storage capacity.
# Snapshots	Number of active snapshots. This link drills down to the Isilon Snapshot report.
Snapshot Used	Sum of the snapshots' sizes.
Performance Summary	

Table 9-9 Array Capacity & Utilization - (EMC Isilon View) *(continued)*

Column name	Description
# Connected Clients	Number of clients that are connected to the cluster, at the time of the last data collection. These clients have unique addresses with established TCP connections to the cluster on known ports.
# Active Clients	Number of clients that are active to the cluster, at the time of the last data collection. This takes into account client addresses generating protocol traffic on the monitored cluster.
Network Throughput	Network I/O rate (bps), at the time of the last data collection.
File Throughput	Filesystem I/O rate (bps), at the time of the last data collection.
Avg CPU	Average CPU as of the last five seconds, at the time of the last data collection. Hover your mouse over the bar to determine the percentage.
License Summary	
# Licenses	Total number of Isilon licenses. This link drills down to the License Summary report.
# Active	Number of licenses in the Active state.
# Evaluation	Number of evaluation licenses.
# Expired	Number of expired licenses.

Array Capacity & Utilization - (EMC Symmetrix View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Use this report to identify storage arrays that are under-utilized or highly utilized.

Raw capacity is the summation of the physical disks. Usable space is the space that is calculated from RAID groups and LUNs.

This array information can be used to reference the array vendor's capacity limits.

Table 9-10 Array Capacity & Utilization - (EMC Symmetrix View)

Column name	Description
Storage Array	The name that you assigned to the array. Available values include all supported storage arrays.
Family	The family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Type	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
# FC Ports	Number of Fibre Channel ports. See " Array FC Port Utilization " on page 241.
# EF Ports	Number of EF - FICON (Fibre-ESCON) director ports.
# RA Ports	Number of RA - RDF director ports. Links to the RA Summary report.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected.
Host Usage	Amount of capacity used by hosts.
Physical Summary	
Capacity	Sum of the capacity of all physical disks from all array groups, allocated, hot spares and unallocated.
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Hot Spares	Total capacity of hot spare drives.
Unallocated	Unallocated storage is storage that is not part of an Array Group
Raw Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.

Table 9-10 Array Capacity & Utilization - (EMC Symmetrix View) *(continued)*

Column name	Description
Save Pools	
# of Pools	Number of Save Pools for this storage array, representing the point-in-time copy of data
Pool Capacity	Total physical capacity for all the Save Pools associated with this array. Note that Save Pool capacity can make the Total FB LUN Capacity appear inflated or duplicated.
Pool Capacity Enabled	Amount of Save Pool capacity that has been enabled for host writes.
Pool Capacity Used	Amount of the Save Pool capacity that has already been consumed by hosts.
Pool Usage	Expressed as a percentage, this is the capacity of the Save Pools that has already been used by hosts.
Raw RAID Group Summary	
# RAID Groups	The number of array groups in the storage array. See " List Array Groups " on page 239.
RAID Groups Capacity	Total capacity of the array groups in the storage array.
RAID Groups Used	Total used capacity of the array groups (by LUNs) in the storage array.
RAID Groups Available	Available capacity of the array groups in the storage array.
RAID Group Usage	Thermometer with a tooltip that displays the percentage used.
Thin Provisioning	
# of SRPs	Links to the EMC Symmetrix SRP Summary for storage resource pool information.
# of Pools	Links to EMC Symmetrix Thin Pool Summary.
Pool Capacity	Total capacity of the thin pool.
Pool Used	Amount of thin pool capacity in the storage pools that is in use, not including LUNs that have not yet been physically allocated.
% Used	Pool used divided by pool capacity = % used.

Table 9-10 Array Capacity & Utilization - (EMC Symmetrix View) *(continued)*

Column name	Description
Pool Subscribed	The physical capacity that has been given to hosts/applications; this is from the host's perspective.
% Subscribed	Amount of capacity that has been subscribed to users.
LUN Summary	
# FB LUNs	Number of FB (fixed block) LUNs. See " LUN Utilization Summary " on page 344.
Total FB LUN Capacity	Total capacity of all the FB (fixed block) LUNs for this array. See the Save Pool section of this report to verify capacity values.
#Allocated FB LUNs	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Alloc FB LUN Capacity	Total storage capacity for the FB (fixed block) LUNs.
# UnAllocated FB LUNs	The number of FB (fixed block) LUNs that have yet to be allocated. See " LUN Utilization Summary " on page 344.
UnAlloc FB LUN Capacity	Total unallocated capacity for FB (fixed block) LUNs.
FB LUN Usage	Mouse over the thermometer to view the percentage of the fixed block LUNs that are allocated.
# CKD LUNs	Total number of Count-Key-Data (CKD) LUNs.
CKD LUN Capacity	Total capacity of Count-Key-Data (CKD) LUNs.
# CKD Allocated LUNs	Total number of allocated Count-Key-Data (CKD) LUNs.
CKD Alloc LUN Capacity	Total allocated capacity of Count-Key-Data (CKD) LUNs.
# CKD UnAllocated LUNs	Total number of unallocated Count-Key-Data (CKD) LUNs.
CKD UnAlloc LUN Capacity	Total capacity of unallocated Count-Key-Data (CKD) LUNs.
CKD LUN Usage	Mouse over the thermometer to view the percentage of the Count-Key-Data (CKD) LUNs that are allocated.

Array Capacity & Utilization (Hitachi Content Platform)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents Hitachi Content Platform (HCP) capacity summaries, with drilldowns to relevant details.

Table 9-11 Array Capacity & Utilization (Hitachi Content Platform)

Column name	Description
Summary	
Cluster	The name that you assigned to the array. Available values include all supported storage devices. For a list of supported devices, see the Certified Configurations Guide.
Model	The model of the cluster/array.
Type	The specific model of the cluster.
Vendor	The name of the vendor that distributes the cluster.
Version	Version of the cluster software.
MIB Version	Version of the HCP controller MIB.
Serial Number	Serial number of the cluster.
Last Poll Time	Date and time the data was collected.
# Nodes	Drill down to the Node Summary to access capacity details.
# Tenants	Drill down to the Tenant Summary to view tenant and namespace capacity details.

Table 9-11 Array Capacity & Utilization (Hitachi Content Platform) *(continued)*

Column name	Description
# File Systems/Volumes/Namespaces	Drill down to the Namespace Summary for capacity and object data.
# Files/Objects	Number of files/objects associated with the cluster.
Logical Summary	
Total Capacity	Total amount of logical capacity for the HCP cluster.
Used	Amount of the logical capacity that is in use.
Available	Amount of the logical capacity that is available to be allocated.
Storage Pool Summary	
# Storage Pools	Number of storage pools associated with this cluster.
Storage Pool Capacity	Storage capacity of the storage pool.
Storage Pool Used	Amount of the storage pool capacity that is in use.
Storage Pool Available	Amount of the storage pool capacity that is available.
Available	The percentage of the storage pool capacity that is available.

Array Capacity & Utilization (Hitachi Arrays & HDP View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents Hitachi Arrays and Hitachi Dynamic Provisioning (HDP) in the following manner:

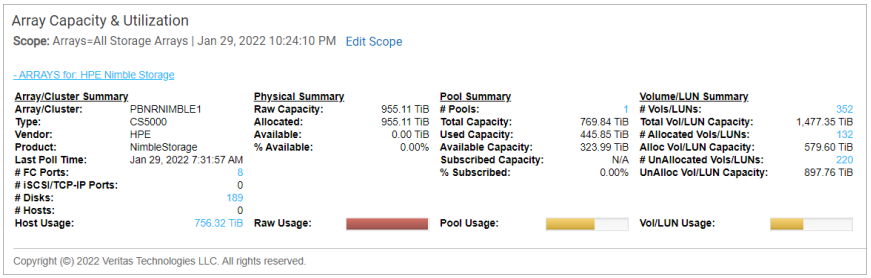


Table 9-12 Array Capacity & Utilization (Hitachi Arrays & HDP View)

Column name	Description
Storage Array	The name that you assigned to the array. Available values include all supported storage devices. For a list of supported devices, see the Certified Configurations Guide.
Family	The family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Type	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
# Ports	The number of ports in the array. Click the value to view the Array Port Utilization.
# Disks	The number of disks in the array.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See " Host Capacity & Utilization " on page 478.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
% Used	The percentage of the storage that is currently used.

Table 9-12 Array Capacity & Utilization (Hitachi Arrays & HDP View) *(continued)*

Column name	Description
Physical Summary	
Raw Capacity	For HDS arrays: Sum of capacity of all PDEVs from all array groups. When calculating raw capacity for HDS arrays, the data collector makes requests to the Device Manager for physical devices (PDEVs) for all Array Groups in a storage array. The data collector assumes that all PDEVs are part of an array group, so PDEVs that are marked as not included in an array group (-1) are included in the capacity calculation. That is, no disks are treated as spares.
Configured RAID	Total raw capacity of the LUNs in the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application.Unallocated storage is storage that is not part of an Array Group.
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the LUNs in the array that have not been assigned to a host and that are available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
# RAID Groups	The number of array groups in the storage array. See " List Array Groups " on page 239.
Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.
Usable Summary (post-RAID)	
Internal Capacity	Amount of usable capacity from physical RAID groups.
External Capacity	Applicable to HDS arrays only - Capacity from block storage, derived from daisy-chained external storage. This is array group capacity that does not have PDEVs in the storage array.
Used Capacity	Amount of internal RAID group capacity allocated to LUNs.
Free Capacity	Amount of available capacity in RAID groups from which to allocate LUNs.
Usage	Thermometer with a tooltip that displays the RAID Group Usage percentage

Table 9-12 Array Capacity & Utilization (Hitachi Arrays & HDP View) *(continued)*

Column name	Description
Dynamic Provisioning	
# Pools	The number of dynamic provisioning pools links to the Dynamic Provisioning Pool Summary
Capacity	The Dynamic Provisioning capacity for this array
Used	Amount of the dynamically provisioned capacity in use
Free	Amount of the dynamically provisioned capacity that is available.
DP Usage	Mouse over this thermometer to view the usage.
LUN Summary	
# Open Reserved LUNs	For Hitachi USP arrays, these are reserved LUNs that cannot be assigned to hosts; for example, dynamic pool capacity. No paths can be set for these LUNs.
Open Resv LUN Capacity	For Hitachi USP arrays, this is the capacity of the reserved LUNs that cannot be assigned to hosts; for example, dynamic pool capacity. No paths can be set for these LUNs.
# Open Allocated LUNs	The number of open LUNs (volumes) that have set paths.
Open Alloc LUN Capacity	Represents the capacity of the open LUNs (volumes) that have set paths.
# Open Unallocated LUNs	The number of open LUNs that do not have set paths.
Open Unalloc LUN Capacity	The capacity of the open LUNs that do not have set paths.
# LUNs	Number of LUNs present in the storage array.
Total LUN Capacity	Sum of the capacity of all LUNs for the storage array.
LUN Usage	The percentage of the LUN that is allocated.

Array Capacity & Utilization (Hitachi NAS Arrays)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents Hitachi NAS (HNAS) arrays in the following manner:

Table 9-13 Array Capacity & Utilization (Hitachi NAS Arrays)

Column name	Description
Storage Array	The name of the Hitachi NAS (HNAS) array.
Model	The model of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Status	The status of the array; for example, online or offline.
Product	The name of the product for this Hitachi NAS array. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
# Storage Servers	Click the value to view the HNAS Storage Server Summary. See "Hitachi NAS Storage Server Summary" on page 304.
# Storage Subsystems	Number of storage sub-systems associated with this array.
# EVS	Number of Enterprise Virtual Servers (EVS). See "Hitachi NAS EVS Summary" on page 297.
# 10G Ports	The number of 10G ports on all Enterprise Virtual Servers (EVS). See "Hitachi NAS Ethernet Port Summary" on page 298.
# Disks	Click the value to view the HNAS Disk Summary. See "Hitachi NAS Disk Summary" on page 296.

Table 9-13 Array Capacity & Utilization (Hitachi NAS Arrays) *(continued)*

Column name	Description
# CIFS Shares	Click the value to view the HNAS CIFS Share Summary. See "Hitachi NAS CIFS Share Summary" on page 295.
# NFS Shares	Click the value to view the HNAS NFS Export Summary. See "Hitachi NAS NFS Export Summary" on page 300.
# FTP Mounts	Sum of the FTP mount points.
# iSCSI LUNs	Sum of the iSCSI Block LUNs.
Physical Summary	
Internal Capacity	The Hitachi NAS (HNAS) filer heads will not have any internal capacity, as they rely on block storage arrays.
External Capacity	Sum of the capacity of all block storage array LUNs being virtualized by the Hitachi NAS filer.
Allocated	Sum of the capacity of all the system drives allocated to Storage Pools.
Available	Sum of the capacity of all system drives not allocated to Storage Pools.
% Available	Expressed as a percentage, Available/Internal capacity.
Raw Usage	Expressed as a percentage, Allocated/Internal capacity.
Storage Pools Summary	
# Storage Pools	Sum of the array's Storage Pools. See "Hitachi NAS Storage Pool Summary" on page 303.
Storage Pool Capacity	Sum the of array's Storage Pool capacity for File Systems (net RAID parity and overhead).
Storage Pool Used	Sum of the array's File System used capacity.
Storage Pool Available	The Storage Pool Capacity minus the Storage Pool Used.
File System Usage	Thermometer with a tooltip that displays % of Used/Capacity.
File System Summary	
# File Systems	Number of array file systems. See "Hitachi NAS File System Summary" on page 298.

Table 9-13 Array Capacity & Utilization (Hitachi NAS Arrays) *(continued)*

Column name	Description
# Thin Prov File Systems	Number of file systems that are thin. See "Hitachi NAS File System Summary" on page 298.
File System Capacity	Sum of the array's file system sizes.
File System Used	Sum of the array's file systems' used capacity.
File System Available	File System Capacity minus File System Used
File System Usage	Percentage: Used/Capacity
LUN Summary	
# LUNs	Number of iSCSI LUNs.
Total LUN Capacity	Sum of the array's iSCSI LUN sizes.
# Allocated LUNs	Sum of the array's iSCSI LUNs allocated to hosts.
Alloc LUN Capacity	Sum of the array's iSCSI LUNs allocated sizes.
# Unallocated LUNs	Sum of the array's iSCSI LUNs not allocated to hosts.
Unalloc LUN Capacity	Sum of the sizes of the array's iSCSI LUNs not allocated to hosts.
LUN Usage	The percentage of the LUN that is allocated, Capacity/Total Capacity.

Array Capacity & Utilization - (Hitachi Vantara All-Flash and Hybrid Flash Storage)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity & Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents Hitachi Vantara All-Flash and Hybrid Flash Storage in the following manner:

Table 9-14 Array Capacity & Utilization (Hitachi Vantara All-Flash and Hybrid Flash Storage view)

Column name	Description
Array/Cluster Summary	
Array Name	The name that you assigned to the array. Available values include all supported storage devices. For a list of supported devices, see the Certified Configurations Guide.
Serial #:	Serial number of the cluster.
IP Address	IP address of the flash storage.
Model	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
Last Poll Time	The date and time of the last successful data collection cycle.
Firmware Version	Current firmware version of the array.
# FC Ports	The number of FC ports in the array. Click the value to view the Array Port Utilization.
# iSCSI Ports	Number of iSCSI ports. Click the link to drill down to the Array iSCSI Port Utilization report.
# Disks	The number of disks in the array. Click the value to view the Disk Summary report.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See “ Host Capacity & Utilization ” on page 478.

Table 9-14 Array Capacity & Utilization (Hitachi Vantara All-Flash and Hybrid Flash Storage view) *(continued)*

Column name	Description
Cache	Cache size of the array. Values are stored as TiB in the database and rendered according to your user profile preferences.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Physical Summary	
Raw Capacity	For HDS arrays: Sum of capacity of all PDEVs from all array groups. When calculating raw capacity for HDS arrays, the data collector makes requests to the Device Manager for physical devices (PDEVs) for all Array Groups in a storage array. The data collector assumes that all PDEVs are part of an array group, so PDEVs that are marked as not included in an array group (-1) are included in the capacity calculation. That is, no disks are treated as spares.
Allocated	Total of all of disk capacities (physical storage).
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the LUNs in the array that have not been assigned to a host and that are available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
# Array Groups	Number of Array Groups. Links to List Array Groups report.
Pool Summary	
# Pools	The number of dynamic provisioning pools links to the Dynamic Provisioning Pool Summary
Total Capacity	The total capacity for the volumes/LUNs.
Used Capacity	Amount of internal RAID group capacity allocated to LUNs.
Available Capacity	Amount of available capacity in RAID groups from which to allocate LUNs.
# Resv Vols/LUNs	The number of LUNs that are reserved on the array. Click the link to view the LUN Utilization Summary.

Table 9-14 Array Capacity & Utilization (Hitachi Vantara All-Flash and Hybrid Flash Storage view) *(continued)*

Column name	Description
# Resv Vols/LUN Capacity	Capacity of the reserved LUNs that cannot be assigned to the hosts. For example, dynamic pool capacity.
Volume/LUN Summary	
# Vols/LUNs	The number of LUNs that have been created on the array. Click the link to view the LUN Utilization Summary.
Total Vol/LUN Capacity	The total capacity for the volumes/LUNs.
# Allocated Vols/LUNs	The number of volumes/LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. Click the link to view the LUN Summary.
Alloc Vol/LUN Capacity	The total storage capacity for the allocated LUNs.
# UnAllocated Vols/LUNs	The number of LUNs that have yet to be allocated. Click the link to view the LUN Summary.
UnAlloc Vol/LUN Capacity	The available LUN capacity.

HP 3PAR Overview

HP 3PAR is classified in a category of disk arrays for Utility Storage. These arrays serve as the storage foundation for utility computing architectures. Utility computing architectures provide a multi-tenant platform on which service providers can deliver both virtualized and scalable enterprise IT as a utility service.

Array Capacity & Utilization (HP 3PAR View)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity and Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The following illustration represents the view available for HP 3PAR storage arrays.

The report is organized into the following areas:

1. Array Summary
2. Physical Summary
3. Chunklets
4. Logical (LD) Summary
5. Common Provisioning Groups (CPG)
6. Virtual Volume (VV) Summary

Array Summary

Use array information to review the HP 3PAR capacity limits.

Table 9-15 Array Summary

Column name	Description
Storage Array	Name assigned to the array. Available values include all supported storage devices. For a supported devices list, see the Certified Configurations Guide.
Family	Family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Type	Specific model of the storage array.
Vendor	Name of the vendor that distributes the storage array.
Product	Name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
O/S Firmware	Version number for operating system firmware.
Serial No	Array's serial number.
Last Updated	Date and time this data was last updated in the database
# Controllers	Number of controllers managing the array.
# FC Ports	Number of fibre channel ports. Links to Array Port Utilization.

Table 9-15 Array Summary *(continued)*

Column name	Description
#iSCSI Ports	Number of ports allocated for Internet Small Computer System Interface (iSCSI), an IP-based storage networking standard for linking data storage facilities. iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances. Click the link to drill down to the Array iSCSI Port Utilization report.
# Disk Drives	Number of disks associated with this array. Click the link to view the Disk Summary for relevant disk drives and their capacities.
# Hosts	Total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See " Host Capacity & Utilization " on page 478.
Host Usage	Represents the amount of array storage being used by hosts and their applications.
% Used	Percentage of total capacity used, calculated as Host In-Use divided by Host Allocated.

Physical Summary

This section of the report summarizes capacity and usage. The percentages, in particular, capture the essential part of your current utilization status.

Table 9-16 Physical Summary

Column name	Description
Capacity	Total raw capacity of the storage array. You can set the capacity units to be displayed, either MB, GB, TB, or PB, by selecting the Advanced option when you generate the report.
Used	Capacity associated with used (written to) chunklets that represent valid data contents available through logical drive mappings to hosts.
Free	Capacity associated with used (written to) chunklets that at one time represented valid data contents available through logical drive mappings to hosts, but do not any longer. Typically resulting from failed drives. All Normal Used OK chunklets on a drive that fail are flagged as Normal Used Failed.

Chunklets

Table 9-17 Chunklets

Column name	Description
Normal Used OK	Capacity associated with used (written to) chunklets that represent valid data contents available through logical drive mappings to hosts.
Normal Used Failed	Capacity associated with used (written to) chunklets that at one time represented valid data contents available through logical drive mappings to hosts, but do not any longer. Typically resulting from failed drives. All Normal Used OK chunklets on a drive that fail are flagged as Normal Used Failed.
Normal Unused Free	Capacity associated with unused and initialized chunklets available to establish new logical drives or expand existing logical drives.
Normal Unused Uninitialized	Capacity associated with replacement drives that have yet to be initialized.
Normal Unused Failed	Capacity associated with unused free disk drives that fail.
Spare Used OK	Capacity associated with used (written to) chunklets that represent valid data contents available through logical drive mappings to hosts, as a result of a redundant RAID group drive rebuild.
Spare Used Failed	Capacity associated with used (written to) chunklets that at one time represented valid data contents available through logical drive mappings to hosts that resulted from a redundant RAID group drive rebuild, but do not any longer. Typically resulting from failed drives. All Spare Used OK chunklets on a drive that fails are flagged as Spare Used Failed.
Spare Unused Free	Capacity associated with unused and initialized chunklets available to redundant RAID groups for drive rebuild on a failed drive.
Spare Unused Uninitialized	Capacity associated with replacement spare drives that have yet to be initialized.
Spare Unused Failed	Capacity associated with unused free spare disk drives that fail.
Chunklet Usage	Normal Used OK + Normal Used Failed/Total in graph form. Roll over to display a usage percentage.

Logical (LD) Summary

Logical Disks are a collection of physical disk chunklets arranged as rows of RAID sets. This section of the report summarizes their capacity and usage.

Table 9-18 Logical (LD) Summary

Column name	Description
# Logical Drives (LD)	Number of logical drives. Refers to an arrangement of rows of RAID sets. Logical drives are mapped to virtual volumes.
LD Capacity	Capacity of all the logical disk space.
LD Used	Logical Disk (LD) space mapped to a Virtual Volume.
LD Free	Logical Disk (LD) Capacity - Logical Disk (LD) Used = LD Free.
LD Usage	Logical Disk (LD) Used/Logical Disk (LD) Capacity in graph form.

Common Provisioning Groups (CPG)

A Common Provisioning Group (CPG) creates a virtual pool of logical disks that allows up to 64 volumes to share the CPG's resources and allocate space on demand.

Table 9-19 Common Provisioning Groups (CPG)

Column name	Description
# CPGs	Number of Common Provisioning Groups. This number includes delta replication pools. Links to HP 3PAR Common Provisioning Group (CPG) Summary.
# Associated LDs	Number logical disks associated with each common provisioning group.
Consumed Capacity	Sum of current space consumed by the storage pool.
Allocated for Snapshot	Sum of each CPG's snapshot allocation.
Allocated for Admin	Space allocated for administrative overhead.

VV Summary

A virtual volume is a representation of a storage area that is not tied to any physical storage allocation. VVs are used in thin provisioning (also known as virtual provisioning) for SAN management. A virtual LUN is also sometimes called a thin LUN. HP 3PAR refers to these as virtual volumes.

Table 9-20 VV Summary

Column name	Description
# VVs	Number of virtual volumes. This report links to the LUN Utilization Summary.
# TPVVs	Number of thinly provisioned virtual volumes (TPVVs).
# CPVVs	Number of concrete provisioned virtual volumes (CPVVs). Links to the LUN Utilization Summary.
Total VV Capacity	Total storage capacity for all of the associated virtual volumes.
# Allocated VVs	Number of virtual volumes on the array that have been mapped to a host. Links to the LUN Utilization Summary - Allocated report.
Alloc VV Capacity	Total storage capacity for the allocated virtual volumes.
# UnAllocated VVs	Number of virtual volumes that have yet to be allocated. This report links to LUN Utilization Summary - Unallocated.
UnAlloc VV Capacity	Available capacity for unallocated virtual volumes.
VV Usage	Percentage of virtual volumes capacity currently in use.
#VLUNs	Number of Storage Paths.

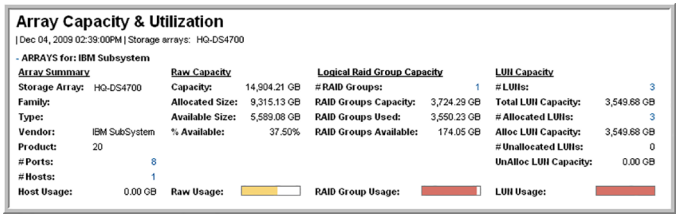
Array Capacity & Utilization (IBM Enterprise Array View)

The following illustration represents an example of the view available for IBM Enterprise Arrays (6K/8K). For a description of the fields displayed refer to IBM Array Detail.

See “ IBM Array Detail ” on page 314.

Array Capacity & Utilization (IBM Modular Array View)

The following examples shows an IBM Modular (3K) array.



Array Capacity & Utilization (IBM SVC View)

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Choose a report scope that includes only IBM SVC storage arrays.

The following illustration represents an example of the view available for IBM SVC.

See "How IBM SVC Capacity Values Are Derived" on page 206.

See "Array Capacity & Utilization (HP 3PAR View)" on page 197.

Table 9-21 Array Capacity & Utilization (IBM SVC View)

Column name	Description
Storage Array	The name that you assigned to the array. Available values include all supported storage arrays.
Family	The family classification of the storage array as assigned by the vendor. Use this name to ascertain the vendor's product specifications.
Type	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number or alphanumeric string. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
# Nodes	Number of SAN Volume Controllers, the providers of virtualization to the SAN.

Table 9-21 Array Capacity & Utilization (IBM SVC View) *(continued)*

Column name	Description
# Ports	The number of ports in the array. Click the value to view the Array Port Utilization.
Total Free Space	The currently available space in the cluster. This is approximately backend storage capacity minus allocated capacity. Indicates the sum of the free capacity for the mdisk group.
% RAID Loss	This value represents the amount of space lost to RAID overhead. It includes mirroring and other overhead usage. The value is derived using the following: (non-primary mirror extents + managed MDisks) - ((MDisk Group real capacity - MDisk Group remaining managed space)/MDisks Total)
# Hosts	The total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See " Host Capacity & Utilization " on page 478.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
% Host Usage	The percentage of the total capacity that is being used by hosts
# Disks	The number of physical disks for this storage array. Links to the Disk Summary report for additional information.
Capacity	Total physical capacity. This is populated as physical capacity from the Issystem CLI. For certain models, when physical capacity is not available, it is calculated as the sum of all the drives capacity.
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity. This is calculated as physical capacity minus physical free capacity as obtained from Issystem CLI when available.
Hot Spare	Total capacity of hot spare drives.
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.

Table 9-21 Array Capacity & Utilization (IBM SVC View) *(continued)*

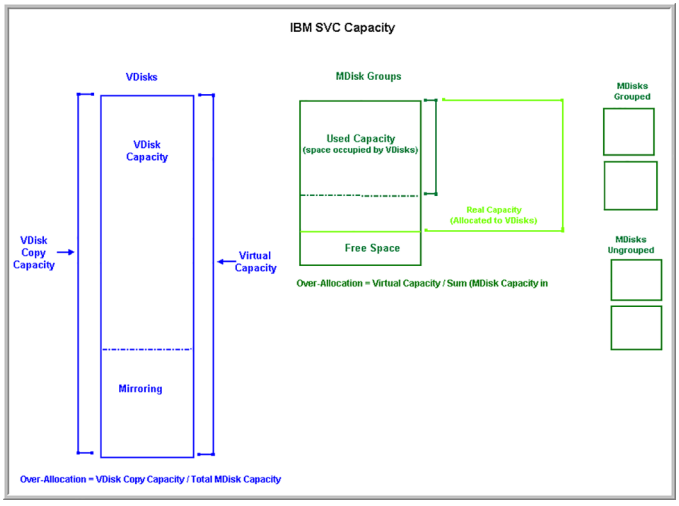
Column name	Description
% Available	Expressed as a percentage, total capacity of the array/cluster that has not been assigned to a host and that is available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
Raw Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.
#MDisks	The number of MDisks associated with this array.
Total MDisk Capacity	Total MDisk (block storage) capacity
Managed Space	MDisks that are in Managed mode; click on the #MDisks to view the IBM SVC MDisk Summary, which lists the MDisks and modes
UnManaged Space	MDisks that are in UnManaged mode; click on the #MDisks to view the IBM SVC MDisk Summary, which lists the MDisks and modes.
MDisk Usage	Expressed as a percentage, total capacity of the LUNs in the array that have been assigned to a host and that are not available for provisioning.
# MDisk Groups	The number of managed disk groups for the VDisks
Space in MDisk Groups	The capacity in the MDisk groups, derived from the SVC cluster pool capacity.
Space Allocated to VDisks	The amount of MDisk capacity that has been allocated to VDisks
Space Available in MDisk Groups	MDisk Group total capacity minus the space allocated to VDisks. Note that some space may be reserved for use by SVC.
MDisk Group Usage	Mouse over the thermometer to view the percentage, derived from the Total MDisk Capacity divided by the Managed Space.
# VDisks	The number of VDisks associated with this array
Total VDisk Copy Capacity	Virtual capacity, including mirrors
Total VDisk Capacity	Virtual capacity, excluding mirrors
Total Used Capacity	Total VDisk used capacity

Table 9-21 Array Capacity & Utilization (IBM SVC View) *(continued)*

Column name	Description
% Over-Allocation	When this value is greater than 100%, it indicates that the VDisk is space-efficient. The percentage represents the ratio of virtual to real capacity.
# VDisks Allocated to Hosts	The number of VDisks allocated to hosts
Alloc VDisk Capacity	The amount of VDisk capacity that has been presented to hosts
# VDisks Unallocated	The number of VDisks that have yet to be presented to hosts
UnAlloc VDisk Capacity	The amount of VDisk capacity that has yet to be presented to hosts
VDisk Usage	Mouse over the thermometer to view the percentage, derived from the Total Used Capacity divided by the Total VDisk Copy Capacity

How IBM SVC Capacity Values Are Derived

The following diagram illustrates how NetBackup IT Analytics represents specific capacity values, such as Copy Capacity and Free Space.



Array Capacity and Utilization (IBM XIV View)

The report is organized into four columns:

1. See “XIV Array Summary” on page 207.
2. See “Physical Summary” on page 208.
3. See “Usable Volumes” on page 208.
4. See “Logical Volumes” on page 209.

See “How Capacity Values are Represented in XIV Reports” on page 334.

XIV Array Summary

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > XIV Storage Pools Summary

Next, click a link from the Storage Array column.

Table 9-22 XIV Array Summary

Column name	Description
XIV Array Summary	
Storage Array	Name of the XIV array
Version	Software version
Array Type	XIV array type
Array Model	XIV array model
Serial #	Array serial number
iSCSI Name	iSCSI qualified name of the array
System ID	ID of the array
# FC Ports	Links to Array Port Utilization report for these Fibre Channel ports
# iSCSI Ports	Links to Array Port Utilization report for these iSCSI ports
# Disks	Links to XIV Drive Detail
# Hosts	Links to Host Capacity & Utilization report

Physical Summary

The Physical Summary provides system-wide capacity status.

Table 9-23 Physical Summary

Column name	Description
Physical Summary	
# Pools	Links to XIV Storage Pools Summary
Soft Limit	The logical amount of space for all storage pools in the array. A soft limit greater than the hard limit indicates thin provisioning.
Hard Limit	Maximum physical capacity that can be used by volumes and snapshots in the pool.
Free Hard	Amount of unused physical space in the array
Free Soft	Unused amount of the logical storage that has been presented to hosts and applications
Hard Utilization	(hard limit - free hard)/hard limit
Spare Disks	# of spare disks in the array
Spare Modules	# of spare modules in the array
Target Spare Disks	# of spare disks on the remote target system
Target Spare Modules	# of spare modules on the remote target system

Usable Volumes

The Usable Volumes fields provide an overview of available storage capacity.

Table 9-24 Usable Volumes

Column name	Description
Usable Summary	
Disk Space	Sum of the number of drives times the drive capacity
Available Disk Space	The value is calculated using the following: [(number of disks – number of spare disks) - (number of spare modules x 12)] x disk capacity x 0.96/2

Table 9-24 Usable Volumes *(continued)*

Column name	Description
Formatted Volume Space	Sum of all volume sizes, or total allocated (or provisioned) storage
Formatted Snapshot Space	Sum of volume sizes minus the sum of all volumes used
Formatted Space	Formatted available and usable storage
Formatted Space Used	The Formatted Space Used value is used to derive Formatted Utilization
Formatted with No Volumes	The unused portion of the provisioned capacity shared by all volumes; volumes do not consume storage until pages are written by the OS; all LUNs are thin
Overall Unavailable Disk Space	Metadata and system reserve used for traces and statistics
Subscription	The thermometer tooltip indicates if storage has been over-subscribed.
Formatted Utilization	The thermometer tooltip indicates the formatted storage utilization %, calculated as Formatted Space Used/Formatted Space

Logical Volumes

Logical Volumes in XIV Arrays represent the amount of logical capacity that has been made available to hosts (and applications on those hosts).

Table 9-25 Logical Volumes

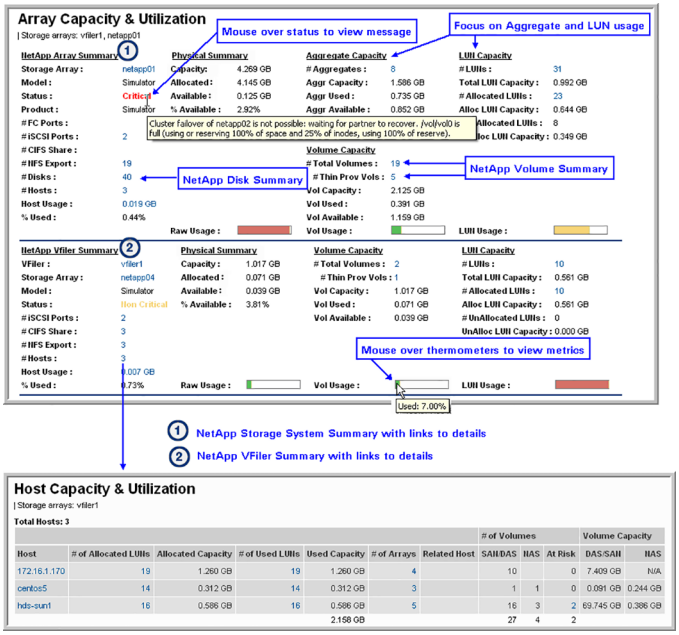
Column name	Description
Logical Volumes	
# of Volumes	Links to the LUN Utilization Summary report
# of Allocated LUNs	Links to the LUN Utilization Summary report, which displays the allocated LUNs

Array Capacity and Utilization (NetApp View)

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Because NetApp utilizes a number of unique constructs--such as Aggregates and Volumes--to organize and access its storage, additional data is available via drilldown links.



The report is organized into four columns:

1. See “NetApp Summary” on page 210.
2. See “Physical Summary” on page 211.
3. See “Aggregate/Volume Capacity” on page 212.

   - Pools of available storage can be found in existing RAID groups. Even though storage may be 100% allocated at the array level, the RAID group utilization will reveal the amount of remaining, usable storage.
4. See “LUN Capacity” on page 213.

NetApp Summary

This section of the report provides a overview of the storage system, with links to relevant components.

Table 9-26 NetApp Summary

Column name	Description
VFiler	The VFiler name
Storage Array	The name that you assigned to the array. Available values include all supported storage devices. For a list of supported devices, see the Certified Configurations Guide.
Model	The specific model of the NetApp storage.
Status	The current status of the storage. Mouse over this status to view details.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
# FC Ports	Number of Fibre Channel ports
# iSCSI Ports	Links to Array iSCSI Port Utilization. See " Array iSCSI Port Utilization " on page 224.
# CIFS Share	Links to NetApp QTree Summary. See " NetApp QTree Summary " on page 389.
# NFS Export	Links to NetApp NFS Summary. See " NetApp NFS Summary " on page 383.
# Disks	Links to NetApp Disk Summary. See " NetApp Disk Summary " on page 378.
# Hosts	The number of hosts that are sharing the capacity. This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
% Used	Percentage of the NetApp storage already utilized.

Physical Summary

This column of information summarizes capacity and usage. The percentages, in particular, capture the essence of your current utilization status.

Table 9-27 Physical Summary

Column name	Description
Capacity	Total raw capacity of the storage array. This value is the sum of all PDEVs (including spares--that is, drives that are not allocated to an Aggregate). If the RAID state of the PDEV is partner or broken, the disk is not included in the raw capacity calculation. You can set the capacity units to be displayed, either GB or MB, by selecting the Advanced option when you generate the report.
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the LUNs in the array that have not been assigned to a host and that are available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
Raw Usage	Expressed as a percentage, total capacity that has been assigned to a host and that is not available for provisioning.

Aggregate/Volume Capacity

This column provides an overview of the aggregate (not for VFilers) and volume (for VFilers) capacity and utilization.

Table 9-28 Aggregate/Volume Capacity

Column name	Description
# Aggregates	Links to the Netpp Aggregate Summary. See “ NetApp Aggregate Summary ” on page 364.
Aggr Capacity	The aggregate’s total size
Aggr Used	Amount of the aggregate’s capacity that is in use
Aggr Available	Available storage in the aggregate
Aggr Usage	Used data storage at the physical layer. Mouse over the thermometer bar to view the % available.

Table 9-28 Aggregate/Volume Capacity *(continued)*

Column name	Description
# Total Volumes	The number of volumes defined for this NetApp unit. See " NetApp Snapshot Summary " on page 392.
# Thin Prov Vols	The number of thin provisioned volumes; Includes volumes where the Space Guarantee value equals "file" or "none."
Volume Capacity	Total capacity for all associated volumes
Volume Used	Total of used space for all the associated volumes
Volume Available	Available capacity from all of the associated volumes
Volume Usage	Used data storage at the logical layer. Mouse over the thermometer bar to view the % available.

LUN Capacity

The LUN fields enable drill-down access to additional details so that you can determine the capacity of each configured LUN.

Table 9-29 LUN Capacity

Column name	Description
# LUNs	The number of LUNs that have been created on the array. Click the Used link or the LUN link to go to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Total LUN Capacity	Total storage capacity for all of the associated LUNs
# Allocated LUNs	The number of LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. For NetApp, these are the LUNs associated with an iGroup. See " LUN Utilization Summary " on page 344.
AllocLUN Capacity	The total storage capacity for the allocated LUNs
# UnAllocated LUNs	The number of LUNs that have yet to be allocated
UnAllocated LUN Capacity	The available LUN capacity

Table 9-29 LUN Capacity *(continued)*

Column name	Description
LUN Usage	% of LUN capacity that is currently in use

Note: The Capacity and Used columns are aggregated and then converted to the units selected in the Advanced option of the report from which the Host Utilization Detail was launched. Note that in the conversion to the displayed units, the totals may appear to be slightly inaccurate. This is due to the division or multiplication by 1024.

Array Capacity and Utilization (NetApp Cluster)

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Because NetApp Cluster mode uses a number of unique constructs--such as Aggregates and Volumes--to organize and access its storage, additional data is available through drilldowns.

The report is organized into four sections:

1. See “NetApp Cluster Summary” on page 214.
2. See “Physical Summary” on page 211.
3. See “Aggregate Summary” on page 216.
4. See “Volume Capacity” on page 217.
5. See “LUN Summary” on page 217.
6. See “Performance” on page 218.

NetApp Cluster Summary

This section of the report provides a overview of the storage system, with links to relevant components.

Table 9-30 NetApp Cluster Summary

Column name	Description
Cluster	NetApp Cluster name.
Last Poll Time	Date/Time of the last successful or partially successful collection cycle.

Table 9-30 NetApp Cluster Summary *(continued)*

Column name	Description
UUID	Product designator of NetApp instance (universally unique identifier).
SN	Serial number of the storage array.
ONTAP Ver	Version of ONTAP as of the last data collection.
Location	User designated. Free-form field describing the relative location of the cluster in the user's environment.
Peer	Peer cluster name.
Contact	Primary contact information for the cluster.
# Nodes	Number of configured nodes in the cluster. See "NetApp Cluster-Mode Node Summary" on page 385.
# Vservers	Number of configured Vservers. Drill down to the Vserver Summary report.
# FC Ports	Sum of all member nodes' FC ports. Drill down to the Array Port Utilization reports for these ports.
# iSCSI Ports	Sum of all member nodes' iSCSI ports. Drill down to the NetApp LUN Utilization Summary.
# CIFS Shares	Sum of all member nodes' CIFS shares. Drill down to the NetApp CIFS report.
# NFS Exports	Sum of all member nodes' NFS exports. Drill down to the NetApp NFS report.
# Disks	Sum of all member nodes' disks. Drill down to the NetApp Disk Summary report for these disks.
# Hosts	Total number of connected hosts via all configured nodes. Drill down to the Host Capacity and Utilization report for these hosts.
Host Usage	Amount of logical capacity associated with host connections. Drill down to LUN Utilization Summary for these LUNs.
% Used	Percentage of used capacity associated with LUNs allocated to hosts.

Physical Summary

This section of information summarizes capacity and usage. The percentages, in particular, capture the essence of your current utilization status.

Table 9-31 Physical Summary

Column name	Description
Internal Capacity	Sum of all disk drive capacities associated with individual member nodes.
External Capacity	For V-Series only, sum of all the usable capacities of volumes served up by external arrays associated with the V-Series array.
Virtualized Capacity	0, Cluster-mode does not virtualize capacity.
Allocated	Sum of all member nodes' configured RAID Group capacities in the cluster.
Available	Represents the physical capacity of the Cluster that potentially could be provisioned in to RAID Groups. Could include customer designated spares. Calculated by subtracting Allocated capacity from Internal Capacity.
% Available	Percentage of Available Capacity relative to Internal Capacity (not External).
Raw Usage	Progress bar representing provisioned RAID Groups (including RAID overhead) relative to Internal Capacity.

Aggregate Summary

This section provides an overview of the aggregate and volume capacity and utilization.

Table 9-32 Aggregate Summary

Column name	Description
# Aggregates	Number of member nodes' aggregates. Drill down to NetApp Aggregate Summary for these aggregates.
Aggr Capacity	The aggregate's total size.
Aggr Used	Amount of the aggregate's capacity in use.
Aggr Available	Available storage in the aggregate.
Aggr Usage	Used data storage at the physical layer. Mouse over the thermometer bar to view the % available.

Volume Capacity

This section provides an overview of the volume capacity and utilization.

Table 9-33 Volume Capacity

Column name	Description
# of Total Vol(ume)s	Sum of all member nodes' volumes. Drill down to NetApp Volume Summary for these volumes.
# FlexVols	Sum of all member nodes' thinly provisioned volumes (FlexVols). Drill down to NetApp Volume Summary for these volumes.
# Inf Vol(ume)s	Sum of all member nodes' infinite volumes. Drill down to NetApp Volume Summary for these volumes.
# Thin Vol(ume)s	Sum of all member nodes' volumes with Space Guarantee of none. Drill down to the NetApp Volume Summary for these volumes.
Vol(ume) Capacity	Used data storage at the physical layer. Mouse over the thermometer bar to view the percent available.
Vol(ume) Used	Total of used space for all the associated volumes.
Vol(ume) Available	Available capacity from all of the associated volumes.
Vol(ume) Usage	Used data storage at the logical layer. Mouse over the thermometer bar to view the percent available.

LUN Summary

The LUN fields enable drill-down access to additional details so that you can determine the capacity of each configured LUN.

Table 9-34 LUN Summary

Column name	Description
# LUNs	Sum of member nodes' FC and iSCSI LUNs. Drill down to LUN Utilization Summary for these LUNs.
Total LUN Capacity	Total storage capacity for all of the associated LUNs.
# Allocated LUNs	The number of LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. Drill down to LUN Utilization Summary for these LUNs.

Table 9-34 LUN Summary *(continued)*

Column name	Description
Alloc LUN Capacity	The total storage capacity for the allocated LUNs.
# UnAllocated LUNs	The number of LUNs that have yet to be allocated.
UnAllocated LUN Capacity	The available LUN capacity.
LUN Usage	% of LUN capacity that is currently in use.

Performance

Table 9-35 Performance

Column name	Description
Avg Processor Busy	Average processor utilization across all processors in the system from last data collection cycle.
CPU Busy	Percentage of time one or more processors is busy in the system from last data collection cycle.
Sys Read Resp Time (ms)	System Read Response Time: Average response time for the Cluster in aggregate as of the last data collection poll. Includes all active Node read operations.
Sys Write Resp Time (ms)	System Write Response Time: Average response time for the Cluster in aggregate as of the last data collection poll. Includes all active Node write operations.
Read OPS	Read operations per second from the last data collection cycle.
Write OPS	Write operations per second from last data collection cycle.
CIFS OPS	CIFS operations per second from last data collection cycle.
NFS OPS	NFS operations per second from last data collection cycle.
FCP OPS	Fibre channel protocol (FCP) operations per second from the last data collection cycle.
iSCSI OPS	iSCSI operations per second from the last data collection cycle.

Array Capacity & Utilization - (Hitachi Vantara All-Flash and Hybrid Flash Storage)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The Array Capacity & Utilization report provides an overview of storage array capacity and utilization. Use this report to identify storage arrays that are under utilized or highly utilized.

The Array Capacity & Utilization report represents Hitachi Vantara All-Flash and Hybrid Flash Storage in the following manner:

Table 9-36 Array Capacity & Utilization (Hitachi Vantara All-Flash and Hybrid Flash Storage view)

Column name	Description
Array/Cluster Summary	
Array Name	The name that you assigned to the array. Available values include all supported storage devices. For a list of supported devices, see the Certified Configurations Guide.
Serial #:	Serial number of the cluster.
IP Address	IP address of the flash storage.
Model	The specific model of the storage array.
Vendor	The name of the vendor that distributes the storage array.
Product	The name of the product, typically shown as a series number. This field provides an indicator of capacity specifications, such as LUNs, hosts per array, and drives per array.
Last Poll Time	The date and time of the last successful data collection cycle.
Firmware Version	Current firmware version of the array.
# FC Ports	The number of FC ports in the array. Click the value to view the Array Port Utilization.

Table 9-36 Array Capacity & Utilization (Hitachi Vantara All-Flash and Hybrid Flash Storage view) *(continued)*

Column name	Description
# iSCSI Ports	Number of iSCSI ports. Click the link to drill down to the Array iSCSI Port Utilization report.
# Disks	The number of disks in the array. Click the value to view the Disk Summary report.
# Hosts	This is the total number of unique hosts that have LUNs assigned by this storage array, and the hosts have been successfully queried using one or more host resource policies. If zero, host resources data has not been collected. See " Host Capacity & Utilization " on page 478.
Cache	Cache size of the array. Values are stored as TiB in the database and rendered according to your user profile preferences.
Host Usage	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Physical Summary	
Raw Capacity	For HDS arrays: Sum of capacity of all PDEVs from all array groups. When calculating raw capacity for HDS arrays, the data collector makes requests to the Device Manager for physical devices (PDEVs) for all Array Groups in a storage array. The data collector assumes that all PDEVs are part of an array group, so PDEVs that are marked as not included in an array group (-1) are included in the capacity calculation. That is, no disks are treated as spares.
Allocated	Total of all of disk capacities (physical storage).
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report.
% Available	Expressed as a percentage, total capacity of the LUNs in the array that have not been assigned to a host and that are available for provisioning. This percentage is a better capacity indicator than the number shown as Available Size.
# Array Groups	Number of Array Groups. Links to List Array Groups report.
Pool Summary	

Table 9-36 Array Capacity & Utilization (Hitachi Vantara All-Flash and Hybrid Flash Storage view) *(continued)*

Column name	Description
# Pools	The number of dynamic provisioning pools links to the Dynamic Provisioning Pool Summary
Total Capacity	The total capacity for the volumes/LUNs.
Used Capacity	Amount of internal RAID group capacity allocated to LUNs.
Available Capacity	Amount of available capacity in RAID groups from which to allocate LUNs.
# Resv Vols/LUNs	The number of LUNs that are reserved on the array. Click the link to view the LUN Utilization Summary.
# Resv Vols/LUN Capacity	Capacity of the reserved LUNs that cannot be assigned to the hosts. For example, dynamic pool capacity.
Volume/LUN Summary	
# Vols/LUNs	The number of LUNs that have been created on the array. Click the link to view the LUN Utilization Summary.
Total Vol/LUN Capacity	The total capacity for the volumes/LUNs.
# Allocated Vols/LUNs	The number of volumes/LUNs on the array that have been mapped to a host. These LUNs are typically configured as volumes on the host and dedicated to an application. Click the link to view the LUN Summary.
Alloc Vol/LUN Capacity	The total storage capacity for the allocated LUNs.
# UnAllocated Vols/LUNs	The number of LUNs that have yet to be allocated. Click the link to view the LUN Summary.
UnAlloc Vol/LUN Capacity	The available LUN capacity.

NetApp Storage System Detail

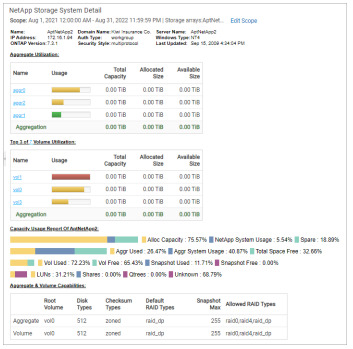
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Click the **Array Name**link.



Current Configuration

Table 9-37 Current Configuration

Column name	Description
Name	Name of the NetApp storage
IP Address	IP address of the storage unit
ONTAP Version	Version of the NetApp ONTAP operating system
Domain Name	The DNS domain name
Auth Type	Indicates the type of authentication: ad - Active Directory, nt4 - Windows NT4, workgroup - Workgroup, passwd - Password file, NIS, or LDAP
Security Style	NTFS or Multiprotocol
Server Name	Server name
Windows Type	Windows software version

Table 9-37 Current Configuration *(continued)*

Column name	Description
Last Poll Time	Last time the portal database was updated with data collected from the NetApp system

Cluster Status

Table 9-38 Cluster Status

Column name	Description
Partner	The associated storage filer in the cluster
Time Master Status	Primary or secondary server within the cluster
State	Cluster State - Connected/Disconnected
Is Enabled	Y or N to indicate enabled state
Is Interconnect Up	Y or N to indicate connectivity

Aggregate Utilization

Aggregate Utilization

The thermometers and capacity information provide a sub-set of the information that is available in the NetApp Aggregate Summary.

See " NetApp Aggregate Summary " on page 364.

Drill down on the "Top 3 of n Aggregate Utilization" to see additional aggregates: Name, Usage, Total Capacity, Allocated Size, and Available Size.

Volume Utilization

Volume Utilization

The thermometers and capacity information provide a sub-set of the information that is available in the NetApp Snapshot Summary.

See " NetApp Snapshot Summary " on page 392.

Drill down on the "Top 3 of n Volume Utilization" to see additional volumes: Name, Usage, Total Capacity, Allocated Size, and Available Size.

Capacity Usage Bar Charts

Capacity Usage Bar Charts	A cascade of bar charts represents usage by: Array - Displays total capacity, the amount used by system overhead, and spare capacity. Aggregate - Of the total capacity, what is used by the aggregate. Volume - Volume usage, including specific snapshot details, and LUN/Shares/QTrees

Aggregate & Volume Capabilities

Table 9-39 Aggregate & Capabilities

Column name	Description
Root Volume	The name of the current root volume of the NetApp storage array
Disk Types	Disks supported by this NetApp array. Possible values include 512 (if all disks are 512 bps), 520 (if all disks are 520 bps), mixed (if disks are a mix of 512 and 520 bps), none (if no disks are 512 or 520 bps).
Checksum Types	Possible values: zoned (fixed VBN), block (block appended), or mixed
Default RAID Types	Possible values include raid0, raid4, raid_dp.
Snapshot Max	Maximum number of snapshots available per aggregate
Allowed RAID Types	List of allowed RAID types

Array iSCSI Port Utilization

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array iSCSI Port Utilization

Or, generate the Array Capacity & Utilization report and click the **# iSCSI Ports**link.

For NetApp Cluster-Mode storage,

See "NetApp Cluster-Mode Array iSCSI Port Utilization" on page 227.

Table 9-40 Array iSCSI Port Utilization

Column name	Description
Storage Array /Cluster	The name of the storage system.
Array Port	The array's port number.
WWN	World Wide Name. EMC VNX (CLARiiON): Host Bus Adapter (HBA) 32-byte unique ID number. EMC Symmetrix: World Wide Name of either HBA port or Array port. NetApp, HP 3PAR: World Wide Name.
Host	Host name associated with the array port.
Host Ports	Comma-separated string of the WWN of all the HBA ports associated with the array port.
Mount Point	Mount point path of the file system that is utilizing the port.
File System	File system name.

Table 9-40 Array iSCSI Port Utilization *(continued)*

Column name	Description
File System Type	File system type. When you know the file system type, you can determine the host's storage requirements and limitations. Examples of file system values: NTFS - Windows NT file system ext2, ext3 - Extended file systems 2 and 3 ufs - Unix file system vxfs - Veritas file system FAT - File Allocation Table hfs - Hierarchical file system (Apple) ctfs - Cooperative temporary file system devfs - Device file system hsfs - High Sierra file system jfs - Journaled file system jfs2 - Journaled file system 2 nfs - Network file system objfs - File system object sfs - Secure file system gpfs - General Parallel file system (IBM) proc - a virtual file system procfs - Process file system tmpfs - a virtual memory file system smbfs - Server Message Block file system Null - empty NA- not applicable
Capacity	Total capacity of the file system. Values are stored as KiB in the database and rendered according to your user profile preferences
Used	Total used capacity of the file system. The amount of the storage array being used by the file system. Values are stored as KiB in the database and rendered according to your user profile preferences.

NetApp Cluster-Mode Array iSCSI Port Utilization

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization (NetApp Cluster-Mode)

Click a **# iSCSI Ports**link.

Table 9-41 NetApp Cluster-Mode Array iSCSI Port Utilization

Column name	Description
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Array Port	The server name.
Host	The name of the host to display in reports.
Port WWM	WWN of the array port. Comma separated string of the WWN of all HBA ports associated with the array port.
Mount Point	The mount point of the file system that is utilizing the port.
File System	File system name.

Table 9-41 NetApp Cluster-Mode Array iSCSI Port Utilization *(continued)*

Column name	Description
File System Type	File System type. When you know the file system type, you can determine the host's storage requirements and limitations. Examples of file system values: NTFS - Windows NT file system ext2, ext3 - Extended file systems 2 and 3 ufs - Unix file system vxfs - Veritas file system FAT - File Allocation Table hfs - Hierarchical file system (Apple) ctfs - Cooperative temporary file system devfs - Device file system hsfs - High Sierra file system jfs - Journaled file system jfs2 - Journaled file system 2 nfs - Network file system objfs - File system object sfs - Secure file system gpfs - General Parallel file system (IBM) proc - a virtual file system procfs - Process file system tmpfs - a virtual memory file system smbfs - Server Message Block file system Null - empty NA- not applicable
Capacity	Total capacity of the file system. Values are stored as KiB in the database and rendered according to your user profile preferences
Used Capacity	Total used capacity of the file system. The amount of the storage array being used by the file system. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-41 NetApp Cluster-Mode Array iSCSI Port Utilization *(continued)*

Column name	Description
Vserver	Name of the Vserver. Drill down to NetApp Cluster-Mode Vserver Summary.
State	Current state of this connection. Possible values: New_Connection, Waiting_Tpgtag_Assignment, Login_Waiting_Req, Login_Req_Rcvd, Login_Waiting_Auth, Login_OK_New_Session_Requested, Login_New_Session_Waiting_Reinstatement, Login_OK_New_Conn_Requested, Login_New_Conn_Waiting_Reinstatement, Login_Send_Final_Resp, Full_Feature_Phase, Shutdown_Start, Shutdown_Waiting_Sockio_Shutdown, Shutdown_Sockio_Shutdown_Done, Shutdown_Waiting_ImmDeliv_FFPCmds_Done, Shutdown_ImmDeliv_FFPCmds_Done, Shutdown_Recovery_Waiting_FFPCmds_Ready, Shutdown_Recovery_Waiting_Logout_Rcvd, Shutdown_Recovery_Logout_Rcvd, Shutdown_Recovery_Waiting_FFPCmds_Reassigned, Shutdown_Terminate_Abort_Seq_FFPCmds, Shutdown_Terminate_Waiting_Seq_FFPCmds_Done, Shutdown_Terminate_Seq_FFPCmds_Done.

Array Capacity and Utilization (OpenStack Swift)

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

The report is organized into these sections:

1. See “Cluster Summary” on page 229.
2. See “Physical Summary” on page 230.
3. See “Logical Summary” on page 230.
4. See “Performance” on page 231.

Cluster Summary

This section of the report provides a overview of the storage system, with links to relevant components.

Table 9-42 Cluster Summary

Column name	Description
Cluster	Name of the OpenStack Swift cluster.
Last Poll Time	Date and time of the last successful data collection cycle.
# of Proxy Servers	Number of front-end proxy servers in the cluster. Drill down to the Host Capacity & Utilization report to view the details for these servers.
# of Storage Nodes	Number of storage nodes in the cluster. Drill down to the OpenStack Swift Storage Node Summary report.
# of Devices	Number of devices across all nodes in the cluster. Drill down to the OpenStack Swift Device Summary report.
# of Zones	Number of zones in the cluster. Drill down to the OpenStack Swift Storage Utilization by Zone report.

Physical Summary

Raw Capacity — Sum of all storage nodes' raw capacity (# of drives * drive sizes). This value includes overhead.

Logical Summary

Table 9-43 Logical Summary

Column name	Description
Capacity	Sum of the storage nodes' allocated capacity (logical)
Used	Sum of all storage nodes' used allocated capacity (logical). This is calculated as (# of objects * object sizes).
Usage	This progress bar indicates usage (red, yellow, or green) based on Used/Capacity.

Table 9-43 Logical Summary *(continued)*

Column name	Description
# of Accounts	Number of accounts defined for the cluster. Drill down to the OpenStack Swift Storage Utilization Account report.
# of Containers	Number of containers defined for the cluster. Drill down to the OpenStack Container Summary report.

Performance

Table 9-44 Performance

Column name	Description
Avg IO/sec	Average I/Os per second, for all storage nodes for the last collection cycle.
Avg Resp Time	Average operation response time for all storage nodes for the last 24 hours. This is the weighted average of all devices, based on I/O per second for the same period.
CPU Usage	This progress bar indicates the CPU utilization of all processors in all storage nodes.

Disk Summary

This report lists a summary of disk device details including capacity, vendor, serial number, model and firmware version. Use this report to determine how disks within an array are being used. This offers the physical view of storage capacity.

Table 9-45 Disk Summary

Column name	Description
Device Name	Name of the physical device. Hitachi Data Systems: Name assigned by the user to the drive. EMC VNX (CLARiiON): This is a combination of the bus number + the enclosure number + the disk number. Example: Bus 0 Enclosure 0 Disk 0. EMC Symmetrix: Internally generated name, combining the Symmetrix ID, Director, Interface, and Target ID of the physical disk. NetApp, NetApp Cluster-Mode: Name of the disk. HP 3PAR: Name of the physical device. Dell Compellent: The name of the disk, as designated by the vendor. NetApp Cluster-Mode: Example: 0a.25. Omitted if excluded by desired-attributes.
Array Name	Name of the storage array. Hitachi Data Systems, NetApp, HP EVA, Dell Compellent, Hitachi NAS, IBM XIV: Name of the storage array. EMC Symmetrix: Name of the storage array. This name is derived from the Symmetrix ID. EMC VNX (Celerra): Worldwide Celerra ID. NetApp Cluster-Mode: The name of the cluster.
MDisk Name	For IBM SVC only. The name of the MDisk of which the drive is a member.

Table 9-45 Disk Summary *(continued)*

Column name	Description
Serial #	Disk serial number. Hitachi Data Systems, EMC VNX (CLARiiON), EMC Symmetrix: Serial number of the disk. Example: 3LN3216B. NetApp: Disk serial number. NetApp Cluster-Mode: Disk serial number. Maximum length of 129 characters. HP 3PAR: Serial number of the disk. Dell Compellent: The serial number of the disk. NetApp Cluster-Mode: Omitted if excluded by desired-attributes.
Firmware Version	Firmware version. Hitachi Data Systems: Microcode version of the drive's firmware. NetApp, HP EVA, NetApp Cluster-Mode: Firmware revision of the disk. HP 3PAR: Firmware version. Dell Compellent: The firmware version of the disk. Hitachi Data Systems: Not yet available for 9900V, USP, USP_V and VSP arrays. NetApp, HP EVA: The format of the firmware revision depends on the type of disk and its vendor. NetApp Cluster-Mode: The format of the firmware revision will vary depending on the type of disk and its vendor. Omitted if excluded by desired-attributes.

Table 9-45 Disk Summary *(continued)*

Column name	Description
Vendor	Vendor/Manufacturer of the disk. Hitachi Data Systems: Drive manufacturer. EMC VNX (CLARiiON): Vendor ID of the disk. EMC Symmetrix, NetApp Cluster-Mode: Vendor/Manufacturer of the disk. NetApp: Vendor of the disk. HP 3PAR: Manufacturer of the disk. Dell Compellent: The vendor of the disk. NetApp Cluster-Mode: Omitted if excluded by desired-attributes.
RPM	Drive's speed in revolutions per minute (RPM). Hitachi Data Systems, EMC VNX (CLARiiON), EMC Symmetrix, HP 3PAR: Drive's speed in revolutions per minute (RPM). NetApp: Rotational speed in revolutions per minute. NetApp Cluster-Mode: Rotational speed in revolutions per minute (RPM). Hitachi Data Systems: Available for the following arrays: USP USP_V VSP AMS/WMS/SMS. -1 or 0 = Unknown. NetApp: Examples: 5400 7200 10000 15000 NetApp Cluster-Mode: Possible values: 5400, 7200, 10000, 15000. Omitted if information is unavailable, if RPM does not apply to this disk type, or if excluded by desired-attributes.
Vendor Disk Type	Drive type of the disk.
Disk Type	For IBM SVC Only. Type of disk, for example: ATA, BD, EATA, EFD, FC, LUN, SAS, SATA, SCSI, SSD, XATA, XSAS.
Hot Spare	Hot spare status: No, if not a hot spare; otherwise, the disk position.
Capacity	Physical disk size.Values are stored as KiB in the database and rendered according to your user profile preferences.
Chassis #	For IBM SVC only. Chassis number.

Table 9-45 Disk Summary *(continued)*

Column name	Description
Slot ID	For IBM SVC only. The slot_id of the drive in the enclosure or node. It can be referred to as the drive bay or location. This can be blank.
Model	For IBM SVC only. Drive model of the physical device.
Drive Role	For IBM SVC only. The current role of the drive. Possible values: unused: the drive which is not configured to be used by anything candidate: the drive is available to be configured spare: the drive is configured as a spare, to be used if the arrays fail members member: the drive is configured as a member of an array failed: the drive has been rejected, and is no longer available for use.

RA Summary

This report is specific to EMC Symmetrix reports and can be accessed via the **# RA Ports** field in the Array Capacity & Utilization (EMC Symmetrix View).

See "Array Capacity & Utilization - (EMC Symmetrix View)" on page 183.

This table lists the RDF Director and port details associated with an array. Note that RDF is also known as SRDF (Symmetrix Remote Data Facility) or RA.

Table 9-46 RA Summary

Column name	Description
Storage Array	The name of the EMC Symmetrix storage array.
Director ID	The ID of the associated director.
Director Numeric	Director numeric associated with the EMC Symmetrix port.
Slot #	The slot number of the director.
Status	Status of the Remote Data Facility (RDF) director, such as online.
Director Symbolic	Director symbolic number associated with the EMC Symmetrix port.

Table 9-46 RA Summary *(continued)*

Column name	Description
Type	The type of Remote Data Facility (RDF) director.
Group Name	The name of the Symmetrix Remote Data Facility (RDF) group.
RDF Group #	Remote Data Facility (RDF) device group number for the device. This is a decimal number.
Remote Symmetrix ID	ID of the remote Symmetrix array associated with the port.
Remote RDF Director Symbolic	The remote director symbolic number.
Remote RDF Port ID	The identifier for the remote Remote Data Facility (RDF) port.
Remote RDF Director ID	Remote Resource Description Framework (RDF) Director ID.

Array Executive Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Executive Summary

The Array Executive Summary provides a single view of the physical array capacity and usage in an enterprise. A pie chart reflects the allocated and available storage, while each row in the table summarizes usage by storage vendor.

This report offers a view of physical storage. From an asset management perspective, this Executive Summary provides instant visibility into how much physical storage is owned, how much has been allocated for use, and how much remains for future use. Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.

Note: Since IBM SVC arrays serve as block storage, they are not included in this report of raw, physical capacity.

In this report, the following definitions apply:

- Thick - Capacity that has been allocated by a traditional array group LUN
- Thin - Capacity that has been allocated by a thin-provisioned LUN

Table 9-47 Array Executive Summary

Column name	Description
Vendor Name	The array vendor name; one row for each array vendor in your report scope.
Raw: Capacity	The sum of all raw capacity associated with the arrays (by vendor); For NetApp vFilers and V-Series, this represents only internal raw capacity and does not include vFiler capacity or gateway external capacity. For EMC Symmetrix arrays, this sum excludes hot spares.
Usage: Frames	The number of arrays for this vendor. Click on the link to view the Array Utilization Summary, which displays additional information for each array.
Usage: Total Usable	Amount of usable capacity from physical RAID groups, after RAID overhead; For NetApp vFilers and V-Series, includes external capacity. For HP 3PAR indicates total logical device usable capacity. For EMC Symmetrix arrays, this is the total capacity of the LUNs (RAID capacity).
Usage: Used	(Total usable capacity - available capacity) / total usable capacity
Usage: Available	Available capacity of the array groups in the storage arrays. The capacity of the open LUNs that do not have set paths; storage in disks that are not in a RAID group.

Table 9-47 Array Executive Summary *(continued)*

Column name	Description
Protection: Primary	Primary capacity; in the case of HDS arrays, it is the amount of the open allocated LUN capacity, the LDEVs with a path to a port (excluding the dynamically provisioned storage).
Protection: Replica	The amount of capacity that has been replicated.
Efficiency: Thick	Amount of usable capacity from physical RAID groups; Thick refers to capacity that has been allocated by a traditional array group LUN. For HP 3PAR, this value represents a value that may not be mapped to a virtual volume.
Efficiency: Thin	Amount of capacity used for thin-provisioned storage; Thin refers to capacity that has been allocated by a thin-provisioned LUN.
Efficiency: Overprov	The percentage of the total capacity, both thick and thin, that has been over-provisioned.
Tier 1	The amount of usable capacity from the tier, defined as a Portal capacity chargeback policy with an associated priority and cost per GB. The name of the tier is dynamic, derived from the chargeback policy name.
Tier 2	The amount of usable capacity from the tier, defined as a Portal capacity chargeback policy with an associated priority and cost per GB. The name of the tier is dynamic, derived from the chargeback policy name.
Tier 3	The amount of usable capacity from the tier, defined as a Portal capacity chargeback policy with an associated priority and cost per GB. The name of the tier is dynamic, derived from the chargeback policy name.

List Array Groups

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Click the **# RAID Groups** link.

This report provides a summary of each of the array groups associated with a storage array, with drill downs to additional details.

Table 9-48 List Array Groups

Column name	Description
Group Name	The name of the array group. See “ Array Group Details ” on page 359.
Group #	The number of the array group.
Storage Array	The array name. See “ Array Capacity and Utilization ” on page 164.
Controller Number	Controller Number.
Chassis Number	The chassis to which the group belongs.
RAID Type	RAID type.
Usable Capacity	Total storage capacity of the group.
Used Capacity	Amount of storage that has been carved into LUNs.
Free Usable Capacity	Available storage.
Raw Capacity	Raw Capacity.

HP EVA Physical Disks

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Click the **# Disks** link.

Table 9-49 HP EVA Physical Disks

Column name	Description
Disk Name	The name of the physical disk.
Disk Group	The group to which the disk belongs
Disk Type	The disk type for this disk group; for example Fibre Channel or solid state.
Formatted Capacity	The amount of storage space on the formatted disk.
Occupancy	How much of the disk is in use.
Quorum Disk	Indicates if this disk is a quorum disk resource for Windows clustering.
Manufacturer	The maker of the disk.
Model Number	The disk model number.
Shelf Number	Indicates the disk enclosure in the array.
Disk Bay Number	The bay in which the disk is located.

Array FC Port Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array FC Port Utilization

This report provides a list of the physical/logical port connections between the host's adapter and the storage array. Use this report to troubleshoot connectivity issues and to re-configure when hosts or arrays are taken out of service. In addition, this report provides information you can use to manage load-balancing across array ports.

Table 9-50 Array FC Port Utilization

Column name	Description
Storage Array	The array links to the Array Capacity and Utilization. You can select multiple storage arrays or a single array when you generate this report.
Array Port	The array's port number.
Host	Name of the host that owns the HBA (host bus adaptor) port. See " Host Utilization Detail " on page 481.
Host Ports	List of the host ports. Comma-separated string of the WWN of all the HBA ports associated with the array port.
Mount Point	The mount point of the filesystem that's utilizing the port.
Filesystem	Name of the filesystem.

Table 9-50 Array FC Port Utilization *(continued)*

Column name	Description
Filesystem Type	When you know the filesystem type, you can determine the host's storage requirements and limitations. Examples of filesystem values: NTFS - Windows NT Filesystem ext2, ext3 - Extended Filesystems 2 and 3 ufs - Linux Filesystem vxfs - Veritas Filesystem FAT - File Allocation Table hfs - Hierarchical Filesystem (Apple) ctfs - Cooperative Temporary Filesystem devfs - Device Filesystem hsfs - High Sierra Filesystem jfs - Journaled Filesystem jfs2 - Journaled Filesystem 2 nfs - Network Filesystem objfs - Filesystem Object sfs - Secure Filesystem gpfs - General Parallel Filesystem (IBM) proc - a virtual filesystem procfs - Process Filesystem tmpfs - a virtual memory filesystem smbfs - Server Message Block Filesystem Null - empty NA - not applicable
Capacity	Capacity of the filesystem.
Used	The amount of the storage array being used by the filesystem.

Storage Array — The array links to the Array Capacity and Utilization. You can select multiple storage arrays or a single array when you generate this report.

Array Port	The array's port number.
Host	Name of the host that owns the HBA (host bus adaptor) port. See " Host Utilization Detail " on page 481.
Host Ports	List of the host ports. Comma-separated string of the WWN of all the HBA ports associated with the array port.
Mount Point	The mount point of the filesystem that's utilizing the port.
Filesystem	Name of the filesystem.
Filesystem Type	When you know the filesystem type, you can determine the host's storage requirements and limitations. Examples of filesystem values: NTFS - Windows NT Filesystem ext2, ext3 - Extended Filesystems 2 and 3 ufs - Linux Filesystem vxfs - Veritas Filesystem FAT - File Allocation Table hfs - Hierarchical Filesystem (Apple) ctfs - Cooperative Temporary Filesystem devfs - Device Filesystem hsfs - High Sierra Filesystem jfs - Journaled Filesystem jfs2 - Journaled Filesystem 2 nfs - Network Filesystem objfs - Filesystem Object sfs - Secure Filesystem gpfs - General Parallel Filesystem (IBM) proc - a virtual filesystem procfs - Process Filesystem tmpfs - a virtual memory filesystem smbfs - Server Message Block Filesystem Null - empty NA - not applicable

Capacity	Capacity of the filesystem.
Used	The amount of the storage array being used by the filesystem.

Data Domain Filesystem Compression

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Data Domain Filesystem Compression

Use this report to view a bar chart of the last day's (24 hours) compression statistics, both data sent before compression and also the actual physical space used. These statistics provide insight into how well compression is working during that time period.

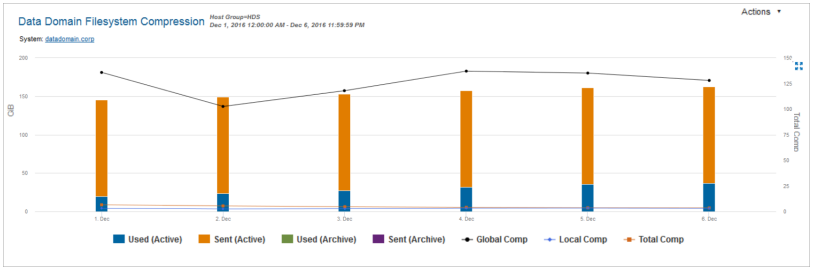


Table 9-51 Data Domain Filesystem Compression

Column name	Description
Used (Active)	Actual amount of space in use for the active tier.
Sent (Active)	Amount of data sent prior to compression is represented by the yellow bar in the graph.
Used (Archive)	Actual amount of space in use for the archive tier.
Sent (Archive)	Amount of data sent for the archive tier is represented by the purple bar in the graph.

Table 9-51 Data Domain Filesystem Compression *(continued)*

Column name	Description
Global Comp	Use this value to determine de-duplication efficiency. Represented by the triangle, the Global Compression factor is derived from the pre-compression (data sent before compression) divided by the size after de-duplication.
Local Comp	The Local Compression factor (file compression) is represented by the circle. This value is derived from the size after de-duplication divided by the space used.
Total Comp	Represented by the square, the Total Compression factor is derived from data sent divided by space used.

Data Domain System Summary - Storage

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Data Domain System Summary - Storage

View a summary of the Data Domain systems.

Use this report to monitor Data Domain system capacities to determine the amount available for data storage. The capacity and used values are derived from Data Domain file system data.

Table 9-52 Data Domain System Summary - Storage

Column name	Description
Array Name	Name of the storage array.
OS Version	A foreign key reference to the domain table.
Last Updated	Last update date of the array.

Table 9-52 Data Domain System Summary - Storage *(continued)*

Column name	Description
Active Capacity	Actual physical space used for compressed data.
Used Active	Amount of active capacity in use.
Archive Capacity	Archive physical disk space available for data. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used Archive	Amount of used archive capacity.
# of MTrees	Number of MTrees associated with the cluster.
# of Disks	Number of physical disks in the array. Click the link to view the Data Domain Disk Summary.
Raw Capacity	Sum of all the capacities of the PDEVs in the array. Values are stored as KiB in the database and rendered according to your user profile preferences.

Data Domain System Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Data Domain System Summary - Storage

This full set of Data Domain System Details is available as a drill down from other Data Domain reports that contain a System link.

Table 9-53 Data Domain System Details

Column name	Description
System	
Internal System ID	Data Domain's internal ID links to the Host Detail report.
System	The Data Domain system name.
Model	The Data Domain system model.
OS Version	The Data Domain operating system version running on this system.
Serial #	Serial number of the Data Domain system.
Filesys Encrypt	Indicates if file system encryption is enabled.
Last Poll Date	Last time the Data Domain system was updated.
Capacity	
Raw Capacity	Total of all disk capacities.
Active Capacity	Amount of capacity that is RAID-configured (used + free).
Used Capacity	Derived from the file system information.
Usage	Mouse over this thermometer to view the active capacity usage %.
Archive Capacity	Amount of archived file system capacity.
Used Archive	Amount of the archived file system capacity that is in use.
Usage	Mouse over the thermometer to view the archive capacity usage %
LSU Summary	
# LSUs	Number of configured logical storage units in this Data Domain system; links to the Storage Unit Detail.

Table 9-53 Data Domain System Details *(continued)*

Column name	Description
# Files	Total number of files in the LSUs.
Compression Factor	De-duplication ratio (original bytes divided by locally compressed).
Original Bytes	Amount of original data before compression and de-duplication.
Globally Compressed	Indicates de-duplication efficiency.
Locally Compressed	File compression.
Metadata	Amount of LSU storage consumed by metadata; this is only collected if SNMP is configured as your data collection method.
MTree Summary	
# MTrees	The number of MTrees associated with this Data Domain system. This number links to the Data Domain MTree Summary report.
Last 24 hrs Pre Compression	Data written before compression within the last 24 hours. Represents the space used and compression achieved for files and directories in the file system.
Last 24 hrs Post Compression	Storage used after compression within the last 24 hours.
Last 24 hrs Compression Factor	The factor used by Data Domain for compression.
VTL Summary	
# VTLs	Number of VTLs available to this Data Domain system.
# Tapes	Number of tapes in the VTLs.
Total Size	Total storage capacity in the VTLs.
Space Used	Amount of VTLs in use.
% Free	% of VTL storage still available.

Table 9-53 Data Domain System Details *(continued)*

Column name	Description
Usage	Mouse over this thermometer to view the usage %.
Average Comp	Average compression factor.
Overall Compression Factor	Overall compression factor is derived from pre-compression size (data written before compression) divided by file system used.

Data Domain MTree Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Data Domain MTree Summary

View Data Domain MTrees pre- and post-compression values for the last 24 hours and the last seven days.

These values are relevant for the last successful data collection cycle.

Table 9-54 Data Domain MTree Summary

Column name	Description
System	Name of the host to display in NetBackup IT Analytics reports.
MTree	Name of the MTree.
Pre Compress	Data written before compression within the last 24 hours. Represents the space used and compression achieved for files and directories in the file system.
Post Compress	Storage used after compression within the last 24 hours.

Table 9-54 Data Domain MTree Summary *(continued)*

Column name	Description
Global CompFactor	Ratio of pre-compression / (size after deduplication) within the last 24 hours.
Local CompFactor	Ratio of (size after deduplication) / post-compression within the last 24 hours.
Pre Compress	Data written before compression within the last 7 days. Values are stored as KiB in the database and rendered according to your user profile preferences.
Post Compress	Storage used after compression within the last 7 days.Values are stored as KiB in the database and rendered according to your user profile preferences.
Global CompFactor	Ratio of pre-compression / (size after deduplication) within the last 7 days.
Local CompFactor	Ratio of (size after deduplication) / post-compression within the last 7 days.

Data Domain Replication History

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Data Domain Replication History

Use this report to determine which destinations are being replicated. The values shown in the report are derived from cumulative statistics since the replication context was created.

Table 9-55 Data Domain Replication History

Column name	Description
System	The Data Domain system

Table 9-55 Data Domain Replication History *(continued)*

Column name	Description
Date	This connection date is relevant only for SNMP data collection; otherwise, the field will be blank
Destination	Indicates the destination that is being replicated
Remaining Pre-Comp	Relevant for directory replication only; the sum of the sizes of the files yet to be replicated for this context; includes the entire logical size of the current file that is in the process of being replicated. If a very large file is currently being replicated, the number may not change for some time as it only changes after the current file replication finishes.
Replicated Pre-Comp	The number of pre-compressed bytes sent by the source. This includes logical bytes associated with the file that is currently being replicated.
Replicated Post-Comp	The actual network data that is sent by the destination to the source
Compression Factor	Calculated as replicated pre-comp divided by replicated post-comp; This is a global compression factor.
Synced As Of Time	The time when the most recently replicated data on the destination was generated on the source. During replication initialization, this value is set to unknown.
Status	Status of the replication
Error	Replication error, if any

Data Domain Snapshot History

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Data Domain Snapshot History

Table 9-56 Data Domain Snapshot History

Column name	Description
System	The Data Domain system
Name	Snapshot name
Pre-Comp	The amount of pre-compression data in this Data Domain snapshot
Created Date	The date and timestamp when the snapshot was created
Retain Until	Indicates when the snapshot expires, either a specific date or Never Expires
Status	Field is either blank or expired Note that an expired snapshot remains available until the next filesystem cleanup.
Expires After (Days)	The time from the current date that the snapshot will expire

EMC Isilon Disks

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon Disks

This report shows Isilon physical disks and associated capacity and other attributes affiliated with an Isilon disk. This report lists Isilon disk details providing Unified Visibility and supporting Capacity Planning and Utilization/Optimization.

Table 9-57 EMC Isilon Disks

Column name	Description
Cluster	Name of the storage array.
Node	The name of the Isilon node.
Chassis #	Chassis number.
Bay #	The physical drive bay number within the node.
Logical #	The disk logical number within the node.
Device	Name of the physical device.
Status	The status of the drive bay or the drive in the bay (if any).
Model	Drive model of the physical device.
Serial #	Disk serial number.
Firmware	Firmware version.
Disk Type	Disk type. Examples: DISK_TYPE, ATA, BD, EATA, EFD, FC, LUN, SAS, SATA, SCSI, SSD, XATA, XSAS.
Size	Physical disk size. Values are stored as KiB in the database and rendered according to your user profile preferences.
OneFS Capacity	Total file system capacity on the drive, including metadata. This may include metadata and other overhead (such as RAID overhead) that do not contribute to total capacity. Values are stored as KiB in the database and rendered according to your user profile preferences.
OneFS Used	Used file system capacity on the drive. Values are stored as KiB in the database and rendered according to your user profile preferences.
OneFS Free	Free file system capacity on the drive. Values are stored as KiB in the database and rendered according to your user profile preferences.

EMC Isilon HDFS Racks

Capacity Manager > Array Capacity & Utilization > EMC Isilon HDFS Racks

This report lists Isilon Hadoop File Systems (HDFS) and associated IP addresses and pools. Isilon HDFS details can be used for Unified Visibility and to support Capacity Planning and Utilization/Optimization.

Table 9-58 EMC Isilon HDFS Racks

Column name	Description
Cluster	Name of the storage array.
HDFS Rack	Name of the HDFS (Hadoop Distributed File Systems) rack.
Zone	Zone name.
Client Addresses	IP addresses and address ranges of clients.
Address Pools	IP pools assigned to the rack.
Last Updated	The date and time the HDFS (Hadoop Distributed File Systems) rack data was last updated in the NetBackup IT Analytics database.

EMC Isilon HDFS Roots

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization

Click the # HDFS Roots link.

This report lists details associated with Isilon Hadoop File Systems (HDFS) root paths in the file system.

Table 9-59 EMC Isilon HDFS Roots

Column name	Description
Cluster	Name of the storage array.
Root Path	HDFS (Hadoop Distributed File Systems) root path in the file system.
Zone	Zone name.
Block Size	HDFS (Hadoop Distributed File Systems) block size (KiB).
Authentication	Authentication method.
# of Threads	Number of threads for HDFS processing.
Log Level	Logging level for HDFS (Hadoop Distributed File Systems).
Is Web HDFS Enabled?	Indicates whether the web is HDFS-enabled.
Version Supported	HDFS (Hadoop Distributed File Systems) version supported.
Usage Without Overhead	Combined used capacity (in KiB) of all paths in the export (not including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage With Overhead	Combined used capacity (in KiB) of all paths in the export (including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
# Files and Directories	Combined count of files and directories of all paths in the NFS export.
Is Usage Exact?	Indicates whether the usage is an exact value or an “at least” value: Yes or No.
Last Updated	The date and time the HDFS (Hadoop Distributed File Systems) root was last updated in the NetBackup IT Analytics database.

EMC Isilon License Detail

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon License Detail

Table 9-60 EMC Isilon License Detail

Column name	Description
Module Name	Optional licensed OneFS software module.
License Status	License Status.
License Expiration	Only relevant for Evaluation licenses.

EMC Isilon NFS Exports

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon NFS Exports

Table 9-61 EMC Isilon NFS Exports

Column name	Description
Cluster	Name of the storage array.
Export ID	The NFS Export ID.
Usage Without Overhead	Combined used capacity (in KiB) of all paths in the export (not including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-61 EMC Isilon NFS Exports *(continued)*

Column name	Description
Usage With Overhead	Combined used capacity (in KiB) of all paths in the export (including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
# Files and Directories	Combined count of files and directories of all paths in the NFS export.
# Paths	Number of NFS export paths in the cluster NFS export.
Description	Description of the NFS export.
Clients	Number of client addresses that have access to this NFS export.
Zone	Zone name. Click the link to drill down to the EMC Isilon Zones report.
Root Clients	Number of client addresses that have root access to this NFS export.
RO Clients	Number of client addresses that have read-only access to this NFS export.
RW Clients	Number of client addresses that have read-write access to this NFS export.
Unresolved	Number of client addresses that have access that is unresolved to this NFS export.
All Dirs	Indicates if mounting subdirectories directly is allowed: Y or N.
Block Size (KB)	Size of block used for the NFS export.
Set Time	Indicates if the time can be set independent of the cluster system time: Y or N.
Commit	Indicates if writes are committed asynchronously: Y or N.
Dir Xfer (KB)	Directory transfer size (Kilobytes).
Encoding	Encoding method (character set).

Table 9-61 EMC Isilon NFS Exports *(continued)*

Column name	Description
UID	Indicates if incoming UNIX UIDs will be looked up locally: Y or N.
Retry	Indicates if failed user-mapping lookups will be retried: Y or N.
All	The identity by which operations by any user will execute.
Root	The identity by which operations by root will execute.
Full	Indicates if full identity mapping resolution will be used for mapped users: Y or N.
Max File Size	Maximum allowed file size on the server (Kilobytes).
RO	Indicates if the NFS export is read-only: Y or N.
Readdir+	Indicates if processing of readdir-plus requests are enabled: Y or N.
Readdir+ Pre	The number of file nodes to prefetch when a readdir-plus request is processed.
32-bit IDs	Indicates if the size of file identifiers returned from readdir is limited to 32-bits: Y or N.
Read Max (KB)	Maximum read transfer size to report to clients.
Read Multi	Suggested multiple read size to report to clients.
Read Xfer (KB)	Preferred read transfer size to report to clients.
Security Type	Security flavor to support for NFS exports.
Setattr Async	Indicates if set-attributes operations are performed asynchronously for NFS exports: Y or N.
Symlinks	Indicates if support is advertised for symlinks: Y or N.

Table 9-61 EMC Isilon NFS Exports *(continued)*

Column name	Description
Time Delta (secs)	Specifies host time granularity (in seconds).
Write DataSync	
Action	Specifies the datasync write method.
Reply	Specifies the datasync reply method.
Write Filesync	
Action	Specifies the filesync write method.
Reply	Specifies the filesync reply method.
Write Unstable	
Action	Specifies the unstable-write method.
Reply	Specifies the unstable-reply method.
Max (KB)	Maximum write transfer size to report to clients
Multi	Suggested multiple write size to report to clients
Size (KB)	Preferred write transfer size to report to clients

EMC Isilon Nodes

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon Nodes

Table 9-62 EMC Isilon Nodes

Column name	Description
Cluster	Name of the storage array.

Table 9-62 EMC Isilon Nodes *(continued)*

Column name	Description
LNN	The logical number of the node.
Device ID	Unique ID assigned when the node joined the cluster.
Name	The name of the Isilon node.
Internal IP Address	Indicates if the interface is internal: Y or N.
External IP Address	Indicates if the interface is internal: Y or N.
Type	The type of Isilon node: Storage, Accelerator, or Unknown.
Health	Overall status/health of the node: OK, Attention, Down, or Invalid.
Read Only	Indicates if the node is read-only: Y or N.
Smart Failed	Indicates if the node is SmartFailed: Y or N.
Serial #	Serial number of the Isilon node.
Capacity	Total file system capacity on the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Total file system capacity free on the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used	Total file system capacity used on the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
# Int Up Interfaces	Number of internal interfaces that are up.
# Int Down Interfaces	Number of internal interfaces that are down.
# Ext Up Interfaces	Number of external interfaces that are up.
# Ext Down Interfaces	Number of external interfaces that are down.

EMC Isilon Quota Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon Quota Summary

This report lists currently established Isilon quotas, including defaults, for all users and folders. It includes capacities for each, if available, with and without overhead. Use this report to validate higher-level roll-up, for chargeback purposes.

Table 9-63 EMC Isilon Quota Summary

Column name	Description
Array Name	Name of the storage array.
Type	Quota type.
User/Group	The user or group name to which the quota applies.
Quota Path	The path to which the quota applies.
Include Snap Usage	Indicates if snapshots are counted against the thresholds.
Hard Threshold	The hard threshold limits disk usage to a size that cannot be exceeded. When an operation attempts to exceed the hard quota, the operation can fail and an alert is logged.
Soft Threshold	The soft threshold allows a limit with a grace period. If the soft threshold is exceeded when the grace period expires, data writes will fail.
Advisory Threshold	An advisory threshold defines an information limit that can be exceeded. Advisory thresholds do not prevent data writes.
Grace Period	The soft threshold grace period (in seconds).
# Files and Directories	The number of files and/or directories under the quota path.

Table 9-63 EMC Isilon Quota Summary *(continued)*

Column name	Description
Usage With Overhead	Combined used capacity (in KiB) of all paths in the export (including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
Enforced	Whether the quota is an enforced quota (true) or an accounting quota (false).
Container	Whether the hard threshold his shown as the capacity of a CIFS share using this quota.
Linked	Whether a user or group quota is linked to the default-user/default-group quota for the path.
Creation Date	Day and time the snapshot was created.
Last Updated	The date and time the Isilon storage system data was last updated in the database.

EMC Isilon SMB Share Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon SMB Share Summary

Table 9-64 EMC Isilon SMB Share Summary

Column name	Description
Cluster	Name of the Isilon cluster.
Name	Display name of the share.
Path	Shared path in the ifs file system.

Table 9-64 EMC Isilon SMB Share Summary *(continued)*

Column name	Description
Usage Without Overhead	Combined used capacity (in KiB) of all paths in the export (not including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage With Overhead	Combined used capacity (in KiB) of all paths in the export (including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
# Files and Directories	Combined count of files and directories of all paths in the NFS export.
Description	Configurable share description.
CSC Policy	Client-side caching policy.
Zone	Zone name. Click the link to drill down to the EMC Isilon Zone Summary report.
Browse	If yes, it makes the share visible in net view and browse lists.
Exec	If yes, allows a user with read-only (RO) access to execute the file.
Strict	If yes, the server checks for and enforces file locks.
Hide	If yes, it hides files that begin with a dot (Unix config files).
Imper	Allow guest access to the share.
Flush	If yes, flush requests are always honored.
Byte Start	Starting point for automatic invalid byte mangling.
Create Perm	Default permissions to apply to sub-directories and files.
Map	Mangle map.
User	Impersonate user name.

Table 9-64 EMC Isilon SMB Share Summary *(continued)*

Column name	Description
Notify	Change notify on the share.
Locks	If yes, allow oplock requests.
RO Del	If yes, allows read-only files to be deleted.
NTFS ACL	If yes, allows ACLs to be stored and edited.

EMC Isilon Snapshot Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC Isilon Snapshot Summary

Table 9-65 EMC Isilon Snapshot Summary

Column name	Description
Cluster	Name of the Isilon cluster.
ID	Snapshot ID.
Name	Name of the snapshot.
Path	Fully qualified path name that was the source of the snapshot.
Locks	If yes, the snapshot is prevented from being deleted.
Schedule	Name of the schedule that created the snapshot.
Alias	Alternate name for the snapshot.
Date/Created	Day and time the snapshot was created.
Date/Expires	Day and time the snapshot is set to expire.

Table 9-65 EMC Isilon Snapshot Summary *(continued)*

Column name	Description
Size	The size of the snapshot (KiB).
Shadow	Amount of data that the snapshot references from shadow stores (cloned files).
% Reserve	Percentage of the snapshot reserve that the snapshot occupies.
% File System	Percentage of the file system that the snapshot occupies.
State	Snapshot state.

EMC Isilon Zones

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization

Click the # Zones link.

Table 9-66 EMC Isilon Zones

Column name	Description
Cluster	Name of the Isilon cluster.
Zone	Zone name.
Isilon Zone ID	Isilon-specified Zone ID.
Zone Path	Zone path.
Cache Size	Cache size of the zone. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-66 EMC Isilon Zones *(continued)*

Column name	Description
Usage Without Overhead	Used capacity (in KiB) of the zone (not including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage With Overhead	Used capacity (in KiB) of the zone (including OneFS overhead). Values are stored as KiB in the database and rendered according to your user profile preferences.
# of Files and Directories	Count of files and directories in the zone.
Is Usage Exact?	Indicates whether the usage is an exact value or an “at least” value: Yes or No.
Last Updated	The date and time the Isilon zone data was last updated in the NetBackup IT Analytics database.

EMC VNX (Celerra) Data Mover Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Data Mover Summary

Table 9-67 EMC VNX (Celerra) Data Mover Summary

Column name	Description
Data Mover	Name of the data mover.
ID	Data mover ID.
Storage Array	Storage array associated with this data mover; links to the EMC VNX (Celerra) Storage System Detail.

Table 9-67 EMC VNX (Celerra) Data Mover Summary *(continued)*

Column name	Description
Type	Data mover type: Primary or Standby.
Slot	Physical slot of the data mover.
State	Actual, online, ok.
Standby	Lists the ID of the Data Mover that is the standby.
Standby Policy	Standby policy type, such as auto.
CPU Utilization	
Average Util %	Average CPU utilization.
Max Util %	Peak CPU utilization.
Memory Utilization	
Average Util %	Average memory utilization.
Max Util %	Peak memory utilization.
Free Memory	Available memory.
Show Chart	Check the boxes and then click Show Chart to produce graphs of the data in the EMC VNX (Celerra) Data Mover Performance report.

EMC VNX (Celerra) Virtual Data Mover Summary

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Data Mover Summary

Click #VDMs.

Table 9-68 EMC VNX (Celerra) Virtual Data Mover Summary

Column name	Description
ID	Number associated with the virtual data mover.

Table 9-68 EMC VNX (Celerra) Virtual Data Mover Summary *(continued)*

Column name	Description
Name	Virtual data mover name.
Server	The data mover to which the virtual data mover is assigned.
Root Filesystem	Root file system that is assigned to the virtual data mover.
I18N Mode	ASCII or Unicode.
Mounted Filesystem	Name of the mounted file system.
Member	Name of the cluster, if a member.
Status Defined	Enabled or Disabled.
State	Loaded ready, loaded active, mounted, temp unloaded, or permanently unloaded.
Interfaces	CIFS or NFS.

EMC VNX (Celerra) File System Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) File System Summary

Table 9-69 EMC VNX (Celerra) File System Summary

Column name	Description
ID	File system ID.
Storage Array	Links to the EMC VNX (Celerra) Storage System Detail report.
Name	File system name.

Table 9-69 EMC VNX (Celerra) File System Summary *(continued)*

Column name	Description
In Use	Indicates if the file system registered in the Data Mover mount table.
Type	File system type such as, uxfs, rawfs, mirrorfs, chpt, mgfs, group, nested.
Volume	Volume on which the file system resides; links to the EMC VNX (Celerra) Volume Detail report.
Pool	Storage pool for the file system; inks to the EMC VNX (Celerra) Storage Pool Detail report.
Allocated	Amount of file system space that is allocated.
Used	Amount of file system space that is in use.
% Util	% of the file system that is being utilized.
Worm	File system (write once read many) state; compliance, enterprise, or blank (disabled).
Worm Clock	Date, time, and time zone of enabled worm filesystems (when the file system is mounted read/write).
Worm Max	Maximum retention for files in the worm filesystem.
Auto Extend	
State	On or blank.
Virtual	Indicates if virtual provisioning is enabled.
HWM	Size threshold to be reached before the filesystem automatically extends.
Max	Maximum retention date for files in the filesystem.
Thin	Yes or blank, to indicate if thin provisioning is enabled.
Deduplication	

Table 9-69 EMC VNX (Celerra) File System Summary *(continued)*

Column name	Description
State	On, off, suspended; when the filesystem is not mounted, the deduplication state is unavailable.
Status	Deduplication-enabled filesystem status: idle, scanning, reduplicating.
Original	Original filesystem size, before deduplication.
Usage	Current space usage for the filesystem.
Files	Number of files scanned.
Deduped	Number of files de duplicated.
Saved	Space saved in the filesystem after deduplication.
Last	Last file system scan time.

EMC VNX (Celerra) SnapSure Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) SnapSure Summary

Management of replication checkpoints, using the SnapSure functionality, provides point-in-time file system images. Use this report to examine related capacity usage.

Table 9-70 EMC VNX SnapSure Summary

Column name	Description
PFS	Production file system; links to the File System Summary report.
Volume	Links to the Volume Summary report.

Table 9-70 EMC VNX SnapSure Summary *(continued)*

Column name	Description
Data Mover	List of data mover servers, with their types noted: RO=Read only; RW=Read/Write.
Storage Array	Links to the Storage System Detail report.
Pool	Storage pool for the checkpoint file system; links to the EMC VNX (Celerra) Storage Pools Summary.
SavVol	
Auto Ext	Auto-extend: Y or blank.
Virt Prov	Virtually provisioned: Y or blank.
Full Mark	Usage point that triggers a SavVol auto-extend, if system space is available.
Volume	Volume where the checkpoint resides; links to the relevant Volume Summary report.
Max Size	Maximum size for auto-extend and virtual provisioning.
Total	SavVol total capacity of the volume.
Used	Amount of SavVol capacity that is in use.
Avail	Available SavVol capacity.
Util %	SavVol Used/Available.
Checkpoint	
Name	Name of the checkpoint.
Total	Size of the PFS when the checkpoint was taken.
Used	Checkpoint amount of storage used.
Avail	Available checkpoint capacity.
Util %	Used/Available capacity.
ID	Lists the checkpoint IDs.

Table 9-70 EMC VNX SnapSure Summary *(continued)*

Column name	Description
In Use	A checkpoint is in use if a file system is registered in the mount table of a Data Mover.
Date	Checkpoint date.

EMC VNX (Celerra) Storage Pools Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Storage Pools Summary

Table 9-71 EMC VNX (Celerra) Storage Pools Summary

Column name	Description
ID	Storage pool ID.
Name	Storage pool name links to the EMC VNX (Celerra) Storage Pool Detail.
Storage Array	Storage Array links to the EMC VNX (Celerra) Storage System Detail.
Description	Storage pool description.
In Use	Indicates if the storage pool is in use: Yes or blank.
Clients	Number of filesystems associated with the pool; links to the EMC VNX (Celerra) File System Summary report.
Thin	Indicates if it is thin-provisioned: Yes or blank.
Slices	Indicates if slices are configured: Yes or blank.

Table 9-71 EMC VNX (Celerra) Storage Pools Summary *(continued)*

Column name	Description
User	Indicates if the storage pool is a User: Yes or blank.
Type	The disk type comprising the storage pool: CLSTD, CLATA, CLSAS, CMSTD, CLEFD, CMEFD or CMATA; See the EMC vendor documentation for a list of disk types.
# Volumes	Number of volumes that comprise the storage pool; links to the Volume Summary report.
Physical	
Used	Mouse over the bar to view the used physical capacity.
Used %	% of the physical capacity in use.
Capacity	Capacity of the physical disks.
Available	Amount of available physical disk space.
Virtual.	
Total	Total virtual capacity.
Allocated	If not thin-provisioned, Virtual=Physical.
Available	Available capacity in the storage pool.
Potential	Potential storage pool capacity, including available space that can be allocated (depending on the associated storage profile).
Subscr	Amount of virtual capacity that has been subscribed (allocated).
Greedy	If Yes, the storage pool tries to create new volumes before consuming capacity storage from existing member volumes.
Dynamic	If Yes, the storage pool can expand or shrink as needed.
# Data Movers	The number of visible data movers in the storage pool; links to the Data Mover Summary report.

Table 9-71 EMC VNX (Celerra) Storage Pools Summary *(continued)*

Column name	Description
Stripe Members	The number of physical drives in the RAID set.
Stripe Size	Stripe size of the current storage pool; a multiple of 8 KB, 0=N/A.

EMC VNX (Celerra) Storage Pools Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization

Click a Storage Pool Name.

Table 9-72 EMC VNX (Celerra) Storage Pools Detail

Column name	Description
Storage Pool Name	Storage pool name.
Description	Storage pool description.
Storage Pool ID	Storage pool ID.
Type	The disk type comprising the storage pool: CLSTD, CLATA, CLSAS, CMSTD, CLEFD, CMEFD or CMATA; See the EMC vendor documentation for a list of disk types.
In Use	Indicates if the storage pool is in use: Yes or blank.

Table 9-72 EMC VNX (Celerra) Storage Pools Detail *(continued)*

Column name	Description
Clients	Number of filesystems associated with the pool; links to the EMC VNX (Celerra) File System Summary report.
Thin	Indicates if it is thin-provisioned: Yes or blank.
Slices	Indicates if slices are configured: Yes or blank.
User	Indicates if it is a User: Yes or blank.
Array Name	Array Name links to the EMC VNX (Celerra) Storage System Detail.
# Data Movers	The number of visible data movers in the storage pool; links to the Data Mover Summary report.
Physical Capacity	Capacity of the physical disks.
Physical Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Physical Available	Amount of available physical disk space.
Virtual Capacity	Total virtual capacity.
Virtual Allocated	Amount of virtual capacity that has been allocated.
Virtual Available	Amount of available virtual capacity.
Subscribed	% of the virtual storage that is subscribed (allocated).
Total Volumes	
Volume ID	Disk volume ID.
Name	Volume name.
Volume Type	Volume type: disk, meta, pool, slice, or stripe.

Table 9-72 EMC VNX (Celerra) Storage Pools Detail *(continued)*

Column name	Description
Total	Total volume capacity.
Allocated	Amount of space on the volume that has been allocated.
Available	Amount of free space on the volume.

EMC VNX (Celerra) Tree Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Tree Summary

Use this report to view quota tree values.

Table 9-73 EMC VNX (Celerra) Tree Summary

Column name	Description
File System	File system name.
Tree ID	Tree ID.
Storage Array	Storage Array links to the Storage System Details report.
Pool	Storage Pool name links to the Celerra Storage Pool report.
Volume	Volume associated with the tree.
Quota Tree Path	The directory path to which tree quotas are applied.
Comment	Tree quota comments.
Capacity	
Used	Used capacity.

Table 9-73 EMC VNX (Celerra) Tree Summary *(continued)*

Column name	Description
Soft	Soft capacity limit, if any.
Hard	Hard capacity limit, if any.
Time Left	Number of days of capacity remaining (aged quota).
Number of Files	
Used	Number of files in the tree.
Soft	Number of files soft limit, if any.
Hard	Number of files hard limit, if any.
Time Left	Number of days left, if any.

EMC VNX (Celerra) Volume Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Volume Summary

User-visible volumes can be many types. When generating this report, the default selection is all volume types: slice, stripe, meta, disk, and pool; however, in the Scope Selector, you can select a sub-set of this list.

Table 9-74 EMC VNX (Celerra) Volume Summary

Column name	Description
ID	Volume ID.
Volume	Volume name links to the EMC VNX (Celerra) Volume Detail report.

Table 9-74 EMC VNX (Celerra) Volume Summary *(continued)*

Column name	Description
Storage Pool	Storage Pool associated with the volume. This field is only populated when the volume is part of a storage pool.
Storage Array	Storage Array links to the EMC VNX (Celerra) Storage System Detail report.
In Use	Indicates if the volume is in use: Yes or blank.
Type	Volume type: disk, meta, pool, slice, or stripe.
Client Type	slice, volume, or file system.
Clients	Links to the relevant summary report that lists the volumes that are clients of this volume.
CKPTs	Number of filesystem checkpoints established; links to the EMC VNX (Celerra) SnapSure Summary.
Vol Set	The associated volume set.
Disks	Links to the Disk Summary report for additional details.
Capacity	
Usage	Mouse over the bar to view the volume's capacity usage.
Total	Total volume capacity.
Allocated	Amount of volume capacity that has been allocated for use.
Available	Available, unused volume capacity.
% Alloc	% of the total volume capacity that is allocated.

EMC VNX (Celerra) Volume Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports

cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC (VNX) Volume Summary

Click a Volume Name.

Table 9-75 EMC VNX (Celerra) Volume Detail

Column name	Description
Volume	Volume name.
Type	Volume type: disk, meta, pool, slice, or stripe.
Alloc Size	Amount of allocated volume capacity.
Clients	Links to the relevant summary report that lists the volumes that are clients of this volume.
# CKPTs	Number of filesystem checkpoints established; links to the EMC VNX (Celerra) SnapSure Summary.
Storage Array	Links to the EMC VNX (Celerra) System Detail report.
Usage	Amount of used volume space.
Available	Amount of available space on the volume.
Used	Amount of used volume space.
Total Member Volumes	
Volume ID	Disk volume ID.
Name	Volume name.
Total	Total volume capacity.
Total Physical Disks	

Table 9-75 EMC VNX (Celerra) Volume Detail *(continued)*

Column name	Description
Disk ID	Disk ID.
Size	Disk capacity.
Used	Disk space already consumed.
Protection	RAID level.

EMC VNX (Celerra) Disk Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Volume Summary

Click the **Disks** link.

Table 9-76 EMC VNX (Celerra) Disk Summary

Column name	Description
Disk ID	Disk identifier.
In Use	Indicates use by a file system or volume.
Size	Total size of the disk.
Used	Amount of disk space currently in use; if a client/file system exists, usage = volume total - volume free space, otherwise the usage = 0.

Table 9-76 EMC VNX (Celerra) Disk Summary *(continued)*

Column name	Description
Usage	Bar provides a graphical representation of used disk space.
Storage ID	Storage system associated with the disk; this is the array group from the back end.
Dev ID	Device associated with the disk.
Protection	RAID level.
Disk Name	Links to the EMC VNX (Celerra) Volume Detail report.
Volume Name	Volume that resides on the disk.
Virt Prov	Indicates if virtual provisioning is enabled; if disabled, this field is blank.
Storage Profile	Storage profile to which the disk belongs.
Mirrored	Indicates if the disk is mirrored.
Storage Array	The storage array for this disk.
Data Movers	List of data movers associated with this disk.

EMC VNX (Celerra) Storage System Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VNX (Celerra) Volume Summary

Click a **Storage Array** link.

Table 9-77 EMC VPLEX (Celerra) Storage System Detail

Column name	Description
Name	Name of the storage system (array)
Domain Name	Domain of the storage system.
CMU	Celerra Management Unit, a unique ID.
Last Poll Time	Last data collection update.
IP Address	Storage array's IP address.
DART Version	Celerra DART OS version (Data Access in Real Time).
Storage Pool Utilization	
Name	Storage pool name.
Usage	Graphical representation of the allocated (consumed) physical capacity.
Total Capacity	Total physical capacity.
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Available	Available physical capacity.
Volume Utilization	
Name	Volume name; To view all volumes, click the above Volume Utilization link to view the Volume Summary report for all the relevant volumes.
Total Capacity	Capacity of the associated volume.

EMC VPLEX Backend Array Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports

cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VPLEX Backend Array Summary

This report lists the backend arrays associated with VPLEX clusters. Link to details in the LUN Utilization Summary and the Array Capacity and Utilization reports.

Table 9-78 EMC VPLEX Backend Array Summary

Column name	Description
Cluster	Name of the VPLEX cluster.
Backend Array	Name of the backend storage array.
Backend Array Ports	List of array ports to which VPLEX is connected. This will contain the Port WWN or iSCSI node name.
Backend Controllers	List of controller names. This field will be empty when the firmware doesn't support the array.
# of Logical Units	Number of logical units exported from the array. This drilldown links to the LUN Utilization Summary. This link may be disabled if collection of backend arrays has not been enabled. If the drilldown is active, there may be a discrepancy in the number of LUNs represented in this report and the corresponding LUN Utilization Summary. This discrepancy can occur when the backend array collection is not as current as the VPLEX collection.

EMC VPLEX Device Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VPLEX Device Summary

This report lists local and global VPLEX devices. View device details such as block size, capacity, and geometry.

Table 9-79 EMC VPLEX Storage View Summary

Column name	Description
Cluster	Name of the storage array.
Device Name	Name of the device.
# of Extents	Number of extents in the device.
# of Devices	Number of complex devices in the device.
Health State	Represents an overview of the health of the extent.
# of Blocks	The number of blocks in the extent.
Block Size	The size of a single block, in kilobytes.
Capacity	The total number of bytes in the extent. This total represents the block size multiplied by the block count.
Geometry	Indicates the geometry or redundancy of this device (similar to RAID level). This does not take into account the geometry of the supporting extents, but indicates how data is stored on the supporting extents.

Table 9-79 EMC VPLEX Storage View Summary *(continued)*

Column name	Description
Visibility	Indicates if this device is being made available to remote clusters. Setting this value changes whether the device is currently available. Possible List of values: local global
Transfer Size	The transfer size, in bytes, during a rebuild.
Virtual Volume	Name of the virtual volume.

EMC VPLEX Extent Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VPLEX Extent Summary

This report lists the Extents that occupy a Storage Volume. Extents associated with a Storage Volume enable capacity allocation, taking advantage of dynamic thin allocation capabilities of the back-end storage array.

Table 9-80 EMC VPLEX Extent Summary

Column name	Description
Cluster	Name of the storage array; drills down to the Array Capacity & Utilization.

Table 9-80 EMC VPLEX Extent Summary *(continued)*

Column name	Description
Storage Volume	The storage array name associated with the storage volume; drills down to the EMC VPLEX Storage Volume Summary.
Extent	Name of the extent.
Device	Name of the VPLEX device; drills down to the EMC VPLEX Device Summary.
IO Status	Indicates the I/O status of the underlying storage-volume. Possible values: alive - I/O is proceeding normally on the extent. dead - The underlying storage volume is marked as hardware-dead. unreachable - The underlying storage volume is unreachable.
Operational Status	The operational status for the extent. This indicates whether the entity is functioning, and, if it is, how well it is functioning. Possible values: degraded ok starting unknown
Health State	Represents an overview of the health of the extent. Indicates the status of the I/O for the extent.
Use	Indicates how the system is currently using the extent.
Capacity	The total number of bytes in the extent. This total is equal to the block size multiplied by the block count.

EMC VPLEX Port Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VPLEX Port Summary

This report lists the status of VPLEX ports associated with a cluster. It includes the number of Storage Views and Virtual Volumes exported by a port.

Table 9-81 EMC VPLEX Port Summary

Column name	Description
Cluster	Name of the VPLEX cluster. Click the link to view the Array Capacity & Utilization report.
Port Name	Name of the target port.
Export Status	This port's export status.
Enabled	Indicates if the port is enabled. The "set" command can be used to enable or disable the port.
# of Views	Number of storage views associated with this port. Click the link to view the EMC VPLEX Storage View Summary.
# of Virtual Volumes	Number of virtual volumes associated with this port. Click the link to view the LUN Utilization Summary.

EMC VPLEX Storage View Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main

report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VPLEX Storage View Summary

This report lists the logical containers consisting of front end ports, registered host initiator ports, and Virtual Volumes. Storage Views determine host access to virtual volumes from VPLEX.

Table 9-82 EMC VPLEX Storage View Summary

Column name	Description
Cluster	Name of the VPLEX cluster.
View Name	Name of the storage view.
Target Ports	Name of the target port. This is a comma-separated list of the target ports.
Initiator Ports	This list of initiator ports is at the storage view level. This is a comma-separated list of initiator ports.
# of Virtual Volumes	Number of virtual volumes in the storage view.

EMC VPLEX Storage Volume Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > EMC VPLEX Storage Volume Summary

This report lists the VPLEX Storage Volumes, which represent the storage that is connected to back-end ports. View the capacity metrics associated with the back-end storage, including the storage volume, number of extents, health and status, and the capacity served to the VPLEX appliance.

Table 9-83 EMC VPLEX Storage Volume Summary

Column name	Description
Cluster	Name of the storage array. Links to the Array Capacity & Utilization report.
Storage Volume	The storage array name associated with the storage volume.
VPD83 ID	Name of the backend LUN.
IO Status	Indicates the I/O status of the storage volume. Possible values: alive = I/O is proceeding normally on the extent. dead = The underlying storage volume is marked as hardware-dead. unreachable = The underlying storage volume is unreachable.
Operational Status	Reason for the operational status, if the status is not OK.
Type	Indicates whether the storage volume has an associated storage volume header, and whether that header is stored on the storage volume or on a metadata volume.

Table 9-83 EMC VPLEX Storage Volume Summary *(continued)*

Column name	Description
Health State	Displays an overview of the extent health. Possible values: Degraded OK Non-recoverable-error Unknown Critical failure
Vendor	The vendor-specific name for the storage volume.
Use	Indicates how the system is currently using the storage volume. Possible values: Use claimed meta-data unclaimed unusable used
Capacity	The total number of bytes in the extent. This total represents the block size multiplied by the block count.
# of Extents	The number of extents in the storage volume.

Object Store Namespace Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Object Store Namespace Summary

List the capacity metrics for namespaces, volumes, or file systems. View total and used capacity, soft quota, and objects in a namespace, along with the amount of metadata that is consuming capacity. Identify objects that have been deleted and are waiting to be shredded. For example, this report is relevant for monitoring Hitachi Content Platform (HCP) capacity.

Table 9-84 Object Store Namespace Summary

Column name	Description
Cluster	Name of the storage array.
Namespace	Namespace name.
Model	Array family.
Tenant/User	Name of the tenant.
Total Capacity	The amount of storage allocated to the namespace. This is the total space available for all data stored in the namespace, including object data, metadata, and redundant data. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used Capacity	The amount of storage currently occupied by all data stored in the namespace, including object data, metadata, and any redundant data. Values are stored as KiB in the database and rendered according to your user profile preferences.
Soft Quota %	The threshold (% of the total capacity) when the system notifies the tenant that free storage space for the namespace is running low.
# of Objects	The total number of objects and versions stored in the namespace.
Shred # of Objects	The number of objects and versions that have been deleted and are waiting to be shredded.
Shred Objects Capacity	The amount of storage occupied by the objects and versions that are waiting to be shredded. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-84 Object Store Namespace Summary *(continued)*

Column name	Description
Meta # of Objects	The number of objects and versions that have one or more annotations.
Capacity Metadata Objects	The amount of storage occupied by custom metadata in all annotations for all objects. Values are stored as KiB in the database and rendered according to your user profile preferences.
Last Updated	The date and time the node data was last updated in the NetBackup IT Analytics database.

Node Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Drilldown on the **Node** link to render this report.

List the capacity metrics for nodes. View the node's type, status, and SAN status. For example, use this report for Hitachi Content Platform (HCP) node capacity metrics.

Table 9-85 Node Summary

Column name	Description
Cluster	Name of the storage array.
Node IP	IP of the node.
Node Type	The type of node.
Hardware Type	Type of hardware used for the node.

Table 9-85 Node Summary *(continued)*

Column name	Description
Total Capacity	Total amount of storage space on all the logical volumes managed by the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used Capacity	Total amount of used space on all the logical volumes managed by the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available Capacity	Available = Total - Used Capacity.
Boot Time	Date and time the node last started.
Node Status	Current status of the node.
SAN Status	Indication of the status of the SAN path. Possible values: 0 = error, 1 = rain, 2 = sanWithoutMultipath, 3 = sanGood, 4 = sanDegradedToOne.
Last Updated	The date and time the node data was last updated in the NetBackup IT Analytics database.

Object Store Tenant Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Object Store Tenant Summary

List capacity metrics for a tenant or user. View object and quota values to understand how space is being allocated and used. For example, this report is relevant for Hitachi Content Platform (HCP) namespace capacity metrics.

Note: Column aggregations are displayed in green on a separate line with the label Aggregation. Rollover the value to display the aggregation type: Avg., Min., Max., or Sum. Column filtering is not supported on reports with aggregated columns.

Table 9-86 Object Store Tenant Summary

Column name	Description
Cluster/Array	Name of the storage array.
Tenant Name	Name of the tenant.
Model	Array family.
# of Namespaces	Number of namespaces enables drilldown to the Namespace Summary.
Description	The description of the tenant.
# of Objects Stored	Number of objects currently stored in all namespaces owned by the tenant.
# of Objects Indexed	Number of indexed objects currently stored in all namespaces owned by the tenant.
Hard Quota	Hard quota for the tenant. This is the total amount of storage available to the tenant.
Soft Quota %	Soft quota for the tenant.
Allocated Quota	Total amount of storage that the tenant has allocated to its namespaces in KiB. Values are stored as KiB in the database and rendered according to your user profile preferences.
Namespace Quota %	Namespace quota for the tenant.
Physical Used	Total amount of storage used by all the namespaces owned by the tenant to store object data and metadata. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-86 Object Store Tenant Summary *(continued)*

Column name	Description
Logical Used	Total size of the stored data and custom metadata before it was added to the namespaces owned by the tenant. Values are stored as KiB in the database and rendered according to your user profile preferences.

Hitachi NAS CIFS Share Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS CIFS Share Summary

This report lists all of the Hitachi NAS (HNAS) Common Internet File System (CIFS) shares exports to hosts. There will be a line item for each share in the scope of the report.

Table 9-87 Hitachi NAS CIFS Share Summary

Column name	Description
Storage Array	The name given to the storage array.
EVS	Name given to the enterprise virtual server (EVS).
Filesystem Name	The instance name of the filesystem. For storage-based snapshots, this is the name assigned when the filesystem was created or renamed.
CIFS Share Name	Name given to the CIFS share.
Comment	Comment given to the share.
Path	Path of the CIFS share.

Table 9-87 Hitachi NAS CIFS Share Summary *(continued)*

Column name	Description
Show Snapshots	Indicates if snapshots are enabled: Y or N.
Follow Links	Indicates if follow symlinks is enabled for the CIFS share: Y or N.
Virus Scan	Indicates if virus scanning is enabled on the filesystem: Y or N.
Cache Options	Cache options for the CIFS share. 1 = Manual local caching for documents, 2 = Automatic local caching for documents, 3 = Automatic local caching for programs, 4 = Local caching disabled.

Hitachi NAS Disk Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Disk Summary

This report lists all of the Hitachi NAS (HNAS) disks discovered as of the last data collection cycle. There will be a line item for each disk in the scope of the report.

Table 9-88 Hitachi NAS Disk Summary

Column name	Description
Storage Array	The name given to the storage array.
Type	Type of HNAS physical disk.
System Drive	Device ID of the SCSI disk.
Capacity	Total capacity of the HNAS physical disk. Values are stored in the database as KiB and rendered according to your user profile preferences.

Hitachi NAS EVS Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS EVS Summary

This report lists all of the Hitachi NAS (HNAS) Enterprise Virtual Servers (EVS) discovered as of the last data collection cycle. There will be a line item for each EVS in the scope of the report. Multiple filers possible. Each Hitachi NAS Cluster can have multiple EVS depending on the amount of front-end connectivity required.

Table 9-89 Hitachi NAS EVS Summary

Column name	Description
Storage Array	The name given to the storage array.
Node ID	Identifier of the Hitachi NAS clustered node.
EVS ID	Unique ID for the enterprise virtual server (EVS) in the HNAS subsystem.
Node Name	Name of the Hitachi NAS clustered node.
EVS Name	Name given to the enterprise virtual server (EVS).
FC Status	Fibre channel port status of the node. One of the following states: OK - All fibre channel interfaces are up on the node. Degraded - One or more fibre channel interfaces are down on the node.
FS Access Status	File system access status of the node. One of the following states: OK - All filesystems that have been made available to this node are accessible.Failures - One or more file systems that had been made available to this node are no longer accessible.
IP Address	IP address for the node.
# FC Ports	Number of fibre channel ports associated with the node.

Hitachi NAS Ethernet Port Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Ethernet Port Summary

This report lists all of the Hitachi NAS (HNAS) Ethernet ports discovered as of the last data collection cycle. There will be a line item for each Ethernet port in the scope of the report.

Table 9-90 Hitachi NAS Ethernet Port Summary

Column name	Description
Storage Array	The name given to the storage array.
IP Address	IP address of the enterprise virtual server (EVS) interface.
EVS	Name given to the enterprise virtual server (EVS).
Port	Port name of the interface aggregation.
Agg	Hitachi NAS interface name.
Status	Status of the port: enabled or disabled.
Cluster Node	Name of the Hitachi NAS clustered node.

Hitachi NAS File System Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS File System Summary

This report lists all of the Hitachi NAS (HNAS) filesystems discovered as of the last data collection cycle. There will be a line item for each filesystem in the scope of the report.

Table 9-91 Hitachi NAS File System Summary

Column name	Description
Storage Array	The name given to the storage array.
EVS	Name given to the enterprise virtual server (EVS).
Type	Provides filesystem characteristics, such as ro (read-only), syslocked (system locked), or WORM. Possible values: Block Size.4KB or 32KB.ro = filesystem has been mounted read-only.syslocked = filesystem has been syslocked.WORM / non-strict WORM = filesystem is WORM or non-strict WORM.File system types = WFS-1 or WFS-2.DSB count = for WFS-2 filesystems, the DSB count is shown.replication target = filesystem is being used as the target for high-speed replication.NDMP recovery target = filesystem is mounted to recover a filesystem image from tape.read resiliency = filesystem is mounted to tolerate disk errors for read-only operations.write resiliency = filesystem is mounted to tolerate disk errors for read and write operations.
# Paths	Number of paths associated with the filesystem.
Capacity	Total capacity of the filesystem. Values are stored in the database as KiB and rendered according to your user profile preferences.
Available	The amount of formatted, available capacity.

Table 9-91 Hitachi NAS File System Summary *(continued)*

Column name	Description
Used	Used capacity of the filesystem. The amount (and percentage) of formatted capacity used by live and snapshot data. If files are queued for deletion, the total size of files awaiting deletion will be shown. If a filesystem is not mounted, the current state of the filesystem will be shown. If the current state shows “Not determined,” the filesystem software has not accessed the underlying storage to determine the state and usage. This can occur when storage is inaccessible or the filesystem has failed a numerous is considered to be too unstable to access. Values are stored as KiB in the database and rendered according to your user profile preferences.
Utilization	Percentage of the filesystem that is used.
Filesystem Name	The instance name of the filesystem. For storage-based snapshots, this is the name assigned when the filesystem was created or renamed.
Confined	The capacity to which the filesystem has been confined. An “-” indicates that the filesystem is not confined.
# Tiers	Number of tiers associated with the filesystem.
Thin Provisioned	Indicates if the filesystem is thinly provisioned.

Hitachi NAS NFS Export Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS NFS Export Summary

This report lists all of the Hitachi NAS (HNAS) Network File System (NFS) exported to hosts as of the last data collection cycle. There will be a line item for each NFS Export in the scope of the report.

Table 9-92 Hitachi NAS NFS Export Summary

Column name	Description
Storage Array	The name given to the storage array.
EVS	Name given to the enterprise virtual server (EVS).
Filesystem	The instance name of the filesystem. For storage-based snapshots, this is the name assigned when the filesystem was created or renamed.
Path	NFS export path, which is the filesystem directory path to the export.
Capacity	Filesystem size. Values are stored in the database as KiB and rendered according to your user profile preferences.
Free Space	Filesystem free space. Values are stored in the database as KiB and rendered according to your user profile preferences.
Thin	Indicates if the filesystem is thinly provisioned: Y or N.
Access Snapshots	Indicates the accessibility and visibility of the snapshot directory. Examples: disabled, hidden, or enabled.
Display Snapshots	Indicates if snapshots are enabled: Y or N.
Read Caching	Status of read caching. Possible values: disable = Read caching is disabled.enable-all = All files are candidates for read caching.enable-cvl = Only cross-volume links are candidates for read caching.
Mounted	Indicates if the filesystem is mounted: Y or N.
Formatted	Indicates if the filesystem is formatted: Y or N.

Table 9-92 Hitachi NAS NFS Export Summary *(continued)*

Column name	Description
Failed	Indicates if the filesystem is in a failed state: Y or N.

Hitachi NAS Snapshot Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Snapshot Summary

This report lists all of the Hitachi NAS (HNAS) snapshots of filesystem as of the last data collection cycle. There will be a line item for each snapshot in the scope of the report relating the snapshot and its capacity to the source file system.

Table 9-93 Hitachi NAS Snapshot Summary

Column name	Description
Storage Array	The name given to the storage array.
EVS Name	Name given to the enterprise virtual server (EVS).
Filesystem Name	The instance name of the filesystem. For storage-based snapshots, this is the name assigned when the filesystem was created or renamed.
Snapshot Name	Name given to the snapshot.
Snapshot Created	Date and time the snapshot was created.
Creation Purpose	Reason for creating the snapshot.
Used Capacity	Space occupied by the snapshot.
Snapshot Status	Status of the snapshot.

Hitachi NAS Storage Pool Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Storage Pool Summary

This report lists all of the Hitachi NAS (HNAS) Storage Pools as of the last data collection cycle. There will be a line item for each Storage Pool in the scope of the report describing the capacities and relating to file systems defined.

Table 9-94 Hitachi NAS Storage Pool Summary

Column name	Description
Storage Array	The name given to the storage array.
Span Instance Name	The name that was assigned to the span when it was created.
Is Healthy	Indicates if the server has I/O with the filesystems. Yes = the server can do I/O to the filesystems on the span. If storage-based mirroring is implemented, a failed secondary does not impact the good I/O health, because a span can support I/O as long as all its primaries are healthy.
# Tiers	Number of tiers associated with the span.
Total Capacity	Total capacity of the span. Values are stored in the database as KiB and rendered according to your user profile preferences.
Used	Used capacity of the span.
Usage	Percentage of the span that is used space, unallocated to filesystems.

Table 9-94 Hitachi NAS Storage Pool Summary *(continued)*

Column name	Description
# Sets	Number of sets associated with the span. See "Hitachi NAS System Drive Summary" on page 306. Note that a single set could have more than one system drive and therefore the drilldown report may return more rows than indicated by the value in this field.
# Filesystems	Number of filesystems associated with the span.
# Mirrored SDs	Number of mirrored system drives. See "Hitachi NAS System Drive Summary" on page 306.
# Chunks	Number of chunks contained in the span. Storage is allocated in chunks (small allocations of each stripe set, of relatively equal size) from the span (storage pool) to filesystems. When a file system is created or expanded, the server always allocates a whole number of chunks. Chunks can be allocated as needed, supporting thin provisioning, or preallocated at file system creation.
Chunk Size	The guideline or approximate size of each chunk. Chunks are close to, but not always exactly the same size. This size controls how much capacity will be allocated to a filesystem at one time.
Auto Expand	Indicates if a span is auto-expandable: Y or N.

Hitachi NAS Storage Server Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Storage Server Summary

This report lists all of the Hitachi NAS (HNAS) Storage Servers configured as of the last data collection cycle. There will be a line item for each Storage Server in the scope of the report. Multiple filers are possible. Each Cluster can have multiple Storage Servers that manage the block storage.

Table 9-95 Hitachi NAS Storage Server Summary

Column name	Description
Storage Array	The name given to the storage array.
Node ID	Identifier of the Hitachi NAS clustered node.
# EVS	Number of Hitachi NAS platform enterprise virtual servers (EVS) associated with the node.
Cluster Node	Name of the Hitachi NAS clustered node.
Enabled	Fibre channel port status of the node. One of the following states: OK - All fibre channel interfaces are up on the node. Degraded - One or more fibre channel interfaces are down on the node.
Status	File system access status of the node. One of the following states: OK - All filesystems that have been made available to this node are accessible. Failures - One or more file systems that had been made available to this node are no longer accessible.
IP	IP address for the node.
# Ports	Number of fibre channel ports associated with the node.

Hitachi NAS System Drive Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS System Drive

This report lists all of the Hitachi NAS (HNAS) System Drives (SD) configured as of the last data collection cycle. There will be a line item for each SD in the scope of the report. Multiple filers are possible. System Drives are special drives required in each Cluster for internal purposes only.

Table 9-96 Hitachi NAS System Drive Summary

Column name	Description
Storage Array	The name given to the storage array.
Label	HNAS system drive number assigned by the server to uniquely identify the system drive.
Status	The system drive's status on the cluster node, as reported by the SCSI driver.
Role	The system drive's mirror role: primary or secondary. If relevant, this will also include the system drive to which this one is mirrored, according to the configuration of the span that includes the system drives. If the system drive has been pegged, the role will be also be shown. Example: Secondary (Scsi: Primary; peg Secondary); configured partner: 3, where Secondary is the effective role, Primary is the role reported by the RAID controller to the servers SCSI driver, and in this case, the system drive's role has been pegged to Secondary.
Access	HNAS system drive access status. Allowed = access to the system drive is allowed, otherwise the value is set to Denied.
Storage Pool	The name that was assigned to the span when it was created.

Table 9-96 Hitachi NAS System Drive Summary *(continued)*

Column name	Description
SD Group	List of all the system drives that are in the same group as this HNAS system drive. This value will be None if the system drive is not in a group.
Capacity	Total capacity of the HNAS physical disk. Values are stored in the database as KiB and rendered according to your user profile preferences.
Manf	Manufacturer of the HNAS physical disk.
Comment/Rank Name	User-assigned comment for the HNAS system drive.
Disk Type	Type of HNAS physical disk.
Disk Size	Total capacity of the HNAS system drive. Values are stored in the database as KiB and rendered according to your user profile preferences.
Superflush	The number of data disks in the system drive, and the RAID5 stripe size. The server uses these details to optimize write operations. Super flush aggregates I/O and is used to improve write performance.
Chunk Size	The number of bytes in an HNAS system drive data block.
Tier	Name of the HNAS tier.
LUN	The logical unit number (LUN) used by the RAID controller to identify a system drive. Within a single RAID rack, every system drive will have a unique LUN. Since a single server can communicate with multiple RAID racks, several system drives can have the same LUN.For Hitachi Data Systems storage, this is the external (host) LUN and the internal LUN is <>.
Site	Site ID that an administrator has assigned to the HNAS system drive.

Table 9-96 Hitachi NAS System Drive Summary *(continued)*

Column name	Description
Serial #	System Drive's serial number, as taken from the SCSI standard inquiry.

Hitachi NAS Tier Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Tier Summary

This report lists all of Hitachi NAS (HNAS) Tiers configured as of the last data collection cycle. There will be a line item for each Tier in the scope of the report. Multiple filers possible. Each Storage Pool can have up to 3 different types of drives configured, each representing a unique Tier.

Table 9-97 Hitachi NAS Tier Summary

Column name	Description
Array Name	The name given to the storage array.
Span Instance	The name that was assigned to the span when it was created.
Tier Name	Name of the HNAS tier.
Total Capacity	Total capacity of the tier. Values are stored in the database as KiB and rendered according to your user profile preferences.
Free Capacity	Free capacity of the span. Values are stored in the database as KiB and rendered according to your user profile preferences.

Hitachi NAS Virtual Volume Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report only report on volumes where Quotas are set. This report is located here:

Capacity Manager > Array Capacity & Utilization > Hitachi NAS Virtual Volume Summary

This report lists all of the Hitachi NAS (HNAS) Virtual Volumes (VV) configured as of the last data collection cycle. There will be a line item for each VV in the scope of the report. Multiple filers possible. Virtual Volumes are allocated from filesystems and appear to clients as a normal filesystem.

Table 9-98 Hitachi NAS Virtual Volume Summary

Column name	Description
Overview	
Storage Array	The name given to the storage array.
Virtual Volume Name	Name of the virtual volume.
EVS Name	Name given to the enterprise virtual server (EVS).
Filesystem Name	The instance name of the filesystem. For storage-based snapshots, this is the name assigned when the filesystem was created or renamed.
Span Instance Name	The name that was assigned to the span when it was created.
Root Path	Path to the root of the virtual volume.
Email List	Email contact list of an existing virtual volume.
Quota Usage	
Used Capacity	Capacity used by virtual volume. Values are stored in the database as KiB and rendered according to your user profile preferences.

Table 9-98 Hitachi NAS Virtual Volume Summary *(continued)*

Column name	Description
Limit	Maximum limit for capacity of virtual volume. For example Usage: 55.03 GB, Limit : 2 TB (Hard). Values are stored as KiB in the database and rendered according to your user profile preferences.
Limit Type	Indicates if the usage limit is a hard or soft value.
Warning Threshold	Warning threshold for capacity used by virtual volume. For example Usage: 55.03 GB, Limit: 2 TB (Hard), Warning: 75% (1.5 TB). Values are stored as KiB in the database and rendered according to your user profile preferences.
Warning Theshold %	Warning threshold percentage for capacity used by virtual volume.
Critical Threshold	Critical threshold for capacity used by virtual volume. For example Usage: 55.03 GB, Limit: 2 TB (Hard), Critical: 1.7 TB. Values are stored as KiB in the database and rendered according to your user profile preferences.
Critical Threshold %	Critical threshold percentage for capacity used by virtual volume.
Reset Threshold	Amount by which usage must fall below a threshold before re-logging its event. For example Usage: 55.03 GB, Limit: 2 TB (Hard), Critical: 1.7 TB. Values are stored as KiB in the database and rendered according to your user profile preferences.
Reset Threshold %	Percentage by which usage must fall below a threshold before re-logging its event.
Quota File Count	
# Files	Number of files associated with the virtual volume.
Limit	Maximum number of files that can be associated with the virtual volume.

Table 9-98 Hitachi NAS Virtual Volume Summary *(continued)*

Column name	Description
Limit Type	Indicates if the file count limit is a hard or soft value.
Warning Threshold	Warning threshold for number of files associated with the virtual volume.
Warning Threshold %	Warning threshold percentage for number of files associated with the virtual volume.
Critical Threshold	Critical threshold for number of files associated with the virtual volume.
Critical Threshold %	Critical threshold percentage for number of files associated with the virtual volume.
Reset Threshold	Amount by which file count must fall below a threshold before re-logging its event.
Reset Threshold %	Percentage by which file count must fall below a threshold before re-logging its event.

Overview of IBM Array Support

NetBackup IT Analytics supports data collection and reporting on the following logical elements.

IBM Array Site Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization> IBM Array Site Summary

Four or eight disk drive modules comprise an Array Site. Then, one or two Array Sites comprise an Array. An Array is an Array Site that has been formatted to a specific RAID format.

Understanding IBM Enterprise Storage Capacity Calculations

- Capacity values for IBM arrays may appear different from what is reported by the IBM CLI or IBM Tools, as IBM uses a decimal gigabyte--10^9 bytes, while NetBackup IT Analytics uses 1024 for a KB when converting to GB.
- For FB open systems hosts, an extent is 1 GB. For CKD IBM zSeries or S/390 systems, an extent is 1 cylinder (.92 GB).

In the Scope Selector, you can narrow the report output by selecting values for the following criteria:

- State: Assigned, Unassigned, Unavailable
- Disk Classes: ENT, NL

Table 9-99 Understanding IBM Enterprise Storage Capacity Calculations

Column name	Description
Array Site	The Array Site ID.
Storage Array	The physical storage system; Drill down to the IBM Array Detail.
DA Pair	Device adapter pair number--the physical I/O enclosure location. If the DA Pair is an even number, it indicates the first DA pair in the enclosure; an odd number indicates the second DA pair in the enclosure.
Capacity	The capacity of all disks belonging to the Array Site.

Table 9-99 Understanding IBM Enterprise Storage Capacity Calculations *(continued)*

Column name	Description
# of Disks	Number of disks associated with this Array Site. See " IBM Disk Summary " on page 323.
Array Name	Links to the IBM Array Summary. See " IBM Array Summary " on page 313.
Disk RPM	The disk revolutions per minute, such as 10000 or 15000.
State	assigned, unassigned, unavailable, initializing. Indicates if the array site is assigned to the array. If the state is unavailable, its capacity is not included in the total array capacity calculation, as shown in the Array Capacity and Utilization report.
Disk Class	Enterprise fibre channel drives (ENT) or Nearline ATA drives (NL).

IBM Array Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > IBM Array Summary

See "Understanding IBM Enterprise Storage Capacity Calculations" on page 312.

In the Scope Selector, you can narrow the report output by selecting values for the following criteria:

- State: Assigned, Unassigned, Unavailable
- RAID Types: 5, 6, 10
- Disk Classes: ENT, NL

Table 9-100 IBM Array Summary

Column name	Description
Array	The Array ID. An Array (also known as a RAID Array) is an Array Site that has been formatted to a specific RAID format.
Storage Array	The physical storage array system
Array State	assigned, unassigned, unavailable. Indicates if the data is assigned to a Rank.
RAID Type	RAID 5, 6, or 10
# of Array Sites	Drill down to the IBM Array Site Summary. An array can have 1 or 2 array sites. See "IBM Array Site Summary " on page 312.
Rank	Drill down to the IBM Extent Pool Summary. See " IBM Extent Pool Summary " on page 321.
DA Pair	Device adapter pair number--the physical I/O enclosure location. If the DA Pair is an even number, it indicates the first DA pair in the enclosure; an odd number indicates the second DA pair in the enclosure
Capacity	The capacity of the Array (also known as the RAID Array)--one or more Array Sites can comprise an array and the capacity shown is the total for all the Array Sites in the array.
Disk Class	Enterprise fibre channel drives (ENT) or Nearline ATA drives (NL)

IBM Array Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

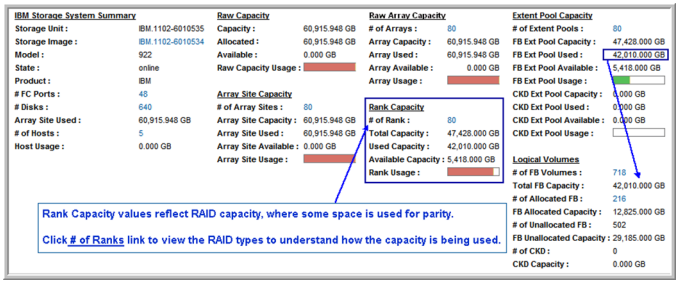
Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > IBM Array Summary

Next, click an Array.

See "Understanding IBM Enterprise Storage Capacity Calculations" on page 312.



IBM Storage System Summary

Table 9-101 IBM Storage System Summary

Column name	Description
Storage Unit	The Storage Unit ID contains the manufacturer, machine type, and serial number.
Storage Image	The Storage Image ID contains the manufacturer, machine type, and serial number.
Model	The storage unit model.
State	State of the storage unit: online, offline, resuming, quiescing, quiesce exception, forced quiescing, fenced, discovery.
Product	Storage Unit vendor.

Table 9-101 IBM Storage System Summary *(continued)*

Column name	Description
# FC Ports	The # of Fibre Channel ports links to the Array Port Utilization report. Note that the # of FC ports may not match the number of rows shown in the drilldown report, as the Port Utilization report lists all the hosts connected to each port.
# Disks	Number of disks for this storage unit
Array Site Used	Amount of array storage used by the Array Site
# of Hosts	Number of hosts mapped to this array. See " Host Capacity & Utilization " on page 478.
Host Usage	Amount of capacity used by hosts

Raw Capacity

Table 9-102 Raw Capacity

Column name	Description
Capacity	Total raw or physical capacity. For IBM Arrays: Total array capacity derived from the physical DDMs (disk drive modules); specifically, the DDM usage in the array site
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use. For IBM Arrays: Amount of capacity that has been allocated is derived from DDM usage for array members and spares

Table 9-102 Raw Capacity *(continued)*

Column name	Description
Available	Total raw capacity of the array that has not been allocated. The size is displayed in the units you selected with the Advanced option when you generated the report. **For IBM Arrays**: Amount of capacity that is free, derived from unassigned and unconfigured DDMs
Raw Capacity Usage	Mouse over the thermometer to view the usage percentage

Array Site Capacity

Table 9-103 Array Site Capacity

Column name	Description
# of Array Sites	The number of array sites to which this array belongs. See "IBM Array Site Summary " on page 312.
Array Site Capacity	Sum of the capacity of all the associated array sites, derived from the capacity of the disks in each array site
Array Site Used	Sum of the array site capacity that is in used--that is, "assigned" to an array
Array Site Available	Array Site storage that has not been allocated--that is, "unassigned" to an array
Array Site Usage	Mouse over the thermometer to view the usage percentage

Raw Array Capacity

Table 9-104 Raw Array Capacity

Column name	Description
# of Arrays	Links to the IBM Array Summary. See " IBM Array Summary " on page 313.

Table 9-104 Raw Array Capacity *(continued)*

Column name	Description
Array Capacity	The capacity of the Array (also known as the RAID Array)--one or more Array Sites can comprise an array and the capacity shown is the total for all the Array Sites in the array. This is raw array capacity.
Array Used	If an array is assigned to a rank, it is “used”
Array Available	Array capacity that is not part of a rank is “available”
Array Usage	Mouse over the thermometer to view the usage percentage.

Rank Capacity

Table 9-105 Rank Capacity

Column name	Description
# of Ranks	The number of ranks links to the IBM Rank Summary
Total Capacity	Sum of the capacity of all ranks
Used Capacity	Capacity assigned to an extent pool
Available Capacity	Capacity not assigned to an extent pool
Rank Usage	Mouse over the thermometer to view the percentage of the ranks that have been allocated.

Extent Pool Capacity

Table 9-106 Extent Pool Capacity

Column name	Description
# of Extent Pools	Links to the IBM Extent Pool Summary; both FB and CKD are included in this number.

Table 9-106 Extent Pool Capacity *(continued)*

Column name	Description
FB Extent Pool Capacity	The sum of the capacity of all the fixed block extents in the extent pool; For FB, an extent is 1 GB.
FB Extent Pool Used	The amount of fixed block extent pool storage in assigned to a logical volume
FB Extent Pool Available	The fixed block extent pool storage that is not assigned to a logical volume
FB Extent Pool Usage	Mouse over the thermometer to view the percentage of the fixed block extent pool that is allocated.
CKD Extent Pool Capacity	The sum of the capacity of all the CKD extents in the extent pool; For CKD, an extent is 1 cylinder (.92 GB).
CKD Extent Pool Used	The amount of the CKD extent pool storage assigned to a logical volume.
CKD Extent Pool Available	The CKD extent pool storage that is not assigned to a logical volume.
CKD Extent Pool Usage	Mouse over the thermometer to view the percentage of the CKD extent pool that is allocated.

Logical Volumes

This section provides the LUN mapping for determining capacity and usage.

Top 3 of n Ranks

Click on the link to view the complete list of Ranks.

Name	Rank name links to the IBM Rank Summary.
Usage	Used Extents in the Rank.
Total Extents	Total Extents in the Rank.

Available	Extents in this Rank that are available for allocation in the Extent Pools; If the available amount equals the total amount, it indicates that the rank is not allocated to any Extent Pools.

Top 3 of n Extent Pools

Click on the link to view the complete list of Extent Pools.

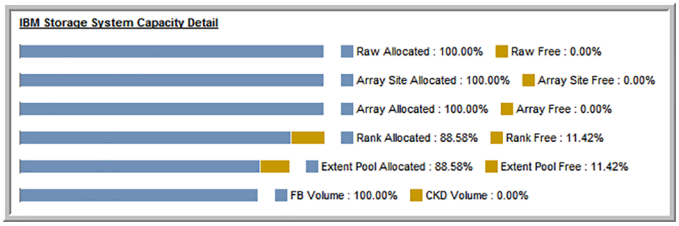
Either 1 or 2 Ranks can comprise an Extent Pool. When it is a 1-to-1 relationship, the Extent Pool capacity values will look identical to the Rank values. Otherwise, they will be double the Rank values. All Ranks in an Extent Pool are of the same RAID type.

Table 9-107 Top 3 of n Extent Pools

Column name	Description
Name	Extent Pool name links to the IBM Extent Pool Summary. See " IBM Extent Pool Summary " on page 321.
Usage	Amount of Extent capacity that is in use.
Total Extents	Total Extent capacity.
Available	Available Extent capacity.

IBM Storage System Capacity Detail

Mouse over the bars to view the capacity totals in GB.



IBM Storage System FB Volumes

A pie chart represents fixed block volumes:

- Ungrouped

- Unmapped
- Mapped

Mouse over the pie segments to view percentages.

Capacity and Object Usage Forecast

Graphs represent forecasted usage for the following objects:

- Array Sites
- Arrays
- Extent Pools
- Ranks

IBM Extent Pool Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > IBM Extent Pool Summary

See “Understanding IBM Enterprise Storage Capacity Calculations” on page 312.

In the Scope Selector, you can narrow the report output by selecting values for the following criteria:

- Storage Types: fb, ckd
- Rank Status: below, exceeded, full

Table 9-108 IBM Extent Pool Summary

Column name	Description
Name	Extent Pool name
Storage System	The physical storage system (storage array) on which the extent pool resides
Rank Group	A rank can only be assigned to one extent pool, but more than one rank can reside in an extent pool.

Table 9-108 IBM Extent Pool Summary *(continued)*

Column name	Description
Storage Type	Fixed block (FB) or Count-Key-Data (CKD).For FB, an extent is 1 GB. For CKD, an extent is 1 cylinder (.92 GB).
Status	Rank Status: below - % of available extents is less than threshold exceeded - % of available extents is greater than threshold full - 0 extents are available
Total Capacity	Sum of the capacity of all extents in this pool
Available Capacity	The amount of space in this extent pool that is currently available; this value is derived from the total capacity minus the allocated capacity
Available Extents	The number of extents that are available to the logical volume
Reserved Extents	The number of extents that are in reserved (this space must be explicitly released before the logical volume can use it)
# of Volumes	The number of logical volumes that have been configured from this extent pool; this is the storage that can be given to a host
# of Ranks	The number of ranks in this extent pool--1 or 2

IBM Rank Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > IBM Rank Summary

An IBM Rank is logical storage, enabling storage provisioning that is not constrained by the size of an array. One or more arrays comprise an IBM Rank, forming logically contiguous storage. Note that IBM often refers to Arrays as RAID Arrays.

See "Understanding IBM Enterprise Storage Capacity Calculations" on page 312.

In the Scope Selector, you can narrow the report output by selecting values for the following criteria:

- State: Normal, Configuring, Unassigned, Reserved, Deconfiguring, Depopulating, Configuration Error, Deconfiguration Error
- RAID Types: 5, 6, 10
- Storage Types: fb, ckd

IBM Disk Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > IBM Disk Summary

In the Scope Selector, you can narrow the report output by selecting values for the following criteria:

- Disk State: Normal, New, Installing, Verifying, Formatting, Initializing, Certifying, Rebuilding, PFSed, Inter failed, Inappropriate, Removed, Failed, Failed - Deferred
- Disk Usage: unassigned, unconfigured, spare required, spare not required, array member
- Disk Classes: ENT, NL

Table 9-109 IBM Disk Summary

Column name	Description
ID	Disk ID
Storage Array	The physical storage array system

Table 9-109 IBM Disk Summary *(continued)*

Column name	Description
Array Site	The array site ID links to the IBM Array Site Summary. See "IBM Array Site Summary " on page 312. If a disk is not in an array site, it is available for other use.
Disk Class	Enterprise fibre channel drives (ENT) or Nearline ATA drives (NL)
Model	The disk model identifier
Disk Usage	How this disk is currently being used: unassigned, unconfigured, spare required, spare not required, array member
Disk Interface	FCAL, SATA, SAS
Disk RPM	The disk revolutions per minute, such as 10000 or 15000
Disk Rate	The minimum disk interface rate for the disks in the array site
State	Normal, New, Installing, Verifying, Formatting, Initializing, Certifying, Rebuilding, PFSed, Inter failed, Inappropriate, Removed, Failed, Failed - Deferred
Capacity	The disk's capacity

IBM SAN Volume Controller Overview

The IBM SAN Volume Controller (SVC) provides virtualized storage--that is, an array of arrays. As such, this technology requires specific terminology to describe physical and logical configurations. Since NetBackup IT Analytics reports on a variety of storage systems, the following table offers a simple glossary of terminology in the context of heterogeneous storage reporting.

Table 9-110 SAN Volume Controller Overview

Column name	Description
Generic Terminology	Corresponding SVC-specific Terminology

Table 9-110 SAN Volume Controller Overview *(continued)*

Column name	Description
Array	Cluster (up to 4 pairs of nodes)
Disk, Hard Disk, or Raw Storage	MDisk (managed disk) or Backend LUN, which is the physical storage carved up into logical units
Storage Pool	MDisk Group
LUN	VDisk (virtual disk); often referred to as front-end storage
Node	One SAN Volume Controller

Figure 6.2 Overview of IBM SVC Architecture

The following diagram illustrates the elements comprising the storage provisioning process, from back-end MDisks to creating RAID disks for VDisk LUNs. The storage in VDisk LUNs is the capacity that is presented to Hosts.

Use this diagram to understand how certain values are displayed in the NetBackup IT Analytics SVC reports:

- Copy Capacity includes both mirrored and un-mirrored LUNs.
- Space-Efficient VDisks (LUNs), also known as thin-provisioned storage, can have VDisk LUN capacity that is greater than the Allocated Disk capacity (displayed as Real Capacity in some reports).

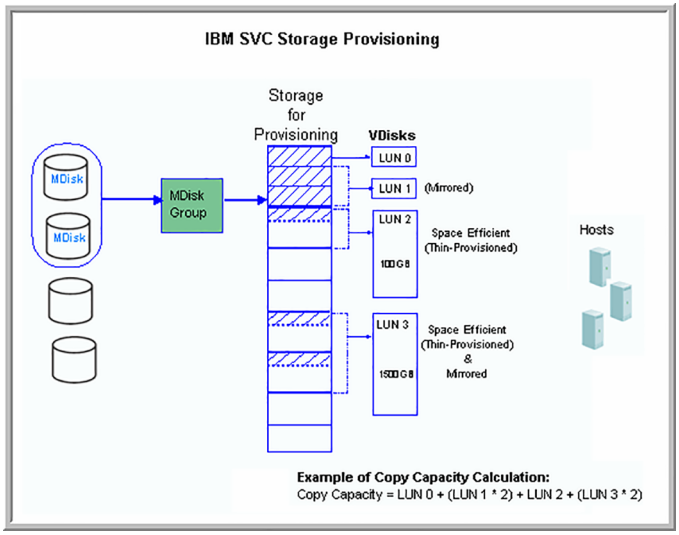


Figure 6.3 Example of IBM SVC Storage Provisioning

Space-efficient VDisks

A Virtual Disk (VDisk) can be configured as space-efficient, which means that it has both virtual and real capacity.

- Virtual Capacity - VDisk storage that is available to hosts
- Real Capacity - Physical storage from an MDisk that is allocated to a VDisk copy

For a fully allocated VDisk (all real capacity), virtual and real capacity are identical; however, in a space-efficient VDisk, virtual capacity can exceed real capacity.

IBM SVC VDisk Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Then, select IBM SVC arrays for the report scope.

Then, click on **# VDisks**or **# VDisks Allocated to Hosts**.

The following view is available when only IBM SVC arrays are selected for the report scope. In addition, several related reports narrow the list of LUNs in the report:

- IBM SVC VDisk Summary - Unallocated: Access this report from the Array Capacity & Utilization report by clicking on **#VDisks Unallocated**.
- IBM SVC VDisk Summary - Allocated: Access this report from the Array Capacity & Utilization report by clicking on **#VDisks Allocated to Hosts**.

To display IBM SVC LUNs, along with LUNs from other storage vendors,

See " LUN Utilization Summary " on page 344.

See "How IBM SVC Capacity Values Are Derived" on page 206.

See "Array Capacity & Utilization (HP 3PAR View)" on page 197.

Table 9-111 IBM SVC VDisk Summary

Column name	Description
LUN #	The LUN number
RAID Type	SVC RAID Type: Striped, Mirrored
Cluster	2-8 nodes comprise a cluster; 2 nodes comprise an I/O group
MDisk Group ID	The name of the MDisk group

Table 9-111 IBM SVC VDisk Summary *(continued)*

Column name	Description
Capacity	Total VDisk (LUN) capacity
Copy Capacity	Copy capacity includes mirrors
Real Capacity	Allocated VDisk capacity; the capacity that is available to hosts
Used Capacity	Amount of the allocated VDisk capacity (real capacity) that is used
Free Capacity	Free = Capacity - Real Capacity
Mirrored	Indicates if this LUN has a mirror
Space Efficient	Indicates if this is a thin-provisioned LUN
IO Group Name	The name of the I/O group to which this VDisk belongs
Estimated Used	From the host's perspective, the amount of storage used
Estimated Available	From the host's perspective, the amount of available storage
Host	Hosts that have storage allocated from this LUN
Host LUN	The name of the host to which storage from this LUN was allocated. Click on a host link to launch the Host Utilization Detail. See " Host Utilization Detail " on page 481.
Mount Point	Mount point of the filesystem
Filesystem	The host's filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts.
Volume Group	Volume group using storage from this LUN

Table 9-111 IBM SVC VDisk Summary *(continued)*

Column name	Description
Host Storage Domain	A logical group of LUNs that can be accessed by multiple hosts through the same physical storage port, with each host accessing its unique LUN. For security purposes, LUN masking is also coupled with this capability. For Hitachi arrays, this is the multi-hosted storage capability.
Array Port WWN	The array port world-wide name

IBM SVC MDisk Groups

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Then, click on **#MDisk Groups**.

Use this report to view raw disk storage for a managed disk group.

See “How IBM SVC Capacity Values Are Derived” on page 206.

See “Array Capacity & Utilization (HP 3PAR View)” on page 197.

Table 9-112 IBM SVC MDisk Group

Column name	Description
Cluster	The physical storage array system

Table 9-112 IBM SVC MDisk Group *(continued)*

Column name	Description
Group ID	The ID of the MDisk group links to the IBM SVC MDisk Group Details. See " IBM SVC MDisk Group Details " on page 330.
Capacity	The total managed MDisk capacity
Real Capacity	The amount of MDisk capacity that has been allocated to VDisk LUNs
Free Space	Free = Capacity - Real
Occupied	Occupied = Real Capacity divided by Capacity
Virtual Capacity	Total VDisk Copy Capacity, which includes both the mirrored and un-mirrored VDisk capacity
Used Capacity	The amount of the allocated capacity that is being used
Virtual Used	Virtual Used = Used Capacity divided by Virtual Capacity
% Over-Allocation	The percentage is derived from the total VDisk Copy Capacity divided by the Space in MDisk Groups

IBM SVC MDisk Group Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

To access this report, in the Navigation pane select:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Click **#MDisk Groups**.

In the IBM SVC MDisk Group Summary, click an **MDisk Group ID**.

Use this report to view the details for block storage in a managed disk group.

See "How IBM SVC Capacity Values Are Derived" on page 206.

See "Array Capacity & Utilization (HP 3PAR View)" on page 197.

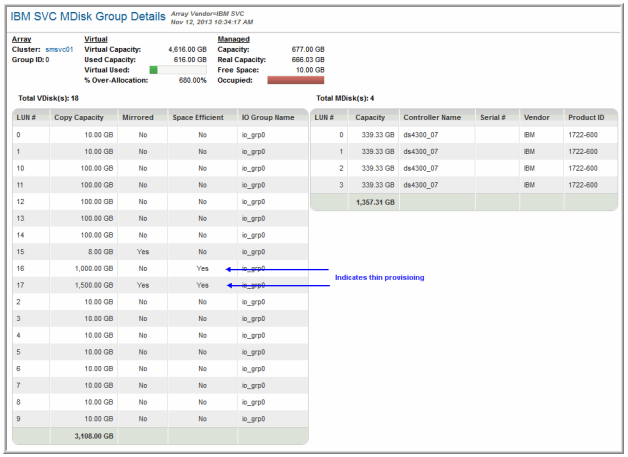


Table 9-113 IBM SVC MDisk Group Details

Column name	Description
Cluster	The physical storage array system
Group ID	The ID of the MDisk group
Virtual Capacity	Total VDisk Copy Capacity, which includes both the mirrored and un-mirrored VDisk capacity
Used Capacity	The amount of the allocated capacity that is being used
Virtual Used	Virtual Used = Used Capacity divided by Virtual Capacity

Table 9-113 IBM SVC MDisk Group Details *(continued)*

Column name	Description
% Over-Allocation	The percentage is derived from the total VDisk Copy Capacity divided by the Space in MDisk Groups
Capacity	The total managed MDisk capacity
Real Capacity	The amount of MDisk capacity that has been allocated to VDisk LUNs
Free Space	Free = Capacity - Real
Total VDisks	
LUN #	LUN number
Capacity	The total managed MDisk capacity
Copy Capacity	Copy capacity includes mirrors
Mirrored	Indicates if the VDisk is mirrored
Space Efficient	Indicates if this is a thin-provisioned LUN
Total MDisks	
LUN #	LUN number
Capacity	Total amount of physical storage for the LUN (MDisk)
Controller Name	The controller name for the MDisk
Serial #	MDisk serial number
Vendor	MDisk vendor
Product ID	MDisk product ID

IBM SVC MDisk Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Next click the **# MDisks** link.

See "How IBM SVC Capacity Values Are Derived" on page 206.

See "Array Capacity & Utilization (HP 3PAR View)" on page 197.

Table 9-114 IBM SVC MDisk Summary

Column name	Description
Controller Name	Name of the I/O controller
LUN #	The VDisk LUN number
Mode	Managed or Unmanaged
Cluster	The cluster (storage array) to which the MDisk belongs
MDisk Group ID	The MDisk group to which the MDisk belongs
Capacity	Capacity of the MDisk
Vendor	The MDisk vendor
Serial #	The MDisk serial number
Product ID	The ID that identifies the physical disk type
Backend Controller Node	Backend Controller Node Worldwide Name
Backend LUN Port	Backend LUN Port Worldwide Name
SSD Node	Indicates if it is a solid-state drive, physically inside a node

How Capacity Values are Represented in XIV Reports

At first glance, it may appear that NetBackup IT Analytics capacity values differ from values shown by other reporting software.

- NetBackup IT Analytics represents capacity values in MB/GB/TB based on: $2^{10}/2^{20}/2^{30}$
- Some other software, such as IBM XIV Storage Systems software, reports capacity values in MB/GB/TB based on: $10^{6}/10^{9}/10^{12}$
- For example, a volume reported by IBM XIV Storage Systems with a size of 17 GB ($17 * 2^{20}$) will appear in NetBackup IT Analytics as: 16 GB ($16 * 10^{9}$)

IBM XIV Consistency Groups Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization> IBM XIV Consistency Groups Summary

See "How Capacity Values are Represented in XIV Reports" on page 334.

Table 9-115 IBM XIV Consistency Groups Summary

Column name	Description
Consistency Group Summary	
Name	Name of the Consistency Group
Storage Array	Array name
Storage Pool	Storage pool associated with this Consistency Group
Mirrored	Yes = storage is mirrored
Sync Type	Mirroring Status, if mirrored

Table 9-115 IBM XIV Consistency Groups Summary *(continued)*

Column name	Description
Current Role	For mirrored volumes and snaps, either Primary or Secondary servers.
Related System	Related storage array
Related Pool	Related storage pool
Related CG	Related Consistency Group

IBM XIV Consistency Group Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization> IBM XIV Consistency Groups Summary

Next, click a link from the Name column.

The Consistency Group ensures that concurrent operations across multiple volumes preserve consistency among the volumes.

See "How Capacity Values are Represented in XIV Reports" on page 334.

Table 9-116 IBM XIV Consistency Group Details

Column name	Description
Consistency Group Details	
Storage Array	Array name
Storage Pool	Storage pool associated with this Consistency Group

Table 9-116 IBM XIV Consistency Group Details *(continued)*

Column name	Description
Mirrored	Yes = storage is mirrored
Sync Type	Mirroring Status, if mirrored
Current Role	For mirrored volumes and snaps, either Primary or Secondary servers.
Related System	Related storage array
Related Pool	Related storage pool
Related CG	Related Consistency Group
List of Volumes/Snapshots	
Name	Snapshot name
Allocated Size	Allocated size in GB
Used	Physical capacity minus Free Hard size
% Util	Used/Hard Limit
Master	Either Volume name for snapshots to volumes or Snapshot Name for snapshots of snaps
Locking Status	If the field is empty, it means it is Not Locked
Serial #	Snapshot serial number
Host	Name of the host that is using the storage
Host LUN	The LUN to which the host is mapped
Mount Point	The host's filesystem mount point for this LUN. It is possible to have zero mount points and filesystems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a filesystem and mounting it.
Filesystem	The host's filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts.
Array Port WWN	The unique identifier, World Wide Name, of the array port address.

XIV Physical Disks

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization

Next, click on a # Disks link for an IBM XIV array.

See "How Capacity Values are Represented in XIV Reports" on page 334.

Table 9-117 XIV Physical Disks

Column name	Description
Component ID	Disk ID
Capacity	Physical capacity of the disk; a total is shown at the bottom of the table
Model	Disk model #
Serial #	Disk serial #
Size	Disk Size; a total is shown at the bottom of the table
Firmware	Firmware Version

XIV Snapshot Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > LUN Utilization Summary

Next, click a link from the Snaps column.

See "How Capacity Values are Represented in XIV Reports" on page 334.

Table 9-118 XIV Snapshot Details

Column name	Description
XIV Snapshot Detail	
Storage Array	Array for the snapshot
Master Volume	Master volume for the snapshot
List of Snapshots	
Name	Snapshot name
Snap Time	Date/Time of Snapshot
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Used	Used in GB or blocks
% Util	Used/Hard Limit
Master	Either Volume name for snapshots to volumes or Snapshot Name for snapshots of snaps
Locking Status	If the field is empty, it means it is Not Locked
Serial #	Snapshot serial number
Host	Name of the host that is using the storage
Host LUN	The Snapshot to which the host is mapped

Table 9-118 XIV Snapshot Details *(continued)*

Column name	Description
Mount Point	The host's filesystem mount point for this Snapshot. It is possible to have zero mount points and filesystems, in the scenario where a Snapshot has been allocated to the host, but the host has not yet made use of it by creating a filesystem and mounting it.
Filesystem	The host's filesystem or drive name to which the Snapshot maps. Snapshots simply appear as disks to hosts.
Host Storage Domain	A logical group of hosts that can be accessed via the same physical storage port, with each host accessing its own unique Snapshot.
Array Port WWN	The unique identifier, World Wide Name, of the array port address.

IBM XIV Storage Pools Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization> IBM XIV Storage Pools Summary

The XIV Storage Pools Summary report provides a view of IBM XIV storage pool consumption, including a hard and soft status indicators to draw attention to when physical and logical capacity has been over-provisioned are reaching capacity limits. For example, for growing LUNs, 30% utilized may not be considered to be over-provisioned. However, over time, it may warrant a second look to determine a more efficient approach to storage allocation.

When generating this report, the Scope Selector enables customization of the hard and soft capacity values used to render the report.

See "How Capacity Values are Represented in XIV Reports" on page 334.

Table 9-119 IBM XIV Storage Pools Summary

Column name	Description
Pool Name	The name of the storage pool links to the XIV Storage Pool Details.
Storage Array	Name of the XIV array links to the Array Capacity & Utilization report
Thin	Y = thin provisioning
Hard Limit	Maximum physical capacity that can be used by volumes and snapshots in the pool
Used	Physical capacity used by the volumes (see Allocated) + Snapshot Used
% Util	Used/Hard Limit
Available	Capacity that is available for allocation
Allocated	Amount of storage used by volumes; Note that Allocated does NOT include Snapshots.
Soft Limit	The logical amount of space for all storage pools in the array. A soft limit greater than the hard limit indicates thin provisioning.
Soft Util %	Soft Util % = Soft limit - Available/Soft Limit
Snapshot Used	Amount of storage in use by snapshots
Lock Status	Locked status
Lock Behavior	Lock behavior, such as Read Only
# CGs	Number of Consistency Groups associated with this pool, links to the XIV Consistency Group Summary report
# Volumes	Number of volumes and snapshots associated with this pool, links to XIV LUN Utilization Summary report
Hard Status	Physical Capacity Status: If % Util < 2% = empty, if % Util > 95% = near capacity, if % Util > 2% and <30% = over-provisioned

Table 9-119 IBM XIV Storage Pools Summary *(continued)*

Column name	Description
Soft Status	Logical Capacity Status: If % Util < 2% = empty, if % Util > 95% = near capacity, if % Util > 2% and <30% = over-provisioned

IBM XIV Storage Pool Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > IBM XIV Storage Pools Summary

Next, click a Pool Name.

The XIV Storage Pool Details offers a graphical representation of pool usage. Note that thin-provisioned pools have a soft limit greater than the hard limit.

See "How Capacity Values are Represented in XIV Reports" on page 334.

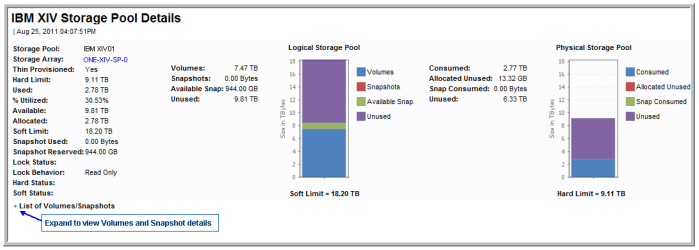


Table 9-120 IBM XIV Storage Pool Details

Column name	Description
XIV Storage Pool Details	
Storage Pool	Storage Pool name
Storage Array	Storage Array name
Thin Provisioned	Yes = thin provisioned
Hard Limit	Maximum physical capacity that can be used volumes and snapshots in the pool.
Used	Physical capacity minus Free Hard size
% Utilized	Used/Hard Limit
Available	Amount of available pool space
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Soft Limit	The logical amount of space for all storage pools in the array. A soft limit greater than the hard limit indicates thin provisioning.
Snapshot Used	Amount of storage used by snapshots
Snapshot Reserved	Amount of storage for snapshot reserves
Lock Status	If this field is empty, it means Not Locked
Lock Behavior	Lock behavior, such as Read Only
Hard Status	Physical Capacity Status: If % Util < 2% = empty, if % Util > 95% = near capacity, if % Util > 2% and <30% = over-provisioned
Soft Status	Logical Capacity Status: If % Util < 2% = empty, if % Util > 95% = near capacity, if % Util > 2% and <30% = over-provisioned
Logical Storage Pool	
Volumes	Sum of the capacity of all master volumes.

Table 9-120 IBM XIV Storage Pool Details *(continued)*

Column name	Description
Snapshots	Amount of storage pool space consumed by snapshots
Available Snap	Total snapshot size minus space used by snapshots
Unused	Amount of free soft (logical) storage in the pool
Physical Storage Pool	
Consumed	Sum of all capacity used by master volumes
Allocated Unused	Amount of storage pool space used by volumes (see Used in first column of this report) minus Consumed (value shown above)
Snap Consumed	Amount of storage pool space used by snapshots
Unused	Free hard (physical) space
List of Volumes/Snapshots	
Name	Volume/Snapshot name
Allocated	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Used	Used capacity in GB
% Util	% of volume or snapshots used
Master	Either Volume name for snapshots to volumes or Snapshot Name for snapshots of snaps
Consistency Group	Consistency group name, if exists, links to CG Detail Report
Locking Status	If this field is empty, it means it is Not Locked.
Serial #	Snapshot serial number

Table 9-120 IBM XIV Storage Pool Details *(continued)*

Column name	Description
Host	Links to the Host Utilization Details
Host LUN	The LUN to which the host is mapped
Mount Point	The host's filesystem mount point for this LUN. It is possible to have zero mount points and filesystems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a filesystem and mounting it.
Filesystem	The host's filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts.
Host Storage Domain	A logical group of hosts that can be accessed via the same physical storage port, with each host accessing its own unique Snapshot.
Array Port WWN	The unique identifier, World Wide Name, of the array port address.

LUN Utilization Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > LUN Utilization Summary

If LUNs are not being used efficiently, you want to know so that you can change their allocation or reassign them altogether. Using the LUN Utilization Summary, you can view all the mapping details, such as mount points and file systems.

- For NetApp-specific LUN Utilization details,
  See " LUN Utilization Summary for NetApp 7-Mode " on page 355.
- For EMC VPLEX-specific LUN Utilization details,
  See "LUN Utilization Summary for EMC VPLEX Arrays" on page 349.
- For IBM XIV-specific LUN Utilization details,

See “LUN Utilization Summary for IBM XIV Arrays” on page 351.

- For IBM SVC-specific LUN Utilization details,
See “IBM SVC VDisk Summary” on page 327.
- For HP-3PAR- specific LUN Utilization details,
See “ LUN Utilization Summary for HP3PAR” on page 358.

This report can be filtered (in the Scope Selector window) to report on the following subsets of data:

- Unallocated - These are simply LUNs that have not been allocated to any host. The array has created the LUN, but there is no storage path.
- Allocated, but undiscovered - Often referred to as orphaned LUNs, these are LUNs that have been assigned to a host, but they have not been discovered and therefore are not seen on the host side. The array has created the LUN and the storage path exists, but there is no access to host port and host information, such as disks, partitions, and volume groups. To make these LUNs visible to the host, a system administrator must execute a LUN scan or, for older OS and volume manager versions, re-boot the system. Otherwise, this storage is not available for mounting filesystems. The LUN scan tells the HBA (host bus adapter) that a LUN has been zoned to the host.
- Allocated, but unused - These are LUNs that have been assigned to a host, but the host has not been placed into a volume group, or a partition has not been created. The array has created the LUN and the storage path exists, making the host and host port information available to Capacity Manager. However, the host storage details--disks, partitions, and volume groups--are not available for the host associated with the storage path.

Note: The list of Allocated, but unused LUNs will always include the Allocated, but undiscovered LUNs, along with the additional LUNs that may not have been discovered.

Table 9-121 LUN Utilization Summary

Column name	Description
	Click on the **arrow**to drill down to a hierarchical detailed view, including mount point details. This arrow is not displayed if the LUN does not have multiple mount points or has no mount points. For HP 3PAR: LUN refers to a Virtual Volume (VV).

Table 9-121 LUN Utilization Summary *(continued)*

Column name	Description
LUN Name	The name of the LUN drills down to the LUN Details report. For HP 3PAR: LUN refers to a Virtual Volume (VV).
RAID Type	The RAID level on the LUN. When the RAID type indicates **Disk**, there is no RAID protection. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Aggregate	(NetApp only).
Volume	(NetApp only).
QTree	QTree associated with the LUN (NetApp only).
Array Group	The group to which the array belongs. See " Array Group Details " on page 359.
Storage Array	The array on which the LUN was created. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Capacity	Size of the volume.
IBM SVC-specific fields for the IBM SVC VDisk Summary report	
Name	The volume's user-friendly name.
Cluster	Name of the storage array.
MDisk Group ID	The name of the MDisk group.
Capacity	Total VDisk (LUN) capacity. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Copy Capacity	Copy capacity includes mirrors.
Real Capacity	Allocated VDisk capacity; the capacity that is available to hosts.
Used Capacity	Amount of the allocated VDisk capacity (real capacity) that is used.
Free Capacity	Free = Capacity - Real Capacity.

Table 9-121 LUN Utilization Summary *(continued)*

Column name	Description
Uncompressed Used Capacity	For compressed volumes, indicates the amount of data written to the volume before compression.
Compressed	Y if the volume is compressed.
Thin Provisioned	Y if the volume is thinly provisioned.
Mirrored	Indicates if this LUN has a mirror. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Space Efficient	Indicates if this is a thin-provisioned LUN. For HP 3PAR: LUN refers to a Virtual Volume (VV).
IO Group Name	The name of the I/O group to which this VDisk belongs.
Estimated Used	LUN estimated used. For HP 3PAR: LUN refers to a Virtual Volume (VV). Dell EMC Unity: This value is same as 'Allocated' attribute in Dell EMC Unity. Amount of pool space reserved for the LUN, may be less than the total size for a thin LUN. For EMC XtremIO: This field is N/A. See "Estimated LUN Usage/Availability" on page 348.
Estimated Available	Estimated available LUN storage. For HP 3PAR: LUN refers to a Virtual Volume (VV).
# of Hosts	Number of hosts that have allocated LUNs.
Host LUN	The LUN to which the host is mapped. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Mount Point	The host's filesystem mount point for this LUN. It is possible to have zero mount points and filesystems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a filesystem and mounting it. For HP 3PAR: LUN refers to a Virtual Volume (VV).

Table 9-121 LUN Utilization Summary *(continued)*

Column name	Description
Filesystem	The host's filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Volume Group	The logical collection or volume group to which the volume belongs. For HP 3PAR: Volume group refers to Common Provisioning Group (CPG).
iGroup	Initiator Group name. For NetApp storage, a LUN is mapped to an iGroup--a method of LUN masking.
Host Storage Domain	The name of the host storage domain, a logical group of LUNs that can be accessed by multiple hosts through the same physical storage port, with each host accessing its unique LUN. For security purposes, LUN masking is also coupled with this capability. A host storage domain is Hitachi's name for multi-hosted storage capability. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Array Port WWN	The unique identifier, World Wide Name, of the array port address.
iSCSI Interface	The iGroup contains an iSCSI qualified name that maps to a LUN. (NetApp only).

Estimated LUN Usage/Availability

Several Capacity Manager reports display certain LUN values as estimated--Estimated Usage, Estimated Available. These values are estimated because of the following logic for deriving the values:

- A volume group has members that are made up of many disks.
- These disks may be local disks or LUNs from one or more arrays.
- Partitions are created from a volume group and then the partition is mounted onto a host, so that the host can use the mounted partition.

- Because of the above sequence, Capacity Manager cannot determine which partition came from which LUN. Therefore, using the available information, the values are estimated.

LUN Utilization Summary for EMC VPLEX Arrays

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > LUN Utilization Summary

Select an EMC VPLEX array scope in the scope selector.

Table 9-122 LUN Utilization Summary for EMC VPLEX Arrays

Column name	Description
Cluster	Name of the storage array.
Virtual Volume	Name of the virtual volume.
Mirror Cluster	Name of the mirror cluster of the distributed virtual volume.
# of Blocks	The number of blocks in the extent.
Block Size	The size of a single block, in kilobytes.
Capacity	The total number of bytes in the virtual volume. This total represents the block size multiplied by the block count.
Used	The total amount of used capacity.
Consistency Group	The name of the consistency group to which this volume belongs, if any.
Expandable	Indicates if the virtual volume is expandable.

Table 9-122 LUN Utilization Summary for EMC VPLEX Arrays *(continued)*

Column name	Description
# of Devices	Number of devices in the virtual volume.
Health State	Health state of the virtual volume. Possible values: Major Failure = One or more of the virtual volume's underlying devices is out-of-date, but will never rebuilt. Minor Failure = One or more of the virtual volumes underlying devices is out-of-date, but will rebuild. Non-Recoverable Error = VPLEX cannot determine the virtual volume's health state. OK = The virtual volume is functioning normally, Unknown = VPLEX cannot determine the virtual volume's health state, or the state is invalid.
Locality	Locality of the virtual volume. Possible values: Local = The virtual volume relies completely on storage at its containing cluster. Remote = The virtual volume is a proxy for a volume whose storage resides at a different cluster. I/O to a remote virtual volume travels between clusters. Distributed = The virtual volume is the cluster local representation of a distributed RAID-1. Writes to a distributed volume travels to all the clusters for which it has storage; reads derive, if possible, from the local leg.

Table 9-122 LUN Utilization Summary for EMC VPLEX Arrays *(continued)*

Column name	Description
Oper State	Operational status of the virtual volume. Possible values: Degraded = The virtual volume may have one or more out-of-date devices that will eventually rebuild. Error = One or more of the virtual volume's underlying devices is hardware-dead. OK = The virtual volume is functioning normally. Starting = The virtual volume is not ready. Stressed = One or more of the virtual volume's underlying devices is out-of-date and will never rebuild. Unknown = VPLEX cannot determine the virtual volume's operational state, or the state is invalid.
Storage Tier	The storage tier for the virtual volume.
Supporting Device	The local, remote, or distributed device underlying this volume.
Volume Type	Volume type, which is always set to virtual volume.
# of Hosts	Number of hosts that allocated LUNs.
Host LUN	The LUN to which the host is mapped.
Array Port WWN	List of Array Port WWN for each logical unit This is the list of Array Port WWN List for all logical unit.

LUN Utilization Summary for IBM XIV Arrays

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > IBM XIV Storage Pools Summary

Next, click a link from the # Volumes column.

See "How Capacity Values are Represented in XIV Reports" on page 334.

Table 9-123 LUN Utilization Summary for IBM XIV Arrays

Column name	Description
Volume Name	Volume name
Allocated Size	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Used	Used storage in GB or blocks consumed
Snaps	Number of snapshots, links to Snapshot Detail Report
Consistency Group	Consistency Group Name, Links to CG Detail Report
Master Volume	The Master volume of the snapshot.
Pool	Storage Pool Name, links to Array display
Thin	Y = Thin Provisioning by Storage Pool
Storage Array	Array Name links to Array Display
Locking Status	The locking status of the volume or snapshot. If the field is empty, it means it is Not Locked
Serial #	Volume serial number
Sync Type	Mirroring Status, if mirrored

Table 9-123 LUN Utilization Summary for IBM XIV Arrays *(continued)*

Column name	Description
Current Role	For mirrored volumes and snaps, either Primary or Secondary servers.
Host	Name of the host that is using the storage
Host LUN	The LUN to which the host is mapped
Mount Point	The host's filesystem mount point for this LUN. It is possible to have zero mount points and filesystems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a filesystem and mounting it.
Filesystem	The host's filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts.
Array Port WWN	The unique identifier, World Wide Name, of the array port address.
Related System	Related storage array
Related Pool	Related storage pool
Related Volume	Related volume

LUN Utilization Summary for Hitachi Vantara

This is a drill down detail report that you can access when you click the **# Resv Vols/LUNs**, **# Vols/LUNs**, **# Allocated Vols/LUNs**, and **# UnAllocated Vols/LUNs** values displayed for the Hitachi Vantara All-Flash and Hybrid Flash Storage in the Array Capacity & Utilization report. This report provides details specific to the values clicked on the main report.

The details displayed in the report are as follows:

Table 9-124 Hitachi Vantara LUN Utilization Summary

Column name	Description
LUN ID	LDEV number.

Table 9-124 Hitachi Vantara LUN Utilization Summary *(continued)*

Column name	Description
LUN Name	Hexadecimal value of LDEV number. Example. LDEV number 239 corresponds to 00:00:EF.
Storage Array	Storage device ID.
Name	LDEV nickname.
Pool	ID of the pool with which the LDEV is associated. ■ For DP volumes, it is the ID of the associated DP pool. ■ For virtual volumes for Thin Image, it is the ID of the pool in which the snapshot data was created.
Volume type	LDEV attributes. Possible values are: CVS, VVOL, CLUN, ALUN, ELUN, OLG, VVOL, HORC, HDP, JNL, POOL, and more.
Status of LDEV	Status of the LDEV. Possible values are: ■ NML: The LDEV is in normal status. ■ BLK: The LDEV is blocked. ■ BSY: The LDEV status is being changed. ■ Unknown: The LDEV status is unknown (not supported).
RAID type	Raid type indicating the drive configuration.
Array Group	Array Group list for each logical unit This is the list of Array Group List for all logical unit.
Virtual Serial #	Serial number of the storage virtual machine.
Capacity	Capacity of the LDEV.
Estimated Used	Used capacity in the pool. This attribute is present when the LDEV is a virtual volume.
Estimated Available	Estimated available LUN storage.
# of Hosts	Number of hosts that have allocated LUNs.
Host Storage Domain	Object ID for the host group or iSCSI target.

LUN Utilization Summary for NetApp 7-Mode

This NetApp-specific report can be accessed via a LUN link in NetApp-specific reports.

See " NetApp Snapshot Summary " on page 392.

A similar report is available for all storage arrays in your environment.

See " LUN Utilization Summary " on page 344.

Using the LUN Utilization Summary, you can view all the mapping details, such as mount points and file systems.

Table 9-125 LUN Utilization Summary for NetApp 7-Mode

Column name	Description
LUN Name	The name of the LUN drills down to the LUN Details report.
RAID Type	The RAID level on the LUN. When the RAID type indicates **Disk**, there is no RAID protection.
Aggregate	Links to the NetApp Aggregate Summary. See " NetApp Aggregate Summary " on page 364.
Volume	Links to the NetApp Snapshot Summary. See " NetApp Snapshot Summary " on page 392.
QTree	Name of the QTree to which this LUN belongs. See " NetApp QTree Summary " on page 389.
Storage Array	Links to the LUN Utilization Summary. See "LUN Utilization Summary for IBM XIV Arrays" on page 351.
Capacity	The total capacity (used and unused) of the LUN.
Estimated Used	LUN estimated used. See "Estimated LUN Usage/Availability" on page 348.

Table 9-125 LUN Utilization Summary for NetApp 7-Mode *(continued)*

Column name	Description
Estimated Available	Amount of LUN estimated to be available
Mount Point Used	The amount of storage used by each mount point. These values are displayed when a LUN is expanded to view mount point details.
Host	The name of the host to which the LUN was allocated. See " Host Utilization Detail " on page 481.
Host LUN	The LUN associated with the host.
Mount Point	The host's file system mount point for this LUN. It's possible to have zero mount points and file systems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a file system and mounting it.
File System	The host's file system or drive name to which the LUN maps. LUNs simply appear as disks to hosts.
Volume Group	The logical collection or volume group to which the volume belongs.
iGroup	Initiator Group name for NetApp storage; a LUN is mapped to an iGroup; that is, a method of LUN masking.
Array Port WWN	The unique identifier, World Wide Name, of the array port address.
iSCSI Interface	The iGroup contains an iSCSI qualified name that maps to a LUN. (NetApp only)
Space Reserved	Indicates if LUN space reservation is enabled.

LUN Utilization Summary for NetApp Cluster-Mode

This report is specific to NetApp Cluster-Mode. For NetApp 7-Mode,

See " LUN Utilization Summary for NetApp 7-Mode " on page 355.

A similar report also is available for all storage arrays in your environment;

See “ LUN Utilization Summary ” on page 344.

Using the LUN Utilization Summary, you can view mapping details, such as mount points and file systems.

Table 9-126 LUN Utilization Summary for NetApp Cluster-Mode

Column name	Description
LUN Name	The name of the LUN drills down to the LUN Details report.
RAID Type	The RAID level on the LUN. When the RAID type indicates **Disk**, there is no RAID protection.
Aggregate	Name of the NetApp aggregate.
Cluster	Name of the storage array.
Capacity	Size of this LUN in bytes in the active file system. This field is unavailable while the LUN is fenced for a restore operation.
Estimated Used	Number of bytes used by this LUN. This field is unavailable while the LUN is fenced for a restore operation.
Estimated Available	Estimated Available = Capacity - Estimated Used.
Mount Point Used	The amount of storage used by each mount point. These values are displayed when a LUN is expanded to view mount point details.
# Hosts	Number of hosts that allocated LUNs.
Host LUN	The LUN associated with the host.
Array Port WWN	List of array port world wide names.
OS Type	OS type of the LUN.
Mapped	Indicates if the LUN is mapped to any initiators: Yes or No.

LUN Utilization Summary for HP3PAR

This HP3PAR-specific report can be accessed via a LUN link in HP3PAR-specific reports. A similar report is available for all storage arrays in your environment.

See " LUN Utilization Summary " on page 344.

Using the LUN Utilization Summary, you can view all the mapping details, such as mount points and file systems.

Table 9-127 LUN Utilization Summary for HP3PAR

Column name	Description
VV Name	The name given to the logical unit.
Array Name	RAID group name.
Type	RAID type of the volume: RAID 0, RAID10, or RAID50.
Allocated Size	Size of the volume. Values are stored as KiB in the database and rendered according to your user profile preferences.
User	Raw user. The total number of logically contiguous blocks * Block size = Total raw size of the User space. Values are stored as KiB in the database and rendered according to your user profile preferences.
Snaps	Used snapshot admin. Values are stored as KiB in the database and rendered according to your user profile preferences.
Admin	Raw admin. The total number of logically contiguous blocks * Block size = Total raw size of the Admin space. Values are stored as KiB in the database and rendered according to your user profile preferences.
Allocation Warning	Allocation warning. Values are stored as KiB in the database and rendered according to your user profile preferences.
Allocation Limit	Allocation limit. Values are stored as KiB in the database and rendered according to your user profile preferences.
CPG	RAID group name.

Table 9-127 LUN Utilization Summary for HP3PAR *(continued)*

Column name	Description
RAID Type	RAID level to display in reports.
Device Type	Device type of the disk. 0 = Unknown 1 = FC 2 = Nearline
Set Size	Number of chunklets in a set.
Availability	Describes the current availability of the volume. This value might be different from preferred availability. 0 = Unknown 1 = Cage 2 = Magazine 3 = Port 4 = Disk 5 = Chunklet
Thin	Indicates if the LUN is thinly provisioned: Y or N.
Storage Array	Name of the storage array.
Host	The name is derived from the host's DNS name.
Host LUN	The LUN to which the host is mapped. For HP 3PAR: LUN refers to a Virtual Volume (VV).
Mount Point	Mount point associated with the host LUN path. For HP 3PAR: LUN refers to a Virtual Volume (VV).
File System	The host's filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts. For HP 3PAR: LUN refers to a Virtual Volume (VV).

Array Group Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > LUN Utilization Summary

Click an **Array Group** link.

For HP 3PAR arrays,

See " HP 3PAR Common Provisioning Group (CPG) Summary " on page 592.

This report lists all the LUNs and the physical devices associated with the selected array group.

Table 9-128 Array Group Details

Column name	Description
Storage Array	Links to the Array Capacity & Utilization.
Chassis Number	The array's chassis number.
Group #	The array group number.
Group Name	The array group name.
Object ID	The object ID, for object-based storage.
RAID Type	RAID type for the array group.
Raw Capacity	Total capacity for this array group.
Usable Capacity	The amount of usable capacity available for creating LUNs; Usable capacity is what is available after storage has been RAIDed.
Used Capacity	Amount of space that has already been used.
Free Usable Capacity	Amount of space available. For HDS arrays, this represents the space from which an LDEV can be carved.
LUN	Links to the LUN Utilization Summary.
Total Capacity	LUN's capacity.
RAID Type	LUN's RAID type.
Device Number	LUN's device number.
Device Name	Physical disk name.

Table 9-128 Array Group Details *(continued)*

Column name	Description
Capacity	Disk's capacity. This capacity value includes the capacity of hot spares. Therefore, the total for this list of values may not be equal to the total raw capacity listed at the top of this report.
Vendor	Disk manufacturer.
Serial #	Disk's serial number.
Model	Disk model.
Firmware Version	Current firmware version of the device.

EMC CLARiiON Storage Array Group Specifics

For EMC CLARiiON storage arrays, the operating system uses a small amount of storage for internal processing. So, in the Array Group Details report, you will notice a difference in the Total (at the top left of the report) and the sum of the Capacity in the List of Physical Devices table (at the bottom right of the report).

Microsoft Azure Endpoint Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Microsoft Azure Endpoint Summary

View Microsoft Azure endpoints associated with an Azure subscription.

Account subscriptions are listed with a summary of related blob, share, table, or queue endpoint data.

Table 9-129 Microsoft Azure Endpoint Summary

Column name	Description
Account	The Account Administrator Email Address of Microsoft Azure Cloud account.
Storage Account	Name of the Storage Account.
Storage Account Type	The Storage Account type, for eg.Standard_RAGRS,Standard_LRS.
Endpoint Type	Specifies the type of Endpoint. Possible values: blob,file,queue,table.
Endpoint Name	Name of the Endpoint.
Container Size	Container size.
Last Updated	Last Updated Date.

Microsoft Azure Storage Account Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Microsoft Azure Storage Account Summary

View Microsoft Azure storage accounts associated with an Azure subscription.

Account information provided in this report includes a summary of blob containers, queues, shares, capacity, and geographical region information.

Table 9-130 Microsoft Azure Storage Account Summary

Column name	Description
Account	The Account Administrator Email Address of Microsoft Azure Cloud account.
Subscription ID	The subscription ID of Microsoft Azure Cloud account.
Storage Account	Name of the Storage Account.
Storage Account Type	The Storage Account type, for eg.Standard_RAGRS,Standard_LRS
Location	Location of the Storage Account.
# of Endpoints	Number of Endpoints for this Storage Account.
# of Containers	Number of the containers for this Storage Account.
Container Size	Container size in bytes.
# of File Shares	Number of the file shares associated with this Storage Account.
# of Tables	Number of the Tables associated with this Storage Account.
# of Queues	Number of the Queues associated with this Storage Account.
Primary Location	Primary Location of the Storage Account.
Primary Status	Status of the primary location of Storage Account. Possible values: Available, Unavailable.
Secondary Location	Secondary Location of the Storage Account.
Secondary Status	Status of the secondary location of Storage Account. Possible values: Available, Unavailable.
Last Updated	Last Updated Date

NetApp Aggregate Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Aggregate Summary

or

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Next, click **# Aggregates**.

View this report for an overview of how your aggregates are configured to manage disks, RAID groups, and Plexes. The Aggregate Utilization Summary displays one row for each aggregate, with links to Volume and Plex details.

Note: If the aggregate is mirrored, two Plexes will be listed.

Table 9-131 NetApp Aggregate Summary

Column name	Description
Aggregate	The Aggregate number/ID
Storage Array	The NetApp system with which the aggregate is associated
% Usage of Aggregate	Mouse over the usage thermometer to view the % available
Total Capacity	The aggregate’s total capacity

Table 9-131 NetApp Aggregate Summary *(continued)*

Column name	Description
Allocated Size	Allocated storage is physical storage consumed, including parity overhead and system overhead. A gap between raw and allocated capacity may indicate unused capacity, such as drives in a rack that are not in use.
Available Size	Remaining capacity available for use
RAID Type	raid0 (mirrored), raid4 (single disk for parity), or raid_dp (double-parity protection)
# Volumes	Links to the NetApp Snapshot Summary. See " NetApp Snapshot Summary " on page 392.
# of Thin Prov Vols	Number of thin provisioned volumes
List of Plexes	Lists 1 or 2 Plexes--2 if SyncMirror is licensed and enabled. See " NetApp Plex Details " on page 373. report.

In addition, the following Disk Type information is displayed:

Table 9-132 NetApp Aggregate Summary

Column name	Description
Raid Status	RAID status
Count	Number of disks
Total Capacity	Total capacity of the disks

NetApp Cluster-Mode Aggregate Summary

For NetApp Cluster Mode, this report lists aggregates with associated capacity metrics along with other related details, such as cache size, volumes, and Plexes.

Note that if the aggregate is mirrored, two Plexes are listed. Use this report for unified visibility in your NetApp ONTAP Cluster Mode environment.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Aggregate Summary

Table 9-133 NetApp Cluster-Mode Aggregate Summary

Column name	Description
Aggregate	Name of the NetApp aggregate. See " NetApp Volumes at Risk " on page 458.
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Node	The textual name of a node. See "NetApp Cluster-Mode Node Summary" on page 385.
Usage	Usage = Allocated Size / Total Capacity
Total Capacity	Total size (in bytes) of the file system. If the file system is restricted or offline, this value is 0.
Allocated Size	Number of bytes used in the file system. If the file system is restricted or offline, this value is 0. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available Size	Number of bytes still available in the file system. If the file system is restricted or offline, this value is 0. Values are stored as KiB in the database and rendered according to your user profile preferences.
RAID Status	RAID status.

Table 9-133 NetApp Cluster-Mode Aggregate Summary *(continued)*

Column name	Description
RAID Type	RAID type. Possible values: raid0 = All the RAID groups in the aggregate are of type RAID 0. raid4 = All the RAID groups in the aggregate are of type RAID 4. raid_dp = All the RAID groups in the aggregate are of type RAID DP.mixed_raid_type The aggregate contains RAID groups of different RAID types (RAID 0, RAID 4, RAID DP).
Aggregate State	Aggregate state. Possible values: creating, destroying, failed, frozen, inconsistent, iron_restricted, mounting, online, offline, partial, quiesced, quiescing, restricted, reverted, unknown, unmounted, unmounting, relocating.
# Disks	Number of disks in the aggregate. See "NetApp Cluster-Mode Disk Summary" on page 379.
# Volumes	Number of volumes in the aggregate. See "NetApp Cluster-Mode Volume Summary" on page 396.
# Thin Volumes	Number of thin volumes in the aggregate. See "NetApp Cluster-Mode Volume Summary" on page 396.
List of Plexes	List of plex names. If the aggregate is mirrored, two Plexes are listed
Hybrid Cache Size	Total cache size (in bytes) for a hybrid aggregate. If the aggregate is restricted or offline, or if it is not a hybrid aggregate, this value is 0.Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-133 NetApp Cluster-Mode Aggregate Summary *(continued)*

Column name	Description
Cache RAID Group Size	Maximum cache RAID group size of a hybrid aggregate. RAID group size is the maximum number of disks that can be added to a RAID group. Values are stored as KiB in the database and rendered according to your user profile preferences.

NetApp Aggregate Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

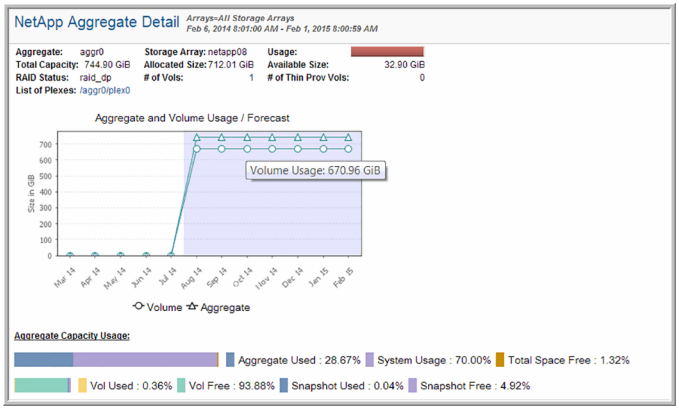
Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Next, click **# Aggregates**.

Click an **Aggregate Name** link.

The Aggregate Detail forecasts usage, based on historical data. Mouse over the Aggregate and Volume usage symbols to view the historical and projected usage.



Aggregate Capacity Usage

At the bottom of the window, the cascade of usage bars illustrates the logical subsets of the aggregate.

- Mouse over the colored bar segments to view the actual values instead of the percentages.

Table 9-134 Aggregate Capacity Usage

Column name	Description
Aggregate	Name of the NetApp aggregate. See " NetApp Volumes at Risk " on page 458.
Total Capacity	Total size (in bytes) of the file system. If the file system is restricted or offline, this value is 0.
RAID Status	RAID status - comma separated list of values.
List of Plexes	Lists 1 or 2 Plexes--2 if SyncMirror is licensed and enabled. See " NetApp Plex Details " on page 373. report.
Storage Array	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.

Table 9-134 Aggregate Capacity Usage *(continued)*

Column name	Description
Allocated Size	Number of bytes used in the file system. If the file system is restricted or offline, this value is 0. Values are stored as KiB in the database and rendered according to your user profile preferences.
# of Vols	Number of volumes in the aggregate. Drill down to NetApp Volume Summary.
% Usage of Aggregate	Usage = Allocated Size / Total Capacity
Available Size	Available bytes in the aggregate.
# of Thin Prov Vols	Number of thin provisioned volumes in the aggregate.

NetApp Cluster-Mode Aggregate Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

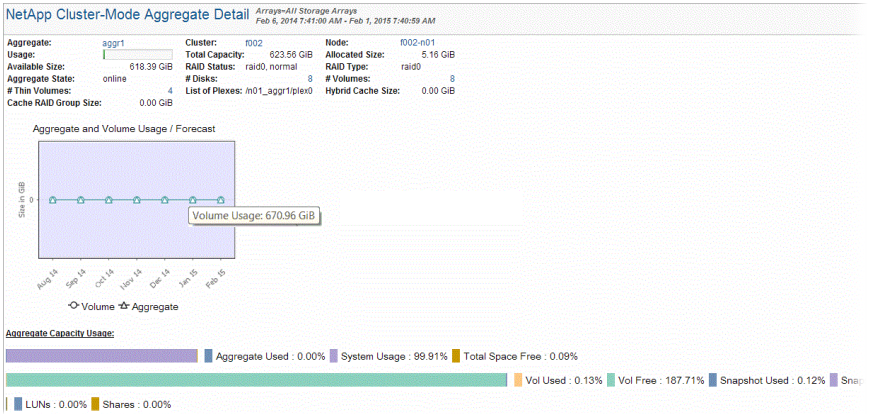
Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Aggregate Summary

Next, click **# Aggregates**.

Click an **Aggregate Name** link.

The NetApp Cluster-Mode Aggregate Detail forecasts usage, based on historical data. Mouse over the Volume and Aggregate symbols to view the historical and projected usage.



Aggregate Capacity Usage

At the bottom of the window, the cascade of usage bars illustrates the logical subsets of the aggregate.

- Mouse over the colored bar segments to view the actual values instead of the percentages.

Table 9-135 Aggregate Capacity Usage

Column name	Description
Aggregate	Name of the NetApp aggregate. See " NetApp Volumes at Risk " on page 458.
Usage	Usage = Allocated Size / Total Capacity
Available Size	Available bytes in the aggregate.
Aggregate State	Aggregate state.
# Thin Volumes	Number of thin provisioned volumes in the aggregate.
Cache RAID Group Size	Maximum cache RAID group size of a hybrid aggregate. RAID group size is the maximum number of disks that can be added to a RAID group. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-135 Aggregate Capacity Usage *(continued)*

Column name	Description
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Total Capacity	Total size (in bytes) of the file system. If the file system is restricted or offline, this value is 0.
RAID Status	RAID status - comma separated list of values.
# Disks	Number of disks in the aggregate.
List of Plexes	Lists 1 or 2 Plexes--2 if SyncMirror is licensed and enabled. See " NetApp Plex Details " on page 373. report.
Node	The textual name of the node. See "NetApp Cluster-Mode Node Summary" on page 385.
Allocated Size	Number of bytes used in the file system. If the file system is restricted or offline, this value is 0. Values are stored as KiB in the database and rendered according to your user profile preferences.
RAID Type	RAID type. Possible values: raid0 = All the RAID groups in the aggregate are of type RAID 0. raid4 = All the RAID groups in the aggregate are of type RAID 4. raid_dp = All the RAID groups in the aggregate are of type RAID DP.mixed_raid_type The aggregate contains RAID groups of different RAID types (RAID 0, RAID 4, RAID DP).

Table 9-135 Aggregate Capacity Usage *(continued)*

Column name	Description
# Volumes	Number of volumes in the aggregate. See " NetApp Volume Summary " on page 395.
Hybrid Cache Size	Total cache size (in bytes) for a hybrid aggregate. If the aggregate is restricted or offline, or if it is not a hybrid aggregate, this value is 0.Values are stored as KiB in the database and rendered according to your user profile preferences.

NetApp Plex Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Next, click **# Aggregates**.

Note: The # of Aggregates is not listed for a VFiler, as it is not relevant.

Then, in the Aggregate Utilization Summary, click a **Plex**link.

This drilldown sequence launches the Plex Details report, which displays a list of details for each of the disks in the Plex's disk pool. Use this report to understand how the RAID groups and disks are organized and configured.

RAID Reconstruction Status

Table 9-136 RAID Reconstruction Status:

Column name	Description
Name	Name of the Plex
Is Reconstructing?	Y or N During RAID reconstruction, as a result of a disk failure, redundant data is read from other disks to recreate the blocks from the failed disk.
Reconstruction %	If reconstruction is in progress, this value shows the completion percentage.

Plex Disk Pool Details

See " NetApp Disk Summary " on page 378.

for a description of these fields.

NetApp CIFS Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp CIFS Summary

or

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Then, click **# CIFS Shares**.

View the volumes that have been exported to Windows hosts through CIFS (Common Internet File System) shares. This report displays the usage of each share.

Table 9-137 NetApp CIFS Summary

Column name	Description
Name	Directory name of the share
Mount Point	Mount point of the share
Max User	The maximum number of simultaneous connections to the share--0 indicates no limit
Is Volume Offline	Indicates if the volume is offline--Yes or No
Force Group	This group overrides the group of the file's owner, so if this field is blank, it defaults to the file owner's group.
Volume	Name of the volume links to the NetApp Snapshot Summary. See " NetApp Snapshot Summary " on page 392.
QTree	Name of the associated QTree links to the NetApp QTree Summary. See " NetApp QTree Summary " on page 389.
Storage Array	Links to the LUN Utilization Summary for NetApp 7-Mode. See " LUN Utilization Summary for NetApp 7-Mode " on page 355.
Description	Description of the share

NetApp Cluster-Mode CIFS Summary

NetApp CIFS Shares are listed in a tabular report, with related objects and configuration metrics for the associated array.

Use this report for Unified Visibility to determine the status of CIFS shares and their associated volumes. Cluster Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode CIFS Summary

Table 9-138 NetApp Cluster-Mode CIFS Summary

Column name	Description
Name	The CIFS share name is a UTF-8 string.
Mount Point	File system path that is shared through this CIFS share. This is the full, user-visible path relative to the Vserver root, and it may cross junction mount points. The path is in UTF8 and uses forward slash as directory separator.
Volume State	Volume state. Possible values: online, restricted, offline, mixed. If an Infinite Volume's constituents are not all in the same state, this state is set to mixed. The mixed state is read-only.
Volume	The name of the volume. See " NetApp Cluster-Mode Volume Detail " on page 401.
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Description	Description of the CIFS share. CIFS clients see this description when browsing the Vserver's CIFS shares.
Vserver	Name of the Vserver. See "NetApp Cluster-Mode Vserver Summary" on page 413.
CIFS Server Name	The NetBIOS name of the CIFS server.
Status	CIFS Server Administrative Status. Possible values: down, up.

Table 9-138 NetApp Cluster-Mode CIFS Summary *(continued)*

Column name	Description
# Qtrees	Number of Qtrees in the volume. See "NetApp Cluster-Mode QTree Summary" on page 390.
CIFS Share Properties	The list of properties for the CIFS share. Possible values: oplocks = Opportunistic locks (client-side caching) are enabled. browsable = Share can be browsed by Windows clients. showsnapshot = Snapshots can be viewed and traversed by clients. changenotify = CIFS clients can request change notifications for directories on the share. homedirectory = Share is added and enabled as part of the CIFS home directory feature. T attributecache = Connections through this share cache attributes for a short time to improve performance. branchcache = Clients connecting to this share can make requests using BranchCache technology, allowing them to cache the content to reduce WAN utilization from a remote office. continuously_available = Clients connecting to the share can open files in a persistent manner. Files opened this way are protected from disruptive events. shadowcopy = Share is pointing to a shadow copy. access_based_enumeration = Access Based Enumeration (ABE) is enabled on the share. ABE-filtered shared folders are only available based on a user's access rights.
Symlinks Properties	This flag controls the accessibility of the symlinks under the shared directory. Possible values: enable, hide, read_only.
ACL	Access Control List (ACL) for the CIFS share. Each list entry is a slash-separated (/) pair of the user or group name and the permissions. Possible effective permissions: Full Control, Change, Read, or No Access. For example, for the Global group with Change permissions, the value would be: Global/Change.

NetApp Disk Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Disk Summary

Table 9-139 Disk Summary

Column name	Description
Name	The Disk name
Storage Array	The name of the NetApp storage system
Aggregate	The name of the associated aggregate. See " NetApp Aggregate Detail " on page 368.
Disk Type	Disk Type: ATA, EATA, FCAL, LUN, SAS, SATA, SCSI, XATA, XSAS, or unknown.
Capacity	The Disk's total capacity
Vendor	Disk Vendor
Disk Model	Disk Model
RAID State	RAID state: broken, copy, partner, pending, present, reconstructing, spare, or zeroing
RAID Type	RAID type: pending, parity, dparity, or data
RAID Status	Possible status: copying, degraded, foreign, growing, initializing, invalid, ironing, mirror degraded, mirrored, needs check, noparity, normal, out-of-date, partial, raid0, raid4, raid_dp, reconstruct, resyncing, snapmirrored, verifying, unrecoverable
Serial #	Disk serial number
RAID Group	The RAID group to which the disk belongs
Firmware Revision	Disk firmware revision number

NetApp Cluster-Mode Disk Summary

Lists NetApp disk attributes in a tabular report, with drilldowns to more detailed information for related NetApp objects. Supported for ONTAP Cluster Mode arrays only.

Use this report to verify disk details, such as capacity, RAID container, and associated aggregates, enabling Unified Visibility in your environment. Use the scope selector to narrow the list to specific disk types and RAID status.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Disk Summary

Table 9-140 NetApp Cluster-Mode Disk Summary

Column name	Description
Name	Name of the disk.
Cluster	Name of the storage array. See “Array Capacity and Utilization (NetApp Cluster)” on page 214.
Node	The textual name of a node. See “NetApp Cluster-Mode Node Summary” on page 385.
Aggregate	Name of the NetApp aggregate. See “ NetApp Volumes at Risk ” on page 458.
Disk Type	Disk interface type. Possible values: ATA, BSAS, EATA, FCAL, LUN, MSATA, SAS, SATA, SCSI, SSD, XATA, XSAS, FSAS, unknown.
Capacity	Physical disk size. Values are stored as KiB in the database and rendered according to your user profile preferences.
Vendor	Vendor/manufacturer of the disk.

Table 9-140 NetApp Cluster-Mode Disk Summary *(continued)*

Column name	Description
Disk Model	Disk model string.
RAID Container Type	Type of overlying disk container. Not relevant if information is unavailable or if excluded by desired_attributes. Possible values: aggregate = Container is an aggregate. broken = Container is broken pool. foreign = Array LUN has been marked foreign. labelmaint = Container is online label maintenance list. maintenance = Container is disk maintenance center. spare = Container is spare pool. unassigned = Disk ownership has not been assigned. unknown = Container is currently unknown. This is the default setting. volume = Container is a traditional volume.
RAID Position	Position of the disk, relative to its container type. Omitted if excluded by desired-attributes. Possible values: copy = RAID group copy destination. data = RAID group data. dparity = RAID group diagonal parity disk. orphan = Disk is orphan of an aggregate or traditional volume. parity = RAID group parity disk. pending = Disk is pending addition to an aggregate or traditional volume. present = Disk is present in system. This is the default setting.
Serial #	Disk serial number. Maximum length of 129 characters.
RAID Group	Name of the RAID group.
Firmware Rev	Firmware revision of the disk. The format of the firmware revision will vary depending on the type of disk and its vendor. Omitted if excluded by desired-attributes.

NetApp Cluster-Mode FC Array Port Utilization

This reports NetApp Fibre Channel (FC) ports and metrics for arrays. Use this report to view associated cluster and capacity usage values.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp FC Array Port Utilization

Table 9-141 NetApp Cluster-Mode FC Array Port Utilization

Column name	Description
Cluster	The name of the storage array.
Vserver	The name of the Vserver.
Array Port	Server name.
Host	The name of the host to display in the reports.
Host Ports	WWN of the array port. Comma-separated string of the WWN of all the HBA ports associated with the array port.
Mount Point	The mount point of the file system that is using the port.
Filesystem	File system name.

Table 9-141 NetApp Cluster-Mode FC Array Port Utilization *(continued)*

Column name	Description
File System Type	File System type. When you know the file system type, you can determine the host's storage requirements and limitations. Examples of file system values: NTFS - Windows NT file system ext2, ext3 - Extended file systems 2 and 3 ufs - Unix file system vxfs - Veritas file system FAT - File Allocation Table hfs - Hierarchical file system (Apple) ctfs - Cooperative temporary file system devfs - Device file system hsfs - High Sierra file system jfs - Journaled file system jfs2 - Journaled file system 2 nfs - Network file system objfs - File system object sfs - Secure file system gpfs - General Parallel file system (IBM) proc - a virtual file system procfs - Process file system tmpfs - a virtual memory file system smbfs - Server Message Block file system Null - empty NA- not applicable
Capacity	Total capacity of the file system. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-141 NetApp Cluster-Mode FC Array Port Utilization *(continued)*

Column name	Description
Used Capacity	Total used capacity of the file system. The amount of the storage array being used by the file system. Values are stored as KiB in the database and rendered according to your user profile preferences
Physical Link State	Physical Status of the adapter. Possible values: link up, link down. This is only available for converged network adapters (CNA).

NetApp NFS Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp NFS Summary

or

Capacity Manager > Array Capacity & Utilization > Array Capacity & Utilization

Then, click a **# NFS Exports**link.

NetApp volumes that are exported via the NFS protocol are listed in this report. Using this report, you can identify who is using the NFS shares.

Table 9-142 NetApp NFS Summary

Column name	Description
Path Name	The path name of the file.
Volume	Name of the volume. See " NetApp Volume Details " on page 398.
QTree	Name of the associated QTree

Table 9-142 NetApp NFS Summary *(continued)*

Column name	Description
Storage Array	Links to the LUN Utilization Summary for NetApp 7-Mode. See " LUN Utilization Summary for NetApp 7-Mode " on page 355.
Host Type	Host NFS permission type: read-only, read-write, root
Root Host Name	The name of the root host
Host	IP address or host name of the host to which the file system is exported
Host Mount Point	The mount point of the host that is using the volume
Total Capacity	Total capacity associated with the path and volume
Used Capacity	Amount of capacity that is used by NFS clients
Free Capacity	Available capacity for NFS client use

NetApp Cluster-Mode NFS Summary

Lists NetApp NFS Exports for the associated array. Details provided in this table include the path, cluster (storage system), and Vserver. Use this report for Unified Visibility to view CIFS shares and their associated volumes. Cluster-Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode NFS Summary

Table 9-143 NetApp Cluster-Mode NFS Summary

Column name	Description
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Vserver	Name of the Vserver. See "NetApp Cluster-Mode Vserver Summary" on page 413.
Volume	The name of the volume. See " NetApp Cluster-Mode Volume Detail " on page 401.
# Qtrees	Number of Qtrees in the volume. See "NetApp Cluster-Mode QTree Summary" on page 390.
Path Name	Junction Path at which the cluster-mode volume is mounted.
Export Policy Name	Export policy name.

NetApp Cluster-Mode Node Summary

This report enables Unified Visibility with a high-level summary of the capacities and performance of clusters in your environment. Use this as a starting point for navigating NetApp Cluster environments. Note that all performance numbers are for the last 24 hours, as of the last data collection. Cluster Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Node Summary

Table 9-144 NetApp Cluster-Mode Node Summary

Column name	Description
Cluster Name	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Node Name	The textual name of a node.
# Disks	Number of disks in the node. See "NetApp Cluster-Mode Disk Summary" on page 379.
Capacity	
Internal	Sum of the file system total capacity of all the aggregates when the disk type value is LUN. Values are stored as KiB in the database and rendered according to your user profile preferences.
External	Sum of the file system total capacity of all the aggregates, when the disk type is not LUN. Values are stored as KiB in the database and rendered according to your user profile preferences.
Allocated	Sum of the file system used capacity of all the aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Sum of the file system available capacity of all the aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage	Usage = Allocated/(Internal + External) Total
Aggregates	
# Aggregates	Number of aggregates in the node. See "NetApp Cluster-Mode Aggregate Summary" on page 365.
Capacity	Capacity = Internal + External

Table 9-144 NetApp Cluster-Mode Node Summary *(continued)*

Column name	Description
Used	Sum of the file system used capacity of all the aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Sum of the file system available capacity of all the aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage	Usage = Used/(Internal + External) Total
Volumes	
# Volumes	Number of volumes in the node. See "NetApp Cluster-Mode Volume Summary" on page 396.
# Flex Volumes	Sum of all member nodes' thinly provisioned volumes (FlexVols). See "NetApp Cluster-Mode Volume Summary" on page 396.
# Infinite Volumes	Sum of all member infinite volumes in the node.
Capacity	Sum of the capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used	Sum of the used capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Sum of the available capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage	Usage = Used / Total Capacity.
LUNs	

Table 9-144 NetApp Cluster-Mode Node Summary *(continued)*

Column name	Description
# LUNs	Number of LUNs in the node.
Capacity	Sum of the allocated capacity of the logical units. Values are stored as KiB in the database and rendered according to your user profile preferences.
# Allocated	Number of allocated LUNs in the node.
Allocated Capacity	Sum of the total capacity of the logical units. Values are stored as KiB in the database and rendered according to your user profile preferences.
# UnAllocated	# UnAllocated = # Luns - # Allocated
UnAllocated Capacity	UnAllocated = Total Capacity - Allocated Capacity
Usage	Usage = Allocated Capacity / Total Capacity
Performance	
Avg Processor Busy	Average processor utilization across all processors in the system from last data collection cycle.
CPU Busy	Percentage of time one or more processors is busy in the system from last data collection cycle.
Sys Read Resp Time (ms)	System Read Response Time: Average response time for the Cluster in aggregate as of the last data collection poll. Includes all active Node read operations.
Sys Write Resp Time (ms)	System Write Response Time: Average response time for the Cluster in aggregate as of the last data collection poll. Includes all active Node write operations.
Read OPS	Read operations per second from the last data collection cycle.
Write OPS	Write operations per second from last data collection cycle.

Table 9-144 NetApp Cluster-Mode Node Summary *(continued)*

Column name	Description
CIFS OPS	CIFS operations per second from last data collection cycle.
NFS OPS	NFS operations per second from last data collection cycle.
FCP OPS	Fibre channel protocol (FCP) operations per second from the last data collection cycle.
iSCSI OPS	iSCSI operations per second from the last data collection cycle.

NetApp QTree Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp QTree Summary

A QTree is the mechanism used to partition a volume, enabling the management of storage by project, user, or group. Use this QTree Summary to view the quota management configurations for the QTrees within a storage system. Then, drill down to view how each volume is partitioned.

Table 9-145 NetApp QTree Summary

Column name	Description
QTree	Name of the QTree, or if the ID is 0, the name of the volume is displayed
Storage Array	Name of the associated NetApp storage system
ID	ID of the QTree, which is unique within a volume. An ID of 0 indicates the volume itself.

Table 9-145 NetApp QTree Summary *(continued)*

Column name	Description
Oplocks	Indicates if CIFS (Common Internet File System) opportunistic locks are enabled/disabled.
Security Style	Linux, NTFS, or mixed
Status	SnapMirror backup status: Normal, Snapvaulted, Snapmirrored, or Read-only (if a SnapMirror destination)
Volume	Name of the volume to which the QTree belongs. See " NetApp Snapshot Summary " on page 392.
Aggregate	Name of the aggregate to which the QTree belongs. See " NetApp Aggregate Summary " on page 364.
# of LUNs	The number of LUNs mapped to the QTree. See " LUN Utilization Summary for NetApp 7-Mode " on page 355.
# of Shares	The number of CIFS shares using this volume. See " NetApp QTree Summary " on page 389.
# of Exports	The number of NFS exports using this volume. See " NetApp NFS Summary " on page 383.

NetApp Cluster-Mode QTree Summary

Lists NetApp QTree attributes, associated configuration metrics, and counts, with access to additional details. Cluster Mode arrays only. Filter the report by security style (unix, ntfs, mixed) and/or status (snap vaulted, snap mirrored, normal, read only).

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode QTree Summary

Table 9-146 NetApp Cluster-Mode QTree Summary

Column name	Description
Qtree Name	Name of the QTree. If the QTree is a volume, this is the volume name.
Cluster	Name of the storage array. See “Array Capacity and Utilization (NetApp Cluster)” on page 214.
ID	ID of the QTree (unique within the volume). This ID is 0 if the QTree is the volume itself.
Oplocks	Indicates whether opportunistic locks are enabled on the QTree. Possible values: enabled, disabled.
Security Style	Security style of the QTree. Possible values: unix, ntfs, or mixed.
Status	Status of the QTree. Possible values: snapvaulted, snapmirrored, normal, and readonly.
Volume	The name of the volume. See “ NetApp Cluster-Mode Volume Detail ” on page 401.
Aggregate	Name of the NetApp aggregate. See “ NetApp Volumes at Risk ” on page 458.
Vserver	Name of the Vserver. When specified as part of a Vserver-create, this is the name of the Vserver. See “NetApp Cluster-Mode Vserver Summary” on page 413.

Table 9-146 NetApp Cluster-Mode QTree Summary *(continued)*

Column name	Description
# LUNs	Number of LUNs in the QTree.
# CIFS Shares	Number of CIFS Shares in the QTree.
# NFS Exports	Number of NFS Exports in the QTree.

NetApp Snapshot Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Snapshot Summary

NetApp snapshots use pointers to reference the blocks where the data is stored, thereby using a minimal amount of capacity. A snapshot captures a point-in-time image of the file system data. When the next snapshot is taken, the pointers will reference other blocks, as the file system data blocks have changed in the interval between snapshots. Therefore, the access time provides the time reference for the snapshot block usage. Snapshots provide a read-only source of data for backups to tape.

Table 9-147 NetApp Snapshot Summary

Column name	Description
Name	Name of the snapshot schedule.
Storage Array	Links to the LUN Utilization Summary for NetApp 7-Mode. See " LUN Utilization Summary for NetApp 7-Mode " on page 355.
Volume	Links to the NetApp Volume Details. See " NetApp Volume Details " on page 398.
Access Time	Time the snapshot was created. This value does not change when the snapshot is accessed.

Table 9-147 NetApp Snapshot Summary *(continued)*

Column name	Description
Total Blocks	Number of 1024-byte blocks in this snapshot
Cumulative Blocks	Number of 1024-byte blocks--cumulative for current and previous snapshots
Dependency	Lists the applications that depend on this snapshot: snapmirror, snapvault, dump, vclone, LUNs, or snaplock
Busy	Indicates if an application is using the snapshot
Cumulative % of Used Blocks	The % of blocks used by this snapshot and all recent snapshots
Used Blocks	Of all the blocks used on the volume, this is the % used by the snapshot
Cumulative % of Total Blocks	The % of blocks owned by this snapshot and all recent snapshots
% of Total Blocks	The % of all the volume's blocks that is owned by this snapshot

NetApp Cluster-Mode Snapshot Summary

Provides insight into NetApp snapshots capacity usage. Use this report for Unified Visibility to view snapshots and their associated volumes. Filter the report scope by busy or non-busy snapshots. Cluster Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Snapshot Summary

Table 9-148 NetApp Cluster-Mode Snapshot Summary

Column name	Description
Name	Name of the snapshot.
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Volume	The name of the volume. See " NetApp Cluster-Mode Volume Detail " on page 401.
Vserver	Name of the Vserver. See "NetApp Cluster-Mode Vserver Summary" on page 413.
Create Time	The volume access time when the snapshot was created, in seconds since Jan 1, 1970. This value will not change even if the snapshot is accessed.
Total Blocks	Number of 1024-byte blocks in the snapshot.
Cumulative Blocks	Cumulative total of 1024-byte blocks of this snapshot and previous snapshots.
Busy	Indicates if the snapshot is being used by an application. Possible values: yes, no.
Cumulative % of Used Blocks	Percentage of blocks owned by this snapshot and all of the more recent snapshots
% of Used Blocks	Percentage of blocks owned by this snapshot.
Cumulative % of Total Blocks	Percentage of blocks owned by this snapshot and all of the more recent snapshots.
% of Total Blocks	Percentage of blocks owned by this snapshot.
State	The state of the snapshot. Possible values: valid = snapshot is complete and consistent. invalid = namespace constituent snapshot is missing partial = One or more data constituent snapshots are missing. Only a snapshot on an Infinite Volume can have a state of partial or invalid. Default value = valid.

NetApp Volume Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Volume Summary

Each volume depends on its aggregate for its physical storage. Both Traditional and FlexVol volumes contain directories and files; they also can contain QTrees and LUNs.

Table 9-149 NetApp Volume Summary

Column name	Description
Volume	Name of the volume. See " NetApp Volume Details " on page 398.
Storage Array	The name of the NetApp physical storage for the volume
Aggr	The name of the associated aggregate
Usage	Percentage of the aggregate that is in use
Space Guarantee	The type of space that has been committed: **volume**, **file**, or **none**. Note that file or none indicates thin provisioning.
Type	**Flex** for flexible volumes, **Trad** for traditional volumes
RAID Status	raid0 (mirrored), raid4 (single disk for parity), or raid_dp (double-parity protection)
Usage	Mouse over the thermometer to view the volume's % available
Total Capacity	The volume's total capacity. If the volume is offline or restricted, 0 is displayed.
Allocated Size	Amount of volume space in use. If the volume is offline or restricted, 0 is displayed.

Table 9-149 NetApp Volume Summary *(continued)*

Column name	Description
Available Size	Amount of space available on the volume. If the volume is offline or restricted, 0 is displayed.
# QTrees	Links to the NetApp QTree Summary. See " NetApp QTree Summary " on page 389.
# LUNs	Links to the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
# Shares	The number of CIFS shares using this volume. See " NetApp QTree Summary " on page 389.
# Exports	The number of NFS exports using this volume. See " NetApp NFS Summary " on page 383.
# Snapshots	The number of snapshots using this volume. See " Array Group Details " on page 359.

NetApp Cluster-Mode Volume Summary

Lists NetApp Volume attributes with related capacity summaries, providing an overview of space utilization. View volume-relevant counts such as # CIFS Shares, # NFS Exports, and # Snapshots. Cluster-Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Volume Summary

Table 9-150 NetApp Cluster-Mode Volume Summary

Column name	Description
Volume	The name of the volume. See " NetApp Cluster-Mode Volume Detail " on page 401.
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Vserver	Name of the Vserver. See "NetApp Cluster-Mode Vserver Summary" on page 413.
Aggregate	Name of the NetApp aggregate. See " NetApp Cluster-Mode Aggregate Detail " on page 370.
% Usage	% Usage = File System Used / File System Total
Space Guarantee	Storage guarantee option associated with the flexible volume. Possible values: none, volume (default option), file.
Type	The type of the volume. Possible values: rw =read-write volume ls = load-sharing volume dp = data-protection volume dc = data-cache volume (FlexCache) tmp = temporary volume (not valid for create) Default value is rw.
RAID Status	RAID status.
Style	The style of the volume. Possible values: flex = FlexVol infinitevol = Infinite Volume Default value is flex.

Table 9-150 NetApp Cluster-Mode Volume Summary *(continued)*

Column name	Description
State	Volume state. Possible values: online, restricted, offline, mixed. If an Infinite Volume's constituents are not all in the same state, this state is set to mixed. The mixed state is read-only.
Usage	Usage = Allocated Size/Total
Total	Total usable size (in bytes) of the volume, not including WAFL reserve or volume snapshot reserve. Valid only when the volume is online. Values are stored as KiB in the database and rendered according to your user profile preferences.
Allocated Size	Size (in bytes) that is used in the volume. Valid only when the volume is online. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available Size	Size (in bytes) that is still available in the volume. Valid only when the volume is online. Values are stored as KiB in the database and rendered according to your user profile preferences.
# QTrees	Number of QTrees in the volume. See "NetApp Cluster-Mode QTree Summary" on page 390.
# LUNs	Number of LUNs in the volume.
# CIFS Shares	Number of CIFS shares in the volume.
# NFS Exports	Number of NFS exports in the volume.
# Snapshots	Number of snapshots in the volume.

NetApp Volume Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports

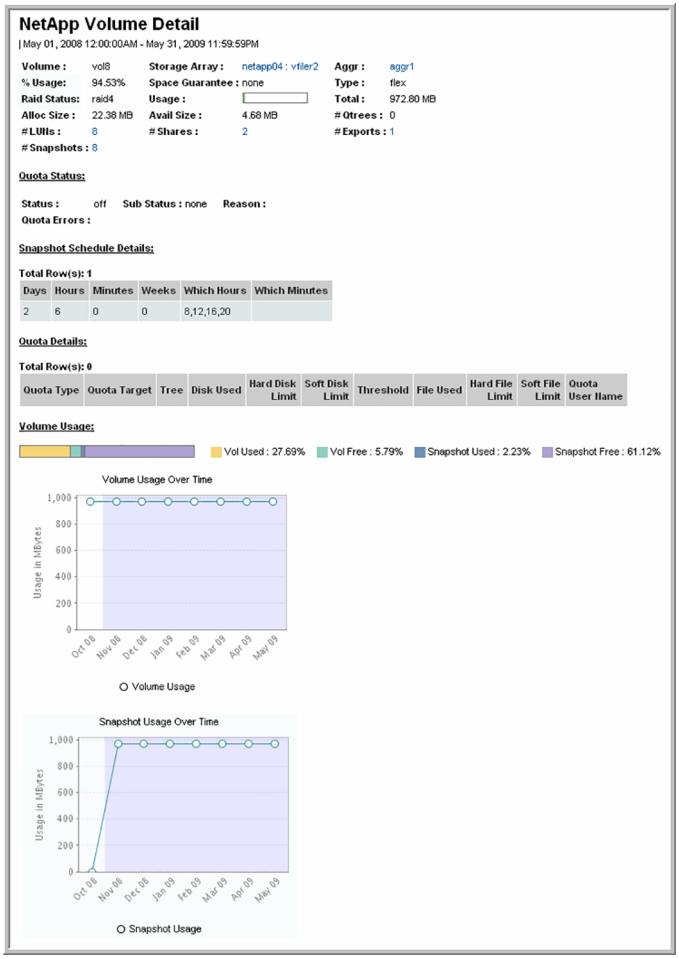
cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Volume Summary

Next, click **Volume Name**.



Volume Summary Status

See “ NetApp Snapshot Summary ” on page 392.

for a description of these fields.

Quota Status

Table 9-151 Quota Status

Column name	Description
Status	Indicates if quotas are enabled: **on** or **off**
Sub-Status	Minor quota status for the volume, valid only when the status is either resizing or initializing. The sub-status could be: scanning, etc scanning, setup, queue scan, done, or none.
Reason	If blank, there were no errors. Otherwise, this field displays the last quota error message.
Quota Errors	If blank, there were no errors. Otherwise, this field displays the collection of quota errors, including the reason code.

Snapshot Schedule Details

This section displays the volume’s current snapshot schedule.

Table 9-152 Snapshot Schedule Details

Column name	Description
Days	Number of daily snapshots taken
Hours	Number of hourly snapshots taken
Minutes	Number of snapshots taken each minute
Weeks	Number of weekly snapshots taken
Which Hours	A list of hours when the hourly snapshots are taken
Which Minutes	A list of minutes when the by-minutes snapshots are taken

Quota Details

Table 9-153 Quota Details

Column name	Description
Quota Type	User, group, or tree
Quota Target	Name, number, or path name, relevant to the quota type
Tree	Name of the QTree for this quota
Disk Used	Current amount of disk space; “-” indicates unlimited
Hard Disk Limit	The amount of disk space reserved for the target; “-” indicates unlimited.
Soft Disk Limit	Both an SNMP trap and a log message are generated when the target exceeds this limit; “-” indicates unlimited.
Threshold	When the target exceeds this amount of disk space, a message is logged
File Used	Current number of files used by the quota target
Hard File Limit	Number of files allowed for the target
Soft File Limit	Both an SNMP trap and a log message are generated when the target exceeds this limit; “-” indicates unlimited.
Quota User Name	Name of the User, Group, or SID (Windows security ID)

NetApp Cluster-Mode Volume Detail

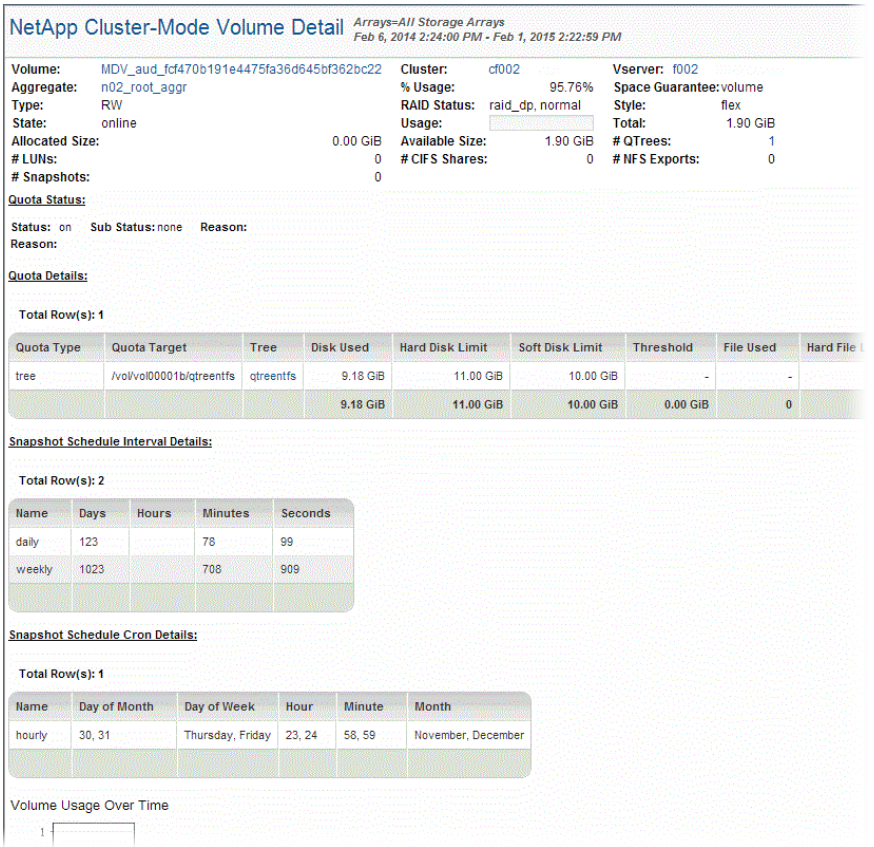
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Volume Summary

Next, click **Volume Name**.



NetApp Cluster-Mode Volume Detail
Arrays=All Storage Arrays
Feb 6, 2014 2:24:00 PM - Feb 1, 2015 2:22:59 PM

Volume:	MDV_aud_fcf470b191e4475fa36d645bf362bc22	**Cluster:**	cf002	**Vserver:**	f002
Aggregate:	n02_root_aggr	**% Usage:**	95.76%	**Space Guarantee:**	volume
Type:	RW	**RAID Status:**	raid_dp, normal	**Style:**	flex
State:	online	**Usage:**		**Total:**	1.90 GiB
Allocated Size:	0.00 GiB	**Available Size:**	1.90 GiB	**# QTrees:**	1
# LUNs:	0	**# CIFS Shares:**	0	**# NFS Exports:**	0
# Snapshots:	0				

Quota Status:

Status: on **Sub Status:** none **Reason:**
Reason:

Quota Details:

Total Row(s): 1

Quota Type	Quota Target	Tree	Disk Used	Hard Disk Limit	Soft Disk Limit	Threshold	File Used	Hard File
tree	/vol/vol00001b/qtreentfs	qtreentfs	9.18 GiB	11.00 GiB	10.00 GiB	-	-	
			9.18 GiB	**11.00 GiB**	**10.00 GiB**	**0.00 GiB**	**0**	

Snapshot Schedule Interval Details:

Total Row(s): 2

Name	Days	Hours	Minutes	Seconds
daily	123		78	99
weekly	1023		708	909

Snapshot Schedule Cron Details:

Total Row(s): 1

Name	Day of Month	Day of Week	Hour	Minute	Month
hourly	30, 31	Thursday, Friday	23, 24	58, 59	November, December

Volume Usage Over Time

1

Quota Status

Table 9-154 Quota Status

Column name	Description
Status	Indicates if quotas are enabled: **on** or **off**
Sub-Status	Minor quota status for the volume, valid only when the status is either resizing or initializing. The sub-status could be: scanning, etc scanning, setup, queue scan, done, or none.
Reason	If blank, there were no errors. Otherwise, this field displays the last quota error message.
Reason	If blank, there were no errors. Otherwise, this field displays the collection of quota errors, including the reason code.

Quota Details

Table 9-155 Quota Details

Column name	Description
Quota Type	User, group, or tree
Quota Target	Name, number, or path name, relevant to the quota type
Tree	Name of the QTree for this quota
Disk Used	Current amount of disk space; “-” indicates unlimited
Hard Disk Limit	The amount of disk space reserved for the target; “-” indicates unlimited.
Soft Disk Limit	Both an SNMP trap and a log message are generated when the target exceeds this limit; “-” indicates unlimited.
Threshold	When the target exceeds this amount of disk space, a message is logged
File Used	Current number of files used by the quota target

Table 9-155 Quota Details *(continued)*

Column name	Description
Hard File Limit	Number of files allowed for the target
Soft File Limit	Both an SNMP trap and a log message are generated when the target exceeds this limit; "-" indicates unlimited.
Quota User Name	Name of the User, Group, or SID (Windows security ID)

Snapshot Schedule Interval Details

This section displays the volume's current snapshot schedule.

Table 9-156 Snapshot Schedule Interval Details

Column name	Description
Name	Name of the schedule.
Days	Number of daily snapshots taken.
Hours	Number of hourly snapshots taken.
Minutes	Number of snapshots taken each minute.
Seconds	Number of snapshots taken each second.

Snapshot Schedule Cron Details

This section displays the cron details for current snapshot schedule.

Table 9-157 Snapshot Schedule Cron Details

Column name	Description
Name	Name of the cron schedule.
Day of Month	Day of the month scheduled to run.
Day of Week	Day of the week scheduled to run.
Hour	Hour the snapshot is scheduled to run.
Minute	Minute the snapshot is scheduled to run.

Table 9-157 Snapshot Schedule Cron Details *(continued)*

Column name	Description
Month	Month the snapshot is scheduled to run.

NetApp Cluster-Mode Summary

This report enables Unified Visibility with a high-level summary of the capacities and performance of clusters in your environment. Use this as a starting point for navigating NetApp Cluster environments. Note that all performance numbers are for the last 24 hours, as of the last data collection. Cluster Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Summary

Table 9-158 NetApp Cluster-Mode Summary

Column name	Description
Cluster Name	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Remote Cluster Name	The current name of the peer cluster.
Associated Objects	
# Nodes	Number of configured nodes in the cluster. See "NetApp Cluster-Mode Node Summary" on page 385.
# Vservers	Number of configured Vservers in the cluster. See "NetApp Cluster-Mode Vserver Summary" on page 413.
# FC Ports	Number of Fibre Channel Ports in the array. See "NetApp Cluster-Mode FC Array Port Utilization" on page 380.

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
# iSCSI Ports	Number of iSCSI Ports in the array. Drill down to NetApp Cluster-Mode Array iSCSI Port Utilization.
# CIFS Shares	Number of CIFS in the cluster. See " NetApp CIFS Summary " on page 374.
# NFS Exports	Number of NFS in the cluster. See "NetApp Cluster-Mode NFS Summary" on page 384.
# Disks	Number of physical disks in the array. See "NetApp Cluster-Mode Disk Summary" on page 379.
# Hosts	Total number of connected hosts via all configured nodes.
Capacity	

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
Internal	Total Capacity after RAID overhead. Usable capacity refers to capacity that has already been carved into RAID groups. This is the physical, usable capacity after RAID overhead. Dell Compellent: Sum of LUN capacity. EMC CLARiiON: Sum of usable capacity of all Array Groups and usable capacity of all storage pools. EMC Symmetrix: Sum of pool usable capacity + sum of thick LUN capacity. HP EVA, IBM SVC: Sum of allocated capacity of all Array Groups HP 3PAR: Sum of allocated capacity of all Array Groups. The total capacity for all Logical Drives (LDs). EMC VNX: If the embedded CLARiiON is not collecting, this usable internal capacity is the total capacity of all Disk Volumes. HDS: Sum of the usable capacity of all Array Groups. HDS H-NAS: Total capacity of all System Drives. IBM Enterprise: Total capacity of all Ranks. An IBM Rank is logical storage, enabling storage provisioning that is not constrained by the size of an array. One or more arrays comprise an IBM Rank, forming logically contiguous storage. Note that IBM often refers to Arrays as RAID Arrays. IBM Subsystem: Total capacity of all Array Groups. IBM XIV: Configure capacity of the array. NetApp: If raw capacity is greater than 0, this will value is the total capacity of all Aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
External	Total Capacity from back-end, external volumes. Usable capacity refers to capacity that has already been carved into RAID groups. This usable external capacity is the total capacity of all the Array Groups, where the array's is_external flag is set to Y. Derived from daisy-chained external storage (back-end systems), this is array group capacity that does not have physical devices (PDEVs) in this storage array. This is the capacity that is allocated to a front-end array. For EMC VNX (CLARiiON) arrays, if the embedded CLARiiON is in use, this value represents the total capacity of all disk volumes. For NetApp, if this is a vFiler, this value represents the total capacity of all volumes. If the raw capacity is 0, this will be the total capacity of all Aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
Allocated	Total Capacity for array groups and storage pools in the array that are carved RAID. Dell Compellent: Sum of LUN active and replay capacity. EMC CLARiiON: Sum of usable capacity of all Array Groups and consumed capacity of all Storage Pools, minus the free capacity of all Array Groups. EMC Symmetrix: Sum of (pool usable capacity - pool free capacity) + sum of thick LUN capacity HP 3PAR, HP EVA, IBM Enterprise, IBM Subsystem: Sum of the total capacity of all Allocated LUNs. EMC VNX: Used capacity of all disk volumes. HDS: (logical capacity + external capacity - external free capacity) of all Array Groups - virtualized capacity. HDS H-NAS: Used capacity of all system drives. IBM SVC: Sum of total capacity of all Allocated LUNs - virtualized capacity. IBM XIV: Allocated capacity of the array. NetApp: If a vFiler, this is the used capacity of all volumes; otherwise, this is the used capacity of all Aggregates, minus the virtualized capacity. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Available = Usable Free Capacity + Available Thin Pool Capacity
Usage	Usage = Allocated/(Internal + External)
# Aggr	Number of aggregates in the cluster.
Capacity	Total Capacity = Internal + External
Used	Sum of the file system used capacity of all the aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Sum of the file system available capacity of all the aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
Usage	Usage = Used/Available
Volumes	
# Volumes	Number of volumes in the storage system.
# Flexes	Number of flex volumes in the storage system.
# Inf	Number of infinite volumes in the storage system.
Capacity	Sum of the capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used	Sum of the used capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Available	Sum of the available capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage	Usage = Used/Capacity
LUNs	
# LUNs	Number of logical units (LUNs) in the storage array. Hosts can write to these LUNs.
Capacity	Total capacity of the logical units in the array. Values are stored as KiB in the database and rendered according to your user profile preferences.
# Allocated	Number of allocated logical units (LUNs) in the storage array. Sum of the capacity of all the allocated LUNs. If a LUN has a storage path to a host, it is an allocated LUN. These LUNs typically are configured as volumes on the host and are dedicated to an application.

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
Alloc Capacity	Total capacity of the logical units in the array that have paths assigned. Values are stored as KiB in the database and rendered according to your user profile preferences.
# Unallocated	# Unallocated = # LUNs - # Allocated
Unalloc Capacity	Unalloc Capacity = Capacity - Alloc Capacity
Usage	Usage = Alloc Capacity/Capacity
Performance	
Avg Processor Busy	Average processor utilization across all processors in the system from last data collection cycle.
CPU Busy	Percentage of time one or more processors is busy in the system from last data collection cycle.
Sys Read Resp Time (ms)	System Read Response Time: Average response time for the Cluster in aggregate as of the last data collection poll. Includes all active Node read operations.
Sys Write Resp Time (ms)	System Write Response Time: Average response time for the Cluster in aggregate as of the last data collection poll. Includes all active Node write operations.
Read OPS	Read operations per second from the last data collection cycle.
Write OPS	Write operations per second from last data collection cycle.
CIFS OPS	CIFS operations per second from last data collection cycle.
NFS OPS	NFS operations per second from last data collection cycle.
FCP OPS	Fibre channel protocol (FCP) operations per second from the last data collection cycle.

Table 9-158 NetApp Cluster-Mode Summary *(continued)*

Column name	Description
iSCSI OPS	iSCSI operations per second from the last data collection cycle.

NetApp Cluster-Mode Thin Provisioning Summary

This report lists NetApp Cluster-Mode aggregate and the net volume capacities of associated volumes relative to the total aggregate size. Use this report to identify over-provisioned aggregates. Cluster Mode arrays only.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Thin Provisioning Summary

Table 9-159 NetApp Cluster-Mode Thin Provisioning Summary

Column name	Description
Cluster	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Aggregate	Name of the NetApp aggregates.
Capacity	Total size (in bytes) of the file system. If the file system is restricted or offline, the value is 0.
# Thin Volumes	The number of thin volumes in the aggregate. See "NetApp Cluster-Mode Volume Summary" on page 396.
Volume Total Capacity	Sum of the capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-159 NetApp Cluster-Mode Thin Provisioning Summary *(continued)*

Column name	Description
Volume Used Capacity	Sum of the used capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Volume Free Capacity	Sum of the available capacity of all the volumes. Values are stored as KiB in the database and rendered according to your user profile preferences.
Aggr Free Space	Aggregate Free Space = Capacity - Volume Used Capacity.
Utilization	Utilization over 100% represents over-committed provisioning.

NetApp Cluster-Mode Vserver Summary

This report identifies the current configuration of Vservers in your environment, including aggregates associated with Vservers, as of the last data collection cycle. Vservers with configuration metrics are listed in a tabular report. Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > NetApp Cluster-Mode Vserver Summary

Supported for ONTAP Cluster-Mode arrays only.

Table 9-160 NetApp Cluster-Mode Vserver Summary

Column name	Description
Cluster Name	Name of the storage array. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.

Table 9-160 NetApp Cluster-Mode Vserver Summary *(continued)*

Column name	Description
Vserver Name	Name of the Vserver.
Status	State of the Vserver. Indicates the data-serving ability of a Vserver; therefore it is applicable only for Data Vservers. Possible values: running, stopped, starting, stopping, initializing, deleting.
Agg List	List of aggregate names for the Vserver.
Allowed	Vserver protocols. Possible values: nfs = NFS protocol cifs = CIFS protocol fcp = FCP protocol iscsi = iSCSI protocol ndmp = NDMP protocol.
Antivirus	Antivirus on access policy for the Vserver. Deprecated NetApp field. This field has no effect on the operation of the Vserver.
Comment	Comments associated with the Vserver.
Disallowed	Protocols not allowed for this Vserver. Possible values: nfs = NFS protocol cifs = CIFS protocol fcp = FCP protocol iscsi = iSCSI protocol ndmp = NDMP protocol.
Inf Vol	Specifies if this Vserver is a Vserver with Infinite Volume. Prior to Data ONTAP 8.2.0, Vserver with Infinite Volume is called Repository Vserver and once a repository Vserver is created in a cluster, additional Vservers cannot be created. Likewise, a Repository Vserver cannot be created if non-repository Vservers already exist in a cluster. Default value is false.
Max Vols	Maximum number of volumes that can be created on the Vserver. When specified as part of a vserver-create, this is the maximum number of volumes that can be created on the Vserver.
Name Map	Name mapping switch configuration details for the Vserver. Possible values: file, ldap. Default value is file.

Table 9-160 NetApp Cluster-Mode Vserver Summary *(continued)*

Column name	Description
Name Server	Name server switch configuration details for the Vserver. Possible values: nis, file, ldap.
NIS Domain	NIS domain configuration details for the Vserver. Attributes: non-creatable, non-modifiable.
Quota	Quota policy for the Vserver. When specified as part of Vserver-create, this field represents the quota policy for the Vserver.
Root/Aggr	Name of the NetApp aggregate.
Root/Vol	The name of the volume.
Snap Policy	Default snapshot policy setting for all volumes of the Vserver. This policy will be assigned to all volumes created in this Vserver unless the volume create request explicitly provides a snapshot policy or the volume is modified later with a specific snapshot policy. A volume-level snapshot policy always overrides the default Vserver-wide snapshot policy.
Type	Type of Vserver. Possible values: data, admin, and node.
UUI	Universal unique identifier (UUID) for the Vserver.

NetApp StorageGrid Active Alerts Summary

This report provides a summarized view of all the active alerts for the NetApp StorageGrid. As a best practice, you are recommended to schedule this report to run daily as a report-based alert that sends out emails or uses SNMP trap to alert your administrators responsible to manage your environment. The report is located here:

Capacity Manager > **Array Capacity & Utilization** > **NetApp StorageGrid Active Alerts Summary**

Table 9-161 NetApp StorageGrid Active Alerts Summary

Column name	Description
Array	Unique name of the StorageGrid system.
Site	Name of the site.
Node	The grid node where this alert is raised.
Alert ID	Unique ID of the alert.
Alert	Name of the alert.
Description	Description of the alert condition.
Rule ID	Rule ID of the alert.
Inhibited	Indicates whether this alert has been suppressed by another alert.
Inhibited By	IDs of the alerts currently suppressing this alert.
Silenced	Indicates whether notifications for this alert are currently suppressed by an active silence.
Silenced By	Silences currently suppressing the alert.
Status	Status of the alert (Active/Suppressed/Unprocessed)
Start Time	The time the alert was triggered.
Monitor	The monitor of the alert.
Severity	The severity of the alert (Critical/Major/Minor).

NetApp StorageGrid Node Summary

Grid node is the basic building block of a StorageGrid system. A data center site is built as a series of interconnected grid nodes. Each type of grid nodes hosts a specific set of services that perform specialized operations. This report provides the summary of grid nodes associated with your data center site.

This report is located here:

Capacity Manager > **Array Capacity & Utilization** > **NetApp StorageGrid Node Summary**

Table 9-162 NetApp StorageGrid Node Summary

Column name	Description
Array	Unique name of the StorageGrid system.
Node	Name of the StorageGrid node.
IP Address	IP address of the grid node.
Type	Node type (Admin, Storage, Archive, and so on).
Status	Connectivity status of the node.
OID	Object ID of the node.
No. of services	Count of services running on the node.
Network Interface	Type of network interface running on the node.
Hardware Address	Hardware address of the node.
Duplex	Duplex value of the node.
Severity	Severity of alert raised on the node status.
Object Usable Space	The total usable space available to store the objects on the node (KB).
Object Used Space	The total space occupied by the objects on the node (KB).
Object Free Space	The total free space left for the objects on the node (KB).
Metadata Usable Space	The total usable space available for metadata on the node (KB).
Metadata Used Space	The total space occupied by the metadata on the node (KB).
Metadata Free Space	The total free space left for the metadata on the node (KB).
Total Logical	Total amount of storage space on all the logical volumes managed by the node.
Used Logical	Total amount of used space on all the logical volumes managed by the node.
Available Logical	Total amount of unused space on all the logical volumes managed by the node.

NetApp StorageGrid Site Summary

StorageGrid environment may consist of a single or multi-data center deployments with many sites across the globe. This report provides the summary of all StorageGrid sites that belong to the StorageGrid arrays. The report lists the number of nodes and the usable, used, and free storage space for each site.

This report is located here:

Capacity Manager > **Array Capacity & Utilization** > **NetApp StorageGrid Site Summary**

Table 9-163 NetApp StorageGrid Site Summary

Column name	Description
Array Name	Unique name of the StorageGrid system.
Site ID	Unique ID of the StorageGrid site.
Site Name	Unique name of the StorageGrid site.
Display Name	Display name of the StorageGrid site.
No. of Nodes	Count of nodes present on the site.
State	Connectivity status of the site.
Severity	The highest current alarm severity affecting the site or any of its descendants.
Object Usable Space	The total usable space available for the object in the site (KB).
Object Used Space	The total space occupied by the object in the site (KB).
Object Free Space	The total free space left for the object in the site (KB).
Metadata Usable Space	The total usable space available for metadata in the site (KB).
Metadata Used Space	The total space occupied by metadata in the site (KB).
Total Free Space	The total free space left for metadata in the site (KB).
Raw Total Space	Stores the total raw space for the site in KB. (Object Usable + Metadata Usable)
Raw Used Space	Stores the total raw used space for the site in KB. (Object Used + Metadata Used)

Table 9-163 NetApp StorageGrid Site Summary *(continued)*

Column name	Description
Raw Free Space	Stores the total raw free space for the site in KB. (Object Free + Metadata Free)

NetApp StorageGrid Tenant Summary

This report provides a summarized view of storage tenant accounts. View the used capacity and objects stored to evaluate the consumed space.

This report is located here:

Capacity Manager > **Array Capacity & Utilization** > **NetApp Storage Grid Tenant Summary**

Table 9-164 NetApp StorageGrid Tenant Summary

Column name	Description
Array ID	Unique ID of the StorageGrid system.
Array Name	Name of the storage array.
Tenant Id	Unique ID of the tenant.
Tenant Name	Name of the tenant.
Logical Used Capacity	Used capacity of the StorageGrid calculated in KiB.
# of Objects Stored	Number of objects currently stored in all namespaces owned by the tenant.
Management Capabilities	Indicates whether the tenant has management capabilities (Y/N).
S3 Capabilities	Indicates whether the tenant has S3 bucket management capabilities (Y/N).
Swift Capabilities	Indicates whether the tenant has Swift management capabilities (Y/N).
Use Account Identity Source	Indicates whether the tenant has the permission to use account identity source (Y/N).
Allow Platform Services	Indicates whether the tenant has the permission to use platform services (Y/N).

Table 9-164 NetApp StorageGrid Tenant Summary *(continued)*

Column name	Description
Users	Indicates the number of users. You can click this tenant user count to drill down and view user details in the NetApp StorageGrid Users Summary report.
Buckets	Indicates the number of managed S3 buckets. You can click this tenant bucket count to drill down and view S3 bucket details in the NetApp StorageGrid Buckets Summary report.

NetApp StorageGrid Users Summary

This report provides a summarized view of tenant. The report also appears as a drill down view when you click the tenant Users count in the NetApp StorageGrid Tenant Summary report.

This report is located here:

Capacity Manager > **Array Capacity & Utilization** > **NetApp StorageGrid Users Summary**

Table 9-165 NetApp StorageGrid Users Summary

Column name	Description
Array ID	The Array ID. An Array (also known as a RAID Array) is an Array Site that has been formatted to a specific RAID format.
Array Name	Name of the storage array.
Tenant Id	Unique ID of the tenant.
Unique Name	The machine-readable name of the user.
User Uniform Resource Name	Unique identifier of the resource.
UUID	Universal unique identifier of the user.
Full Name	Full name of the user.
Group Memberships	Names of host groups associated with the users.
Prevent Login	Indicates whether user login is restricted (Y/N).
Federated	Indicates whether the user has federated access.

NetApp StorageGrid Buckets Summary

This report provides a summarized view of tenant buckets. The report also appears as a drill down view when you click the tenant Buckets count in the NetApp StorageGrid Tenant Summary report.

This report is located here:

Capacity Manager > **Array Capacity & Utilization** > **NetApp StorageGrid Buckets Summary**

Table 9-166 NetApp StorageGrid Buckets Summary

Column name	Description
Array ID	The Array ID. An Array (also known as a RAID Array) is an Array Site that has been formatted to a specific RAID format.
Tenant Id	Unique ID of the tenant.
Bucket Name	S3 bucket name.
Creation Time	S3 bucket creation time.
Region	Region assigned to the S3 bucket.
Object Count	Count of objects per storage array.
Logical Size	Logical size of the S3 bucket.
Legal Hold	Indicates whether the S3 bucket is enabled for legal hold.
Retention Period	Number of days S3 bucket will hold the data.
Last Access Time	The last access time of the S3 bucket.
Consistency	S3 bucket or Swift container consistency setting.

OpenStack Swift Container Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > OpenStack Swift Container Summary

List OpenStack Swift containers to identify used capacity, quotas, and relevant policies. Drill down to details associated with storage utilization by account.

Table 9-167 OpenStack Swift Container Summary

Column name	Description
Cluster Name	Name of the cluster (storage array)
Name	The name of the container. The same name can be used for containers in different accounts.
Account	The unique ID of the OpenStack Swift account. An account is also known as the project or tenant. An account is actually a storage area.
# Objects	The number of objects, as of the last data collection cycle.
Used Capacity	Amount of space currently used by the container.
Quota	If the meta name for the container has a quota string, then this is the meta value. This is collected in bytes and converted to KiB.
Create Date	The date and time the OpenStack Swift container data was created in the NetBackup IT Analytics database.
Storage Policy	Name of the storage policy. Drill down to the OpenStack Swift Policy Summary report.
Version	The value of this flag indicates the container where the versions are stored. Object versioning is implemented by setting a flag on the container to tell OpenStack Swift to version all objects in the container.

Table 9-167 OpenStack Swift Container Summary *(continued)*

Column name	Description
Sync to	Indicates the cluster to which the container is syncing. OpenStack Swift can be configured so that the contents of a container will be mirrored to another container via background synchronization. Clusters can allow/accept sync requests to/from other clusters. This requires a configuration that specifies where to sync the container, along with a secret synchronization key.

OpenStack Swift Policy Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > OpenStack Swift Policy Summary

List OpenStack Swift policies associated with accounts and containers. Capacity metrics are collected for each policy.

Table 9-168 OpenStack Swift Policy Summary

Column name	Description
Cluster Name	Name of the storage array.
Account Name	The unique ID of the OpenStack Swift account. An account is also known as the project or tenant. An account is actually a storage area.
Policy Name	Name of the storage policy.
# Containers	Per-policy container count.
Used	Amount of space currently used by the storage policy.

Table 9-168 OpenStack Swift Policy Summary *(continued)*

Column name	Description
# Objects	Per-policy object count.

OpenStack Swift Storage Utilization by Account

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > OpenStack Swift Storage Utilization by Account

View OpenStack Swift storage capacity usage by cluster account, including the account's quota, number of objects, policies, and containers.

Drill down to a policy or container summary for additional related details.

Table 9-169 OpenStack Swift Storage Utilization by Account

Column name	Description
Cluster Name	Name of the storage array.
Account Identifier	The unique ID of the OpenStack Swift account. An account is also known as the project or tenant. An account is actually a storage area.
# Policies	The number of policies associated with the OpenStack Swift account.
# Containers	The number of containers in the OpenStack Swift account.
Usage Capacity	Amount of space currently used by the object store.
Quota	If the meta name for the account has a quota string, then this is the meta value. This is collected in bytes and converted to KiB.

Table 9-169 OpenStack Swift Storage Utilization by Account *(continued)*

Column name	Description
# Objects	The number of objects in the OpenStack Swift account.
Create Date	The time and date, in UNIX Epoch timestamp format, when the account was initially created.
Tenant	Name of the tenant/project.

OpenStack Swift Storage Utilization by Zone

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > OpenStack Swift Storage Utilization by Zone

List OpenStack Swift zones to view raw capacity for storage nodes. Quickly identify storage usage by zone.

Table 9-170 OpenStack Swift Storage Utilization by Zone

Column name	Description
Cluster Name	Name of the storage array.
Storage Node	IP of the storage/proxy node.
Zone #	Zones are used to isolate data and to define device failure boundaries. A zone can be a single drive or multiple racks, on a larger scale.
Raw Total	Physical disk or partition capacity from the host. Values are stored as KiB in the database and rendered according to your user profile preferences.

OpenStack Swift Storage Device Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > OpenStack Swift Storage Device Summary

View physical device information for a storage node. Identify drive size and the name of the OpenStack zone of which the device is a member.

Table 9-171 OpenStack Swift Storage Device Summary

Column name	Description
Cluster Name	Name of the storage array. Drill down to the Array Capacity & Utilization report for this array/cluster.
Storage Node	IP address of the storage/proxy node. Drill down to the OpenStack Swift Storage Node Summary.
Host Name	The name is derived from the host's DNS name. Drill down to the Host Utilization Detail report.
Device Name	Name of the device. Example: sda. A device is a drive, as it is seen by the OpenStack Swift system.
Zone	Zones are used to isolate data and to define device failure boundaries. A zone can be a single drive or multiple racks, on a larger scale. Drill down to the OpenStack Swift Storage Utilization by Zone report.
Vendor	Vendor name.
# LUNs	Total number of LUNs. Drill down to the LUN Utilization Summary report.
Size	Size in KiB.
Raw Disk Path	Raw disk path.

Table 9-171 OpenStack Swift Storage Device Summary *(continued)*

Column name	Description
Storage Type	The type of storage, such as DAS, SAN, or NAS.

OpenStack Swift Storage Utilization by Node

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Array Capacity & Utilization > OpenStack Swift Storage Utilization by Node

List OpenStack Swift storage nodes with capacity and performance statistics. Drill down to details associated with storage utilization by node.

Table 9-172 OpenStack Swift Storage Utilization by Node

Column name	Description
Cluster	Name of the storage array. Drill down to the Array Capacity & Utilization report.
Storage Node	IP of the storage/proxy node.
# Devices	Number of devices associated with the node. Drill down to the OpenStack Swift Storage Device Summary report.
# Zones	Number of zones. Drill down to the OpenStack Swift Storage Utilization by Zone report.
Raw Total	Physical disk or partition capacity from the host. Values are stored as KiB in the database and rendered according to your user profile preferences.
Logical Total	Total capacity of the file system. Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 9-172 OpenStack Swift Storage Utilization by Node *(continued)*

Column name	Description
Logical Used	Used capacity of the file system. Values are stored as KiB in the database and rendered according to your user profile preferences.
Usage	Logical Used / Logical Total
Avg Resp Time	Average I/O response time, represented in milliseconds. Average operation response time for all storage nodes for the latest collection cycle. This is the weighted average of all devices, based on I/O operations per second for the same period.
Avg IO/sec	Average total I/O transactions per second. This is based on the average I/O transactions per second, for all storage nodes for the latest collection cycle.
CPU Usage	CPU utilization of all processors in all storage nodes.

Hitachi Vantara Efficiency Summary

NetBackup IT Analytics allows to explore the data center using customizable report templates or by using parts of infrastructure as entry points.

Use **Search** to find reports, templates, and dashboards across the portal.

Use the **Reports** tab to examine the catalog of templates, dashboards, and reports which are organized by the products along with user-created, and system folders.

Navigate to *Reports* > *Capacity Manager* > *Array Capacity & Utilization* and then click *Hitachi Vantara Efficiency Summary*. Hitachi Vantara Efficiency Summary report is displayed.

This report displays the summarized view of the Hitachi Vantara operational efficiency.

Note: User can define the scope of the array and configure the report parameters using filters.

Table 9-173 Hitachi Vantara Efficiency Summary report details

Column name	Description
Storage Device Id	Storage Array Id of APS_STORAGE_ARRAY table.
Storage Array Name	Name of the storage system.
Total Efficiency Ratio	The total efficiency of the entire storage system.
Data Reduction Ratio	The efficiency of capacity saving performed by using the capacity saving function (dedupe and compression) or accelerated compression.
Snapshot Efficiency Ratio	The efficiency of capacity saving performed by using snapshots to back up data.
Provisioning Efficiency Percentage	The percentage of saving efficiency for consumption of capacity achieved by using Dynamic Provisioning to virtualize capacity.
Calculation Start Time	The date and time when the calculation of the total efficiency began (UTC), format(YYYY-MM-DDThh:mm:ssZ).
Calculation End Time	The date and time when the calculation of the total efficiency ended (UTC), format(YYYY-MM-DDThh:mm:ssZ).
Software Saving Ratio	Total efficiency of capacity saving performed by using the capacity saving function.
Software Compression Ratio	Efficiency of capacity saving performed by using compression.
Software Deduplication Ratio	Efficiency of capacity saving performed by using deduplication.
Software Pattern Matching Ratio	Efficiency of capacity saving performed by reclaiming the specified data pattern.
FMD Pattern Matching Ratio	The efficiency of capacity saving performed by using accelerated compression are displayed.
FMD Saving Ratio	
FMD Compression Ratio	
Creation Date	Created date
Last Updated	Last updated date

Chapter 10

Available/Reclaimable Capacity Reports

This chapter includes the following topics:

- Allocated but Undiscovered LUNs
- Allocated but Unused LUNs
- Array Utilization Dashboard
- Array Utilization Summary
- Reclamation Summary
- Reclaimable Space by Host Group
- Reclaimable Space by Host Group Details
- Reclaimable Space by Host
- LUNS Added to Volume Group but Unused
- Suspected De-provisioned Hosts
- Unallocated LUNs
- Unused Disks/Volume Groups
- Unused Partitions/Logical Volumes

Allocated but Undiscovered LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

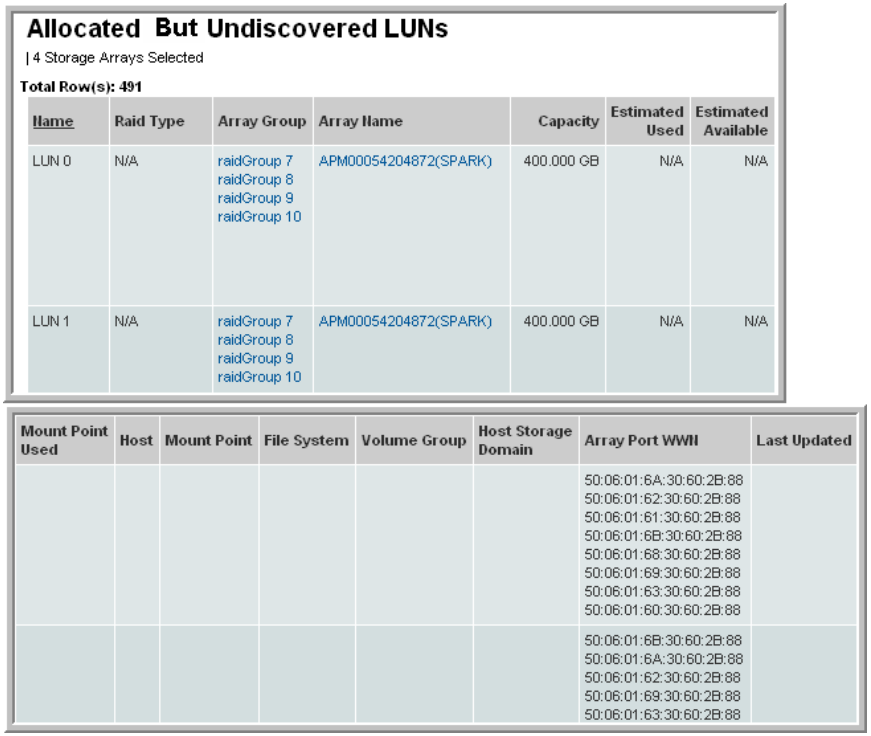
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage> Allocated but Undiscovered LUNs

Often referred to as orphaned LUNs, these are LUNs that have been assigned to a host, but they have not been discovered and therefore are not seen on the host side. The array has created the LUN and the storage path exists, but there is no access to host port and host information, such as disks, partitions, and volume groups. To make these LUNs visible to the host, a system administrator must execute a LUN scan or, for older OS and volume manager versions, re-boot the system. Otherwise, this storage is not available for mounting filesystems. The LUN scan tells the HBA (host bus adapter) that a LUN has been zoned to the host.

For HP 3PAR: LUNs refer to Virtual Volumes and Volume Groups refer to Common Provisioning Groups.

Note: The list of Allocated, but unused LUNs includes LUNs that have been allocated, discovered, but not yet mounted. The list of Allocated, but undiscovered LUNs includes LUNs hat have been allocated, but undiscovered.



Allocated But Undiscovered LUNs

| 4 Storage Arrays Selected

Total Row(s): 491

Name	Raid Type	Array Group	Array Name	Capacity	Estimated Used	Estimated Available
LUN 0	N/A	raidGroup 7 raidGroup 8 raidGroup 9 raidGroup 10	APM00054204872(SPARK)	400.000 GB	N/A	N/A
LUN 1	N/A	raidGroup 7 raidGroup 8 raidGroup 9 raidGroup 10	APM00054204872(SPARK)	400.000 GB	N/A	N/A

Mount Point Used	Host	Mount Point	File System	Volume Group	Host Storage Domain	Array Port WWN	Last Updated
						50:06:01:6A:30:60:2B:88 50:06:01:62:30:60:2B:88 50:06:01:61:30:60:2B:88 50:06:01:6B:30:60:2B:88 50:06:01:68:30:60:2B:88 50:06:01:69:30:60:2B:88 50:06:01:63:30:60:2B:88 50:06:01:60:30:60:2B:88	
						50:06:01:6B:30:60:2B:88 50:06:01:6A:30:60:2B:88 50:06:01:62:30:60:2B:88 50:06:01:69:30:60:2B:88 50:06:01:63:30:60:2B:88	

Additional details can be accessed via the following drill-down links in the report:

- Array Group
  See “ Array Group Details ” on page 359.
- Array Name
  See “ Array Capacity and Utilization ” on page 164.

For descriptions of the displayed fields, refer to the following report.

See “ LUN Utilization Summary ” on page 344.

Allocated but Unused LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Allocated but Unused LUNs

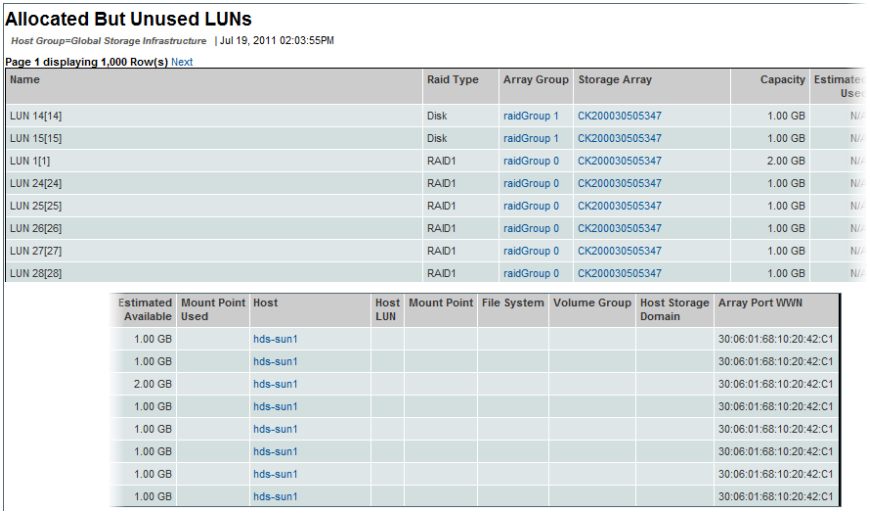
These are LUNs that have been assigned to a host, but the host has not been placed into a volume group, or a partition has not been created. The array has created the LUN and the storage path exists, making the host and host port information available to Capacity Manager. However, the host storage details--disks, partitions, and volume groups--are not available for the host associated with the storage path.

For HP 3PAR: LUNs refer to Virtual Volumes and Volume Groups refer to Common Provisioning Groups.

Note: The list of Allocated, but unused LUNs includes LUNs that have been allocated, discovered, but not yet mounted. The list of Allocated, but undiscovered LUNs includes LUNs hat have been allocated, but undiscovered.

Table 10-1 Allocated but Unused LUNs

Column name	Description
R1C1	R1C2



Additional details can be accessed via the following drill-down links in the report:

- Array Group
  See “ Array Group Details ” on page 359.

- Array Name
  See “ Array Capacity and Utilization ” on page 164.
- Host
  See “ Host Utilization Detail ” on page 481.

Table 10-2 Alloccated but Unused LUNs

Column name	Description
Name	Name of the LUN that has been allocated, but unused. For HP 3PAR: LUN refers to Virtual Volume (VV).
RAID Type	The RAID level on the LUN. When the RAID type is Disk, there is no RAID protection. For HP 3PAR: LUN refers to Virtual Volume (VV).
Array Group	The group to which the array belongs.
Storage Array	The storage array for this LUN. For HP 3PAR: LUN refers to Virtual Volume (VV).
Capacity	The total capacity (used and unused) of the LUN. For HP 3PAR: LUN refers to Virtual Volume (VV).
Estimated Used	LUN estimated used. For HP 3PAR: LUN refers to Virtual Volume (VV).
Estimated Available	Estimated available LUN storage. For HP 3PAR: LUN refers to Virtual Volume (VV).
Mount Point Used	The amount of storage used by each mount point.
Host	The name of the host to which the LUN was allocated. For HP 3PAR: LUN refers to Virtual Volume (VV).
Host LUN	The host LUN. For HP 3PAR: LUN refers to Virtual Volume (VV).
Mount Point	The host’s filesystem mount point for this LUN. It is possible to have zero mount points and filesystems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a filesystem and mounting it. For HP 3PAR: LUN refers to Virtual Volume (VV).
Filesystem	The host’s filesystem or drive name to which the LUN maps. LUNs simply appear as disks to hosts. For HP 3PAR: LUN refers to Virtual Volume (VV).
Volume Group	The logical collection or volume group to which the volume belongs. For HP 3PAR: Volume group refers to Common Provisioning Group (CPG).

Table 10-2 Alloccated but Unused LUNs *(continued)*

Column name	Description
Host Storage Domain	The name of the host storage domain, a logical group of LUNs that can be accessed by multiple hosts through the same physical storage port, with each host accessing its unique LUN. For security purposes, LUN masking is also coupled with this capability. A host storage domain is Hitachi's name for multi-hosted storage capability. For HP 3PAR: LUN refers to Virtual Volume (VV).
Array Port WWN	Array port world wide name.

Array Utilization Dashboard

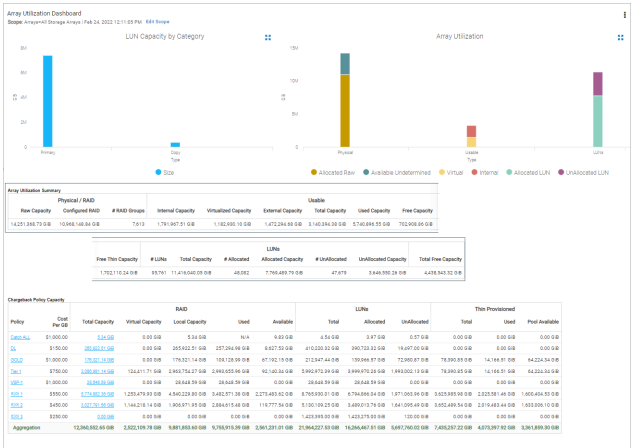
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Array Utilization Dashboard

This report provides array utilization statistics for storage arrays.

The Usable capacity values show a summary of overall array storage consumption from various viewpoints: virtual, internal, external, used, and free. Likewise, the chargeback policy list breaks out capacity usage to help you determine if you have full coverage for your chargeback policies. As a prerequisite for this report, capacity chargeback policies must be configured.



In the LUN Capacity by Category chart, the bars represent the following:

- Primary - The LUN capacity for the primary array. This value excludes open reserved LUNs.
- Copy - The LUN capacity for a copy stored on the primary array.
- Replication - The LUN capacity for a replication of the data on an external storage array.

Click the bars in this chart to launch the Array Utilization Summary.

See “ Array Utilization Summary ” on page 437.

In the Array Utilization bar chart, the bars link to the Array Utilization Summary displaying the capacity usage that corresponds to the chart legend colors.

See “ Array Utilization Summary ” on page 437.

- Physical: Allocated Raw - This varies by vendor. See the description in the Array Utilization Summary.
  See “ Array Utilization Summary ” on page 437.
- Physical: Available Undetermined - This value is undetermined because it represents raw capacity minus configured RAID capacity, but could include spares and failed drives that aren’t necessarily available for consumption.
- Usable: Virtual - This represents the capacity on the back-end array that has been given to the front-end array.
- Usable: Internal - This is the physical, usable capacity after RAID overhead. In the case of Symmetrix arrays, this is raw, usable capacity. For virtualized arrays (HDS, NetApp, and IBM SVC), this is the front-end internal storage.

- Allocated LUN - The total storage capacity for the allocated LUNs. For HDS arrays, it includes Dynamically Provisioned capacity.
- UnAllocated LUN - Mount of unallocated LUN capacity. For HDS arrays, this number includes promised, Dynamically Provisioned capacity, but no reserved space.

In the Array Utilization Summary table, the # Groups links to the List Array Groups. From there, you can access the Array Group Details.

See " List Array Groups " on page 239.

See " Array Group Details " on page 359.

For details about the Chargeback Policy Capacity table, refer to the Aggregated Chargeback report.

See " Aggregated Chargeback " on page 464.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Array Utilization Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Array Utilization Summary

This report can be used to identify the percentage of raw storage that has been allocated to RAID groups and, more importantly, the percentage of raw capacity that is available.

Table 10-3 Array Utilization Summary

Column name	Description
Storage Array	The name of the array.

Table 10-3 Array Utilization Summary *(continued)*

Column name	Description
Virtualized By	For virtualized storage arrays only, for example HDS AMS, VSP; Parent arrays are listed, accounting for capacity that is collected from both the front and back ends of a virtualized array (HDS, IBM SVC, and NetApp). Each array links to the Array Capacity & Utilization report.
Physical/RAID: Raw Capacity	For Dell Compellent arrays: Sum of the capacity of all physical disks. For EMC Symmetrix arrays: Sum of the capacity of all physical disks from all array groups--no spares. Note: When viewing raw capacity in EMC Control Center, disk group 0 (hot spares) is not included in the total capacity calculation--accounting for the difference between the EMC and Capacity Manager reported capacity values. For EMC CLARiiON arrays: Sum of the capacity of all physical disks from all array groups--including spares (drives that are not allocated to an array group) For HDS arrays: Sum of capacity of all PDEVs from all array groups. Capacity Manager assumes that all PDEVs are part of an array group, so PDEVs that are marked as--Not included in an array group--are included in the capacity calculation. That is, no disks are treated as spares. For NetApp V-Series and vFilers: External capacity and vFiler capacity are excluded. For HP 3PAR: Total capacity for all HP 3PAR disk drives associated with the 3PAR array.
Physical/RAID: Configured RAID	The amount of physical (raw) capacity that has been allocated to RAID groups. For Dell Compellent: Sum of (Total Capacity - Free Capacity) of the managed physical disk. For EMC Symmetrix: Sum of all the capacities of the PDEVs in the arrays that are assigned to the Array Group. NetApp: Raw capacity - Raw available capacity. For IBM XIV: Amount of raw, physical storage in the storage array.

Table 10-3 Array Utilization Summary *(continued)*

Column name	Description
Physical/RAID: # RAID Groups	The number of RAID groups for this array drills down to the List Array Groups. See " List Array Groups " on page 239. report. For NetApp storage, this refers to Aggregates. For HP 3PAR: Number of Common Provisioning Groups (CPGs). See " HP 3PAR Common Provisioning Group (CPG) Summary " on page 592.
Usable: Internal Capacity	This is the physical, usable capacity after RAID overhead. In the case of ■ For EMC Symmetrix: Raw, usable capacity. ■ For virtualized arrays (HDS, NetApp, and IBM SVC): This is the front-end internal storage. ■ For HP 3PAR: Internal capacity. The total capacity for all Logical Drives (LDs). ■ For NetApp Cluster-Mode: This value does not include spares.
Usable: Thick Virtualized Capacity	Represents the capacity on the back-end array that has been given to the front-end array.
Usable: Virtualized Capacity	Total Capacity for volumes (Includes thick and thin volumes) that are allocated for the virtualized array. For HDS, IBM SVC, and NetApp only. Total Capacity (includes thick and thin volumes) that is allocated to the virtualized array (front-end array). Values are stored as KiB in the database and rendered according to your user profile preferences.

Table 10-3 Array Utilization Summary *(continued)*

Column name	Description
Usable: External Capacity	Total Capacity from back end external volumes. Usable capacity refers to capacity that has already been carved into RAID groups. This usable external capacity is the total capacity of all the Array Groups, where the array's is_external flag is set to Y. Derived from daisy-chained external storage (back-end systems), this is array group capacity that does not have physical devices (PDEVs) in this storage array. This is the capacity that is allocated to a front-end array. If the raw capacity is 0, this will be the total capacity of all Aggregates. Values are stored as KiB in the database and rendered according to your user profile preferences. ■ For EMC VNX (CLARiiON): Total capacity of all disk volumes. ■ For NetApp V-series: Sum of aggregated capacity ■ For NetApp vFilers: Total capacity of all volumes ■ For HP 3PAR: Always blank. HP 3PAR does not virtualize storage. ■ For NetApp Cluster-Mode: Value does not include Unassigned.
Usable: Total Capacity	Total Capacity = Internal - Virtual + External Capacity For thin-provisioned capacity: Total Capacity = Used + Free + Free Thin Capacity For HP 3PAR: Internal capacity. The total capacity for all Logical Drives (LDs).
Usable: Used Capacity	Amount of internal RAID group capacity allocated to LUNs. The total capacity of physical disks that are associated with a RAID group, minus the total free capacity of the RAID groups; does not include virtualized capacity. For HP 3PAR: Total capacity of Logical Drives (LDs) where the LD is mapped to a virtual volume.
Usable: Free Thick Capacity	Available capacity of the array groups in the storage array; usable uncarved RAID capacity. For NetApp E-Series systems, Free Thick Capacity = Total Usable Capacity - Used Usable Capacity.

Table 10-3 Array Utilization Summary *(continued)*

Column name	Description
Usable: Free Thin Capacity	Dell Compellent: Sum of LUN capacity - sum of LUN active and replay capacity. EMC Symmetrix: Free Thin = Pool Capacity - Pool Used from Thin Pool Summary. EMC VNX (CLARiiON): Free Thin = Available Capacity from Thin Pool Summary. IBM Enterprise: Free Thin = N/A. Thin provisioning for DS8000 via Extent Space Efficient (ESE) Volumes and Track Space Efficient (TSE) Volumes not collected. Hitachi: Free Thin = DP pool capacity - used DP pool capacity. HP EVA: Free Thin = N/A, no thin provisioning.
LUNS: # LUNs	The number of LUNs that have been created on the array. Click on the Used link or the LUN link to go to the LUN Utilization Summary Report. For HP 3PAR: Total number of Virtual Volumes.
LUNS: Total Capacity	The total capacity for the LUNs. For HP 3PAR: Total capacities of all Virtual Volumes (VVs) (fully provisioned and thinly provisioned).
LUNs: # Allocated	The number of LUNs from which storage has been allocated. For HP 3PAR: Total number of Virtual Volumes (VVs).
LUNs: Allocated Capacity	The total storage capacity for the allocated LUNs. For HDS arrays, it includes Dynamically Provisioned capacity. For HP 3PAR: Total capacities of all Virtual Volumes (VVs) exported to hosts.
LUNs: # Unallocated	The number of LUNs that have yet to be allocated. These are LUNs that do not have a path to a host. For HDS arrays, this number includes promised Dynamically Provisioned capacity, but no reserved space. For HP 3PAR: Total number of Virtual Volumes (VVs) not exported to hosts.
LUNs: UnAllocated Capacity	Amount of unallocated LUN capacity. For HDS arrays, this number includes promised, Dynamically Provisioned capacity, but no reserved space. For HP 3PAR: Total capacities of all Virtual Volumes (VVs) not exported to hosts.
LUNs: Total Free Capacity	Total Free = Usable Free Capacity + Usable Free Thin Capacity + LUNs UnAllocated Capacity

Reclamation Summary

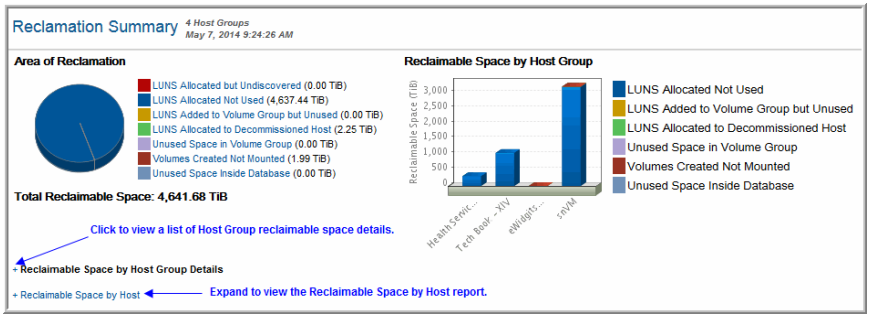
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Reclamation Summary

Within this single report, you can visualize all the possible types of storage that potentially could be reclaimed. The Areas of Reclamation charts provide drilldowns to relevant details.

A pie chart illustrates potential reclamation categories.

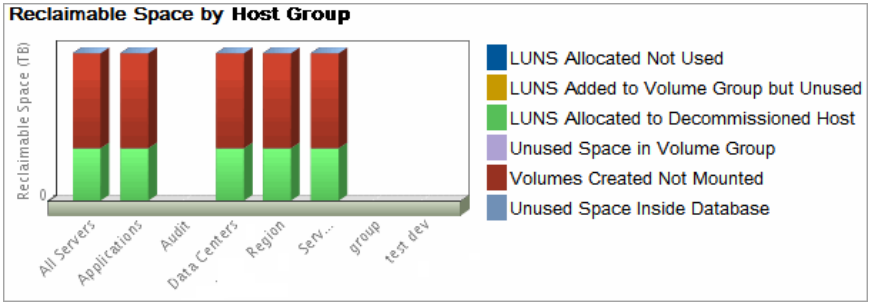


The LUNS allocated but undiscovered link takes you to the Allocated But Undiscovered LUNs report. All other pie chart links are host related and therefore are described in the following tables.

For HP 3PAR: LUNs refer to Virtual Volumes and Volume Groups refer to Common Provisioning Groups.

Reclaimable Space by Host Group

A bar chart illustrates additional areas that should be considered for reclamation. Click a section of a bar to view details.



Reclaimable Space by Host Group Details

Click the + to expand the list of details.

Table 10-4 Reclaimable Space by Host Group Details

Column name	Description
Host Group	Host Group name links to the Host Utilization Detail report.
LUNs Allocated Not Used	Hosts can see these LUNs, but have done nothing with them; this links to the Allocated But Unused LUNs report. For HP 3PAR: LUNs refer to Virtual Volumes (VVs).
LUNs Added to Volume Group but Unused	LUNs were added to the Volume Group, but the LUNs are still partly unused. For HP 3PAR: LUNs refer to Virtual Volumes (VVs).
LUNs Allocated to Decommissioned Host	Links to the Suspected De-provisioned Hosts report. In the Report Designer, select a time period for “No host update for” to help identify decommissioned hosts. For HP 3PAR: LUNs refer to Virtual Volumes (VVs).
Unused Space in Volume Group	Additional space is available in the Volume Group to give to filesystems. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
Volumes Created Not Mounted	Links to the Unused Partitions/Logical Volumes report. See “ Unused Partitions/Logical Volumes ” on page 448.
Unused Space Inside Database	Links to the Over-Provisioned Applications report. See “ Over-Provisioned Applications” on page 156.
Total	The total amount of space that could be a candidate for reclamation.

Reclaimable Space by Host

Click the Reclaimable Space by Host link to list the details.

Note: Due to rounding, some values may appear as zeros when in fact, they are not. By changing the units--for example, from TB to GB--smaller values will be displayed.

Table 10-5 Reclaimable Space by Host

Column name	Description
Host	Host name links to the Host Utilization Detail report.
LUNs Allocated Not Used	Hosts can see these LUNs, but have done nothing with them; this links to the Allocated But Unused LUNs report. For HP 3PAR: LUNs refer to Virtual Volumes (VVs).
LUNs Added to Volume Group but Unused	LUNs were added to the Volume Group, but the LUNs are still partly unused. For HP 3PAR: LUNs refer to Virtual Volumes (VVs).
LUNs Allocated to Decommissioned Host	Links to the Suspected De-provisioned Hosts report. In the Report Designer, select a time period for “No host update for” to help identify decommissioned hosts. For HP 3PAR: LUNs refer to Virtual Volumes (VVs).
Unused Space in Volume Group	Additional space is available in the Volume Group to give to filesystems. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
Volumes Created Not Mounted	Links to the Unused Partitions/Logical Volumes report.
Unused Space Inside Database	Links to the Over-Provisioned Applications report.
Total	The total amount of space that could be a candidate for reclamation.

LUNS Added to Volume Group but Unused

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Available or Reclaimable Storage > Reclamation Summary

Click LUNs Added to Volume Group by Unused.

Access this report from the Reclamation Summary report.

Table 10-6 LUNS Added to Volume Group but Unused

Column name	Description
Host Name	Host associated with the Volume Group. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
Volume Group	Volume Group name. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
Capacity	Capacity of the Volume Group. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
Used	Amount of Volume Group space that is in use for this host. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
Available	Amount of Volume Group space for this host that is available. For HP 3PAR: Volume Group refers to Common Provisioning Group (CPG).
RAID Type	RAID type.

Suspected De-provisioned Hosts

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Suspected De-provisioned Hosts

Sometimes LUNs have been assigned to a host, but later, that host is re-assigned or de-provisioned. Often when hosts are retired, their storage is not given back to the pool and that storage remains unused.

Capacity Manager identifies suspected de-provisioned hosts by determining data staleness--that is, finding hosts for which data has not been collected within the specified time period. The default period is 14 days, however, you can choose from 3 days to a year for the report's time period. Hosts in this list deserve further investigation to determine if they truly have been de-provisioned.

Table 10-7 Suspected De-provisioned Hosts

Column name	Description
Name	The name of the logical unit (LUN).
RAID Type	The RAID type of the volume group.
Array Group	The array group from which this host originated. See " Array Group Details " on page 359.
Storage Array	The name of the array on which the host's application is consuming storage.
Capacity	Capacity of the LUN.
Estimated Used	Estimated capacity that is in use.
Estimated Available	Estimated capacity available for this host.
Mount Point Used	Amount of storage used by the mount point.
Host	Host name. See " Host Utilization Detail " on page 481.
Host LUN	The LUN on which this storage is located
Mount Point	The mount point for the filesystem.
Filesystem	The filesystem that is consuming storage.
Volume Group	The volume group associated with this host.
Host Storage Domain	The storage domain associated with this host.
Array Port WWN	The worldwide names that identify the array ports.

Unallocated LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Unallocated LUNs

These are simply LUNs that have not been allocated to any hosts. The array has created the LUNs, but there are no storage paths.

For HP 3PAR: This report refers to Virtual Volumes (VVs).

Additional details can be accessed via the following drill-down links in the report:

- Array Group
  See “ Array Group Details ” on page 359.
- Array Name
  See “ Array Capacity and Utilization ” on page 164.

For descriptions of the displayed fields,

See “ LUN Utilization Summary ” on page 344.

Unused Disks/Volume Groups

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Unused Disks/Volume Groups

Use this report to view a list of hosts that have available space. These hosts have space from which more partitions or logical volumes can be created. This report can be used to identify storage that has been requested, but that is not being used. Then, the list can be emailed to the group that has requested the storage.

In the report’s scope selector, the Host Type option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other. Note

that host disks, not volume groups, are displayed for virtual hosts. Also, there are no unused disks for VM Guests, as virtual disks are created by the VM Guest.

Note: In this report, if the RAID type = 1, the capacity is divided by 2, to reflect true usable capacity.

For HP 3PAR, volume groups refer to Common Provisioning Groups.

Additional details can be accessed via the following links in the report:

- Host Name
  See " Host Utilization Detail " on page 481.

Table 10-8 Unused Disks/Volume Groups

Column name	Description
Host Name	Host identified as having available space.
Disk Vol Grp	The volume group that has available space. For HP 3PAR: Volume group refers to Common Provisioning Group (CPG).
Type	Type of Disk/Volume Group. For HP 3PAR: Volume group refers to Common Provisioning Group (CPG).
Capacity	Capacity of the host.
Used	Used capacity.
Available	Capacity available for this host.
RAID Type	The RAID type of the volume group. For HP 3PAR: Volume group refers to Common Provisioning Group (CPG).

Unused Partitions/Logical Volumes

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Available or Reclaimable Storage > Unused Partitions/Logical Volumes

This report lists all hosts with logical volumes that have no mounted filesystems. A logical volume may have been created in the volume group, but the logical volume is not being used by the host. These unused volumes provide available storage.

In the report's scope selector, the Host Type option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other. Note that host disks, not volume groups, are displayed for virtual hosts. Also, there are no unused disks for VM Guests, as virtual disks are created by the VM Guest.

Table 10-9 Unused Partitions/Logical Volumes

Column name	Description
Host Name	The name of the host.
Volume Group/ Disk Name	The volume group and/or disk name.
Volume Group/ Disk Size	The volume group and/or disk size.
Volume Group/ Disk Avail	The amount of volume group storage that is available for use.
Logical Volume	The name of the logical volume (partition).
LV Size	The logical volume size.
RAID Type	The RAID type of the volume.

Additional details can be accessed via the following drill-down links in the report:

- Host
  See " Host Utilization Detail " on page 481.

Chapter 11

Capacity at Risk Reports

This chapter includes the following topics:

- Exception Summary
- Hosts at Risk
- LUNs at Risk
- NetApp Aggregates at Risk
- NetApp Cluster-Mode Aggregates at Risk
- NetApp Volumes at Risk
- NetApp Cluster-Mode Volumes at Risk
- Over-Provisioned Hosts
- Over-Provisioned LUNs

Exception Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity at Risk > Exception Summary

This report provides a list of filesystem mount points that may require attention for one of the following reasons:

- Near Capacity: > 90%

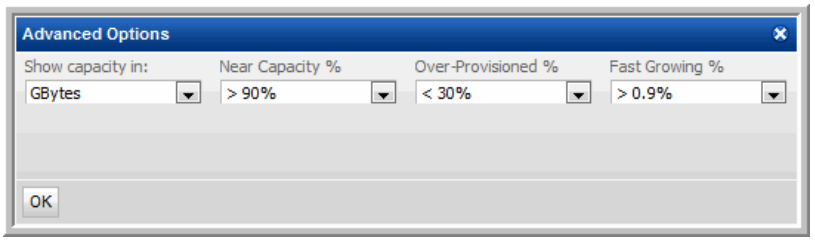
- Over-Provisioned: < 30%
- Empty: < 5% (negative Rate)
- Fast-Growing: > .9%

Note: To provide a realistic forecast, the time span for determining when capacity may be at risk has been limited to five years.

See "Exception Summary Report Details" on page 451.

Advanced Options

In the scope selector, click **Advanced** to override the default parameters shown above:



These Advanced Options are used to determine capacity exceptions, such as near capacity and over-provisioned.

Exception Summary Report Details

Table 11-1 Exception Summary Report Details

Column name	Description
Host	Host name links to the Host Utilization Detail report.
Mount Point	The mount point.
Current Capacity	Total capacity of the mount point.
Current Used	Current allocated storage in mount point.
Current Available	Current Capacity - Current Used.
Current % Used	Current Used / Current Capacity.

Table 11-1 Exception Summary Report Details *(continued)*

Column name	Description
Avg Daily Growth	The value is calculated using the following: capacity sum (used – previously used)/total number of days for the time interval. If the average daily growth is less than 0, the columns, Days Left and Date of Concern are not populated.
Std Dev	Standard Deviation for the Avg Daily Growth.
Max Daily Growth	Maximum daily growth from the selected period of time. Provides an indication of the average validity. Growth must be at least 1 MB to be considered.
75% - Days Left	The value is calculated using the following: (current capacity/avg daily growth) x .75.
75% - Date of Concern	Date of the report + Days Left.
90% - Days Left	((.9 * Current Capacity/Avg Daily Growth).
90% - Date of Concern	Date of the report + Days Left.
100% - Days Left	Current Capacity/Avg Daily Growth.
100% - Date of Concern	Date of the report + Days Left.
Status	Near Capacity if Current % Used > specified criteria; Over-Provisioned if Current % Used < specified criteria; Out of Space if Current Available < Avg Daily Growth; Empty if Current % Used < 5%.
Rate	Growth rate. Fast-growing if Avg Daily Growth/Current Capacity > specified criteria; Negative if Avg Daily Growth < .05.

Hosts at Risk

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > Hosts At Risk

When you generate this report, select a capacity threshold: Low, Warning, or Critical. Typically, for at risk reports, you will want to view the Critical hosts--that is, those that are at risk of breaching the critical threshold.

In the report's scope selector, the **Host Type** option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other.

This report lists only the hosts that require your immediate attention. Use the host link to view details.

Note: When reviewing the values in this report, note that hosts at risk are determined by evaluating current and projected usage.

Table 11-2 Hosts at Risk

Column name	Description
Host Name	Click a host name to access details in the Host Utilization Detail report. See " Host Utilization Detail " on page 481.
Risk Date (Projected)	The date that the host is projected to exceed its threshold and will no longer be considered over-provisioned.
Usage (Projected)	Projected usage associated with the risk date.
High-Water Mark (Projected)	The high-water mark represents the maximum usage for a period. Note that this size may represent temporary usage.
FileSystem Capacity	The filesystem capacity
Capacity	The host's storage limit.
Current Usage	Storage already consumed.
Threshold	A host becomes at risk when this tolerance threshold is crossed. The default thresholds are set at: Low = 30%; Warning = 70%; Critical = 90%. To modify thresholds for a particular host, go to the Portal's Admin toolbar: **Admin > Threshold Policies**.

LUNs at Risk

Pro-actively increase LUN capacity and avoid outages or performance issues associated with insufficient capacity.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > LUNs At Risk

When you generate this report, select a capacity threshold: Low, Warning, or Critical. Typically, for at risk reports, you will want to view the Critical LUNs--that is, those that are at risk of breaching the critical threshold and running out of space.

Note: When reviewing the values in this report, note that LUNs at risk are determined by evaluating current and projected usage.

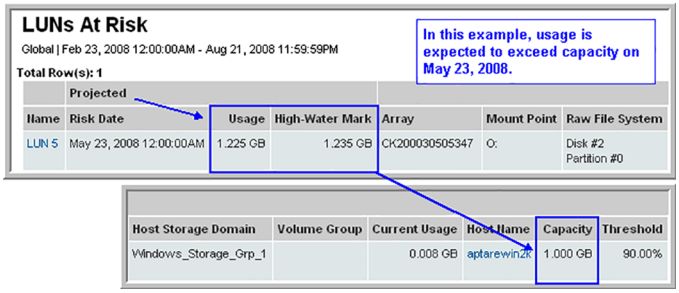


Table 11-3 LUNs at Risk

Column name	Description
Name	The specific logical unit number that is at risk of running out of room. Click on this link to access details in the LUN Details report.
Risk Date (Projected)	The date that the LUN is projected to exceed its capacity.
Usage (Projected)	Usage based on the period that you've selected.
High-Water Mark (Projected)	The projected high-water mark represents the maximum usage, based on historical data. Note that this size may represent temporary usage.
Array	The group to which this LUN belongs.
Mount Point	The host's filesystem mount point for this LUN.
Raw Filesystem	The explicit full directory path for the filesystem.

Table 11-3 LUNs at Risk *(continued)*

Column name	Description
Host Storage Domain	The name of the host storage domain, a logical group of LUNs that can be accessed by multiple hosts through the same physical storage port, with each host accessing its unique LUN. For security purposes, LUN masking is also coupled with this capability. Host storage domain is Hitachi's name for its multi-hosted storage capability.
iGroup	The iGroup associated with the LUN.
Volume Group	The logical collection or volume group to which the volume belongs.
Current Usage	The amount of storage already committed.
Host Name	Click on a host name to launch the Host Utilization Detail report. See " Host Utilization Detail " on page 481.
Capacity	The total capacity (used and unused) of the LUN.
Threshold	A LUN is determined to be at risk when this tolerance threshold is crossed. The default thresholds are set at: Low = 30%; Warning = 70%; Critical = 90%. To modify thresholds for a particular LUN, go to ADMIN > Reports > Threshold Policies.

NetApp Aggregates at Risk

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > NetApp Aggregates at Risk

For thin provisioning, it's important to know when host usage is about to exceed the storage that was promised. The details shown in this aggregate report can help you understand the current physical capacity status. When an aggregate is at risk--the amount of used storage in the FlexVol approaches the aggregate's capacity--it can be re-sized to add physical disk space to the aggregate.

See " NetApp Storage System Detail " on page 221.

Note: The Total Capacity value represents the amount of storage that has been committed or promised to hosts. Therefore, host resources data collection must be implemented to determine which Aggregates are at risk and an associated host must be at risk in order for the Aggregate to be considered at risk.

When you generate this report, you can select the threshold to report: Low, Warning, or Critical.

Table 11-4 NetApp Aggregates at Risk

Column name	Description
Aggregate	The Aggregate number/ID links to details. See " NetApp Aggregate Detail " on page 368.
Storage Array	The NetApp system with which the aggregate is associated, which drills down to details. See " NetApp Storage System Detail " on page 221.
Usage	Mouse over the usage thermometer to view the % available
Total Capacity	The aggregate's total capacity
Allocated Capacity	Amount of capacity that has been allocated for use
Available Capacity	Remaining capacity available for use
RAID Type	raid0 (mirrored), raid4 (single disk for parity), or raid_dp (double-parity protection)
# Volumes at Risk	Links to the NetApp Volumes at Risk. See " NetApp Volumes at Risk " on page 458.
# of Hosts at Risk	Number of hosts at risk.
List of Plexes	Lists 1 or 2 Plexes--2 if SyncMirror is licensed and enabled. Links to the details. See " NetApp Plex Details " on page 373.
Risk Date	The point at which the aggregate is at risk

NetApp Cluster-Mode Aggregates at Risk

For thin provisioning, it's important to know when host usage is about to exceed the storage that was promised. The details shown in this aggregate report can help you understand the current physical capacity status. When an aggregate is at

risk--the amount of used storage in the FlexVol approaches the aggregate's capacity--it can be re-sized to add physical disk space to the aggregate. This is for Cluster Mode arrays.

See " NetApp Storage System Detail " on page 221.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > NetApp Cluster-Mode Aggregates at Risk

---

Note: The Total Capacity value represents the amount of storage that has been committed or promised to hosts. Therefore, host resources data collection must be implemented to determine which Aggregates are at risk and an associated host must be at risk in order for the Aggregate to be considered at risk.

When you generate this report, you can select the threshold to report: Low, Warning, or Critical.

Table 11-5 NetApp Cluster-Mode Aggregates at Risk

Column name	Description
Aggregate	The Aggregate number/ID links to NetApp Cluster-Mode Volumes at Risk. See " NetApp Cluster-Mode Volumes at Risk " on page 459.
Storage Array	Drills down to the Array Capacity and Utilization (NetApp Cluster) report. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Usage	Mouse over the usage thermometer to view the % available
Total Capacity	The aggregate's total capacity
Allocated Capacity	Amount of capacity that has been allocated for use
Available Capacity	Remaining capacity available for use
RAID Type	raid0 (mirrored), raid4 (single disk for parity), or raid_dp (double-parity protection)

Table 11-5 NetApp Cluster-Mode Aggregates at Risk *(continued)*

Column name	Description
# Volumes at Risk	Links to the NetApp Cluster-Mode Volumes at Risk report. See " NetApp Cluster-Mode Volumes at Risk " on page 459.
# of Hosts at Risk	Number of hosts at risk.
List of Plexes	Lists 1 or 2 Plexes--2 if SyncMirror is licensed and enabled. See " NetApp Plex Details " on page 373.
Risk Date	The point at which the aggregate is at risk

NetApp Volumes at Risk

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > NetApp Volumes at Risk

When you configure thin provisioning, you enable multiple volumes to access the same pool of free storage. Since thin provisioning "commits" more storage to hosts than is actually available, it is important to identify the remaining, usable storage and the point at which the aggregate is at risk of not having available storage.

Note: Host resources data collection must be implemented to determine which NetApp Volumes are at risk.

See " NetApp Storage System Detail " on page 221.

When you generate this report, you can select the threshold to report: Low, Warning, or Critical.

Table 11-6 NetApp Volumes at Risk

Column name	Description
Volume	Volume name links to details. See " NetApp Volume Details " on page 398.

Table 11-6 NetApp Volumes at Risk *(continued)*

Column name	Description
Storage Array	The name of the NetApp storage system, which links to details. See " NetApp Storage System Detail " on page 221.
Aggregate	The name of the associated aggregate, with a link to the details. See " NetApp Aggregate Detail " on page 368.
Aggregate Risk Date	The date at which the aggregate is expected to not have storage available for use
Space Guarantee	The type of space that has been committed: volume, file, or none. Note that file or none indicates thin provisioning.
Usage	Mouse over the thermometer to see the % available.
Total Capacity	Total capacity of the volume
Allocated Capacity	Amount of the volume's capacity that has been allocated
Available Capacity	Amount of the volume's capacity that is available
Host Name	The host that is at risk of consuming more than the space guaranteed to it.
Host Usage	Mouse over the thermometer bar to view the estimated % used by the host.
Mount Points	The host's mount points that are using the storage
Host Capacity	The capacity of the host
Estimate Used	Estimated amount of storage used by the host
Estimate Available	Estimated storage available for host consumption

NetApp Cluster-Mode Volumes at Risk

When you configure thin provisioning, you enable multiple volumes to access the same pool of free storage. Since thin provisioning commits more storage to hosts than is actually available, it is important to identify the remaining, usable storage and the point at which the aggregate is at risk of not having available storage. Host resources data collection must be implemented to determine which NetApp Volumes are at risk. Host resources data collection must be implemented to determine which NetApp Volumes are at risk. This is for Cluster Mode arrays.

Note: Host resources data collection must be implemented to determine which NetApp Volumes are at risk.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > NetApp Cluster-Mode Volumes at Risk

When you generate this report, you can select the threshold to report: Low, Warning, or Critical.

Table 11-7 NetApp Cluster-Mode Volumes at Risk

Column name	Description
Volume	Volume name. See " NetApp Cluster-Mode Volume Detail " on page 401.
Storage Array	The name of the NetApp storage system. See "Array Capacity and Utilization (NetApp Cluster)" on page 214.
Aggregate	The name of the associated aggregate. See " NetApp Cluster-Mode Volumes at Risk " on page 459.
Aggregate Risk Date	The date at which the aggregate is expected to not have storage available for use.
Space Guarantee	The type of space that has been committed: volume, file, or none. Note that file or none indicates thin provisioning.
Usage	Mouse over the thermometer to see the % available.
Total Capacity	Total capacity of the volume.
Allocated Capacity	Amount of the volume's capacity that has been allocated.
Available Capacity	Amount of the volume's capacity that is available.
Host Name	The host that is at risk of consuming more than the space guaranteed to it.
Host Usage	Mouse over the thermometer bar to view the estimated % used by the host.

Table 11-7 NetApp Cluster-Mode Volumes at Risk *(continued)*

Column name	Description
Mount Points	The host's mount points that are using the storage.
Host Capacity	The capacity of the host.
Estimate Used	Estimated amount of storage used by the host.
Estimate Available	Estimated storage available for host consumption.

Over-Provisioned Hosts

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > Over Provisioned Hosts

By determining which hosts are over-provisioned, you can take steps to reclaim the space before committing resources to acquire additional storage. Over-provisioning leads to increased costs as a result of under-utilized storage.

In the report's scope selector, the **Host Type** option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other.

When you execute this report, you can select the Threshold for the list that will be displayed. Typically, you'll be interested in the Low or Warning thresholds. Since this report shows over-provisioning, the list will show usage that has not reached the threshold.

Table 11-8 Over-Provisioned Hosts

Column name	Description
Host Name	Click on a host name to access details. See " Host Utilization Detail " on page 481.
As of (Projected)	This date indicates when the host became over-provisioned.
Usage (Projected)	Projected consumption by the host.

Table 11-8 Over-Provisioned Hosts *(continued)*

Column name	Description
High-Water Mark (Projected)	The projected highest level of usage, based on historical data; the high-water mark represents the maximum usage for a period.
Disk Capacity (Current)	The host's current storage limit.
FileSystem Capacity (Current)	The current (not projected) capacity of the filesystem.
Usage (Current)	Storage already consumed.
High Water Mark (Current)	The high-water mark represents the maximum usage for a period. Note that this size may represent temporary usage.
Threshold (Current)	A host becomes over-provisioned when the usage is under the threshold. The default thresholds are set at: Low = 30%; Warning = 70%; Critical = 90%. To modify thresholds for a particular host, go to the Portal's Admin toolbar: **Admin > Threshold Policies**.

Over-Provisioned LUNs

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Capacity At Risk > Over Provisioned LUNs

When you execute this report, you can select the Threshold for the list that will be displayed. Typically, you'll be interested in the Low or Warning thresholds. Since this report shows over-provisioning, the list will show usage that has not reached the threshold.

Table 11-9 Over-Provisioned LUNs

Column name	Description
Name	The LUN number/ID drills down to the LUN Details report.
As of	This date indicates when the LUN became over-provisioned.

Table 11-9 Over-Provisioned LUNs *(continued)*

Column name	Description
High-Water Mark (Projected)	The projected high-water mark represents the maximum usage, based on historical data. Note that this size may represent temporary usage.
Host Name	Click on a host name to launch the Host Utilization Detail report. See " Host Utilization Detail " on page 481.
Capacity	The total capacity (used and unused) of the LUN.
Current Usage	Amount of storage currently used.
Available	Capacity - Current Usage.
Array	The array to which the LUN belongs.
Mount Point	Location of the filesystem.
Raw Filesystem	Explicit disk location of the LUN.
Host Storage Domain	The name of the host storage domain, a logical group of LUNs that can be accessed by multiple hosts through the same physical storage port, with each host accessing its unique LUN. For security purposes, LUN masking is also coupled with this capability. A host storage domain is Hitachi's name for multi-hosted storage capability.
iGroup	The iGroup associated with the LUN
Volume Group	The logical collection or volume group to which the volume belongs.
Threshold	A LUN is determined to be at risk when this tolerance threshold is crossed. The default thresholds are set at: Low = 30%; Warning = 70%; Critical = 90%. To modify thresholds for a particular LUN, go to **Admin > Reports > Threshold Policies**.

Chapter 12

Capacity Chargeback Reports

This chapter includes the following topics:

- Aggregated Chargeback
- Aggregated Chargeback by Host Groups
- Aggregated Chargeback by Attribute
- Chargeback Detail
- Chargeback Array Capacity
- Chargeback By Host
- Chargeback by LUN Detail
- Chargeback Policy Capacity
- Policy Detail
- Object Store Chargeback
- Object Store Chargeback Detail

Aggregated Chargeback

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Aggregated Chargeback

---

Note: This is a host-based report. Host data collection is required before data will appear in this report.

This Capacity Chargeback report enables you to aggregate the data by Host Groups or by a particular host attribute (for example, Department).

The Report Designer provides both a filtering and a grouping mechanism.

- Host Groups are used to filter or narrow the scope and content of the report.
- Host Groups and Attributes are used to aggregate the data--that is, a grouping mechanism. Note that Attributes are from the host's perspective.
- If you choose only Host Groups (without an Attribute), the report will be filtered on the selected Host Groups and also will be grouped by Host Group.
- Allocate Duplicate Host to All, when checked, reports duplicate host usage in all Host Groups, if hosts are in more than one host group.When unchecked, reports duplicate host usage only in the first alphabetical host group.

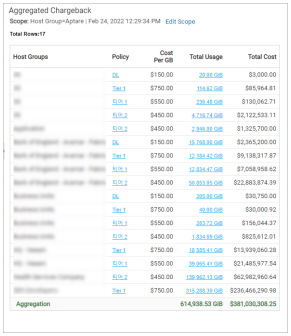
Capacity Chargeback Policies are configured by the NetBackup IT Analytics Administrator. The administrator uses these policies to assign a cost-per-GB to hosts in a host group. See examples in the following.

See "Aggregated Chargeback by Host Groups" on page 465.

See "Aggregated Chargeback by Attribute" on page 466.

Aggregated Chargeback by Host Groups

When you do not select an Attribute in the Report Designer, the report is organized by Host Groups.


Aggregated Chargeback

Scope: Host Group=Aptare | Feb 24, 2022 12:29:34 PM Edit Scope

Total Rows:17

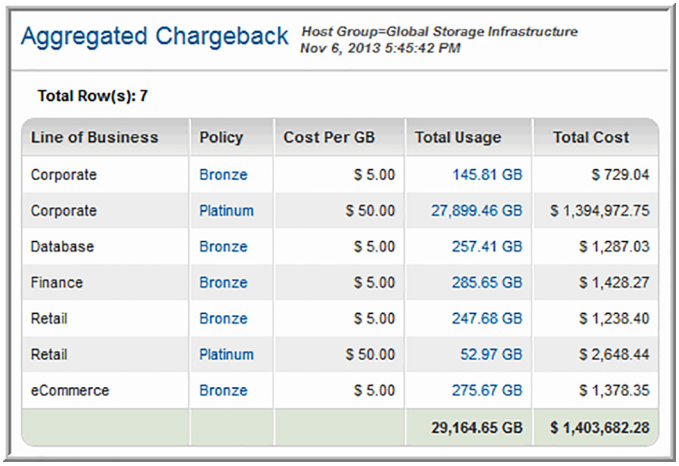
Host Groups	Policy	Cost Per GB	Total Usage	Total Cost
[redacted]	DL	$150.00	20.98 GiB	$3,000.00
[redacted]	Tier 1	$750.00	114.62 GiB	$85,964.81
[redacted]	FLR-1	$550.00	236.48 GiB	$130,062.71
[redacted]	FLR-2	$450.00	4,716.74 GiB	$2,122,533.11
[redacted]	FLR-2	$450.00	2,946.86 GiB	$1,325,700.00
[redacted]	DL	$150.00	15,769.86 GiB	$2,365,200.00
[redacted]	Tier 1	$750.00	12,184.42 GiB	$9,138,317.87
[redacted]	FLR-1	$550.00	12,834.47 GiB	$7,058,958.62
[redacted]	FLR-2	$450.00	50,853.95 GiB	$22,883,874.39
[redacted]	DL	$150.00	205.66 GiB	$30,750.00
[redacted]	Tier 1	$750.00	40.88 GiB	$30,000.92
[redacted]	FLR-1	$550.00	283.72 GiB	$156,044.37
[redacted]	FLR-2	$450.00	1,834.68 GiB	$825,612.01
[redacted]	Tier 1	$750.00	18,585.41 GiB	$13,939,060.28
[redacted]	FLR-1	$550.00	39,995.41 GiB	$21,485,977.54
[redacted]	FLR-2	$450.00	139,962.13 GiB	$62,982,960.64
[redacted]	Tier 1	$750.00	315,288.39 GiB	$236,466,290.98
Aggregation			614,938.53 GiB	$381,030,308.25

Table 12-1 Aggregated Chargeback by Host Groups

Column name	Description
Host Groups OR <Attribute Name>	Host group to which the policy applies; if an Attribute was selected in the Report Designer the values are displayed in Attribute alphabetical order.
Policy	Links to Policy Detail.
Cost Per GB	The unit cost configured in the policy.
Total Usage	Links to Chargeback Detail.
Total Cost	Cost-per-GB * Total Usage.

Aggregated Chargeback by Attribute

When you select an Attribute in the Report Designer, the report is organized by that Attribute.



Aggregated Chargeback *Host Group=Global Storage Infrastructure*
Nov 6, 2013 5:45:42 PM

Total Row(s): 7

Line of Business	Policy	Cost Per GB	Total Usage	Total Cost
Corporate	Bronze	$ 5.00	145.81 GB	$ 729.04
Corporate	Platinum	$ 50.00	27,899.46 GB	$ 1,394,972.75
Database	Bronze	$ 5.00	257.41 GB	$ 1,287.03
Finance	Bronze	$ 5.00	285.65 GB	$ 1,428.27
Retail	Bronze	$ 5.00	247.68 GB	$ 1,238.40
Retail	Platinum	$ 50.00	52.97 GB	$ 2,648.44
eCommerce	Bronze	$ 5.00	275.67 GB	$ 1,378.35
			29,164.65 GB	**$ 1,403,682.28**

Chargeback Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Chargeback and Billing > Aggregated Chargeback

Click a Total Usage link.

Table 12-2 Chargeback Detail

Column name	Description
Host	Links to the Host Utilization Detail.
Policy	Links to the Chargeback Policy.
Array	Links to the Array Capacity & Utilization.
LUN	Links to the LUN Utilization Summary.

Table 12-2 Chargeback Detail *(continued)*

Column name	Description
Unit Cost	The per-unit amount configured in the Capacity Chargeback Policy.
Total Usage	Total amount of storage in use.
Total Cost	Cost-per-GB * Total Usage.

Chargeback Array Capacity

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Chargeback Array Capacity

Refer to the *NetBackup IT Analytics User Guide* for details on configuring a Capacity Chargeback Policy. Capacity Chargeback Policies are configured by the NetBackup IT Analytics Administrator. The administrator uses these policies to assign a cost-per-GB to hosts in a host group.

Table 12-3 Chargeback Array Capacity

Column name	Description
Array	
Name	Click the array name link to the Array Capacity and Utilization report.
Policy	Links to the capacity chargeback policy details.
Cost Per GB	The cost per unit, as defined by the capacity chargeback policy.
RAID	
Total Capacity	Total capacity of the array groups in the storage array.
Virtual Capacity	Applicable to HDS arrays only - Capacity from a back-end system, derived from daisy-chained external storage (back-end systems). This is array group capacity that does not have PDEVs in the storage array.
Local Capacity	All the capacity from the attached disks.

Table 12-3 Chargeback Array Capacity *(continued)*

Column name	Description
Used	Total used capacity of the array groups (by LUNs) in the storage array.
Available	Available capacity of the array groups in the storage array. Available Capacity = Total Capacity - Allocated Capacity.
LUNs	
Total	Total LUN capacity for the array tier. The sum of these values of each array should match the Total Capacity listed in the Array Capacity & Utilization report. For HP 3PAR: LUN refers to Virtual Volume.
Allocated	Allocated LUN capacity for the array tier. The sum of these values of each array should match the Allocated Capacity listed in the Array Capacity & Utilization report. For HP 3PAR: LUN refers to Virtual Volume.
UnAllocated	Unallocated LUN capacity for the array tier. The sum of these values of each array should match the Unallocated Capacity listed in the Array Capacity & Utilization report. For HP 3PAR: LUN refers to Virtual Volume.
Thin Provisioned	
Capacity	The Hitachi Dynamic Provisioning capacity associated with this array. Or, consumed thin provisioned capacity for: EMC Symmetrix, EMC CLARiiON, or IBM SVC.
Used	The consumed Hitachi Dynamic Provisioning, EMC Symmetrix, EMC CLARiiON, or IBM SVC capacity for this array.
Available	The total unused thin provisioned capacity. For Hitachi arrays, this is the total Journal Pool Available, NOT just for this array.

Chargeback By Host

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

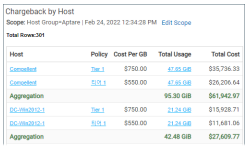
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Chargeback By Host

Note: Column filtering is not supported for this tabular report.

This is a host-based report. Host data collection is required before data will appear in this report.

Capacity Chargeback Policies are configured by the NetBackup IT Analytics Administrator. The administrator uses these policies to assign a cost-per-GB to hosts in a host group.


Chargeback by Host

Scope: Host Group=Aptare | Feb 24, 2022 12:34:28 PM Edit Scope

Total Rows:301

Host	Policy	Cost Per GB	Total Usage	Total Cost
Comptlnd	Tier 1	$750.00	47.95 GiB	$35,736.33
Comptlnd	R/R 1	$550.00	47.95 GiB	$26,206.64
Aggregation			95.90 GiB	$61,942.97
DC-Win2012-1	Tier 1	$750.00	21.24 GiB	$15,928.71
DC-Win2012-1	R/R 1	$550.00	21.24 GiB	$11,681.06
Aggregation			42.48 GiB	$27,609.77

Table 12-4 Chargeback By Host

Column name	Description
Host	Host name links to the Host Utilization Detail.
Policy	Links the capacity chargeback policy details.
Cost Per GB	The cost per unit, as defined by the capacity chargeback policy.
Total Usage	Total capacity that will be charged for this host; links to Chargeback Detail.
Total Cost	Usage * Unit Cost = Total Cost.

Chargeback by LUN Detail

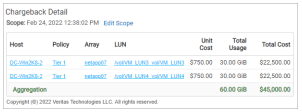
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Chargeback and Billing > Chargeback By Host

Click Total Usage.


Chargeback Detail
Scope: Feb 24, 2022 12:38:02 PM Edit Scope

Host	Policy	Array	LUN	Unit Cost	Total Usage	Total Cost
DC-Win2K8-2	Tier 1	netapp07	/vol/VM_LUN3_vol/VM_LUN3	$750.00	30.00 GiB	$22,500.00
DC-Win2K8-2	Tier 1	netapp07	/vol/VM_LUN4_vol/VM_LUN4	$750.00	30.00 GiB	$22,500.00
Aggregation					60.00 GiB	$45,000.00

Copyright (R) 2022 Veritas Technologies LLC. All rights reserved.

Table 12-5 Chargeback Detail

Column name	Description
Host	Host name links to the Host Utilization Detail.
Policy	Links the capacity chargeback policy details.
Array	Links to the Array Capacity & Utilization.
LUN	Links to the LUN Utilization Summary.
Unit Cost	The cost per unit, as defined by the capacity chargeback policy.
Total Usage	Amount of storage used for this host.
Total Cost	Usage * Unit Cost = Total Cost.

Chargeback Policy Capacity

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Chargeback Policy Capacity

The Chargeback Policy Capacity reports on all LUNs associated with the a Capacity Chargeback Policy.

Refer to the *User Guide* for details on configuring a Capacity Chargeback Policy. Capacity Chargeback Policies are configured by the NetBackup IT Analytics Administrator. The administrator uses these policies to assign a cost-per-GB to hosts in a host group.

Table 12-6 Chargeback Policy Capacity

Column name	Description
Policy	The name of the capacity chargeback policy links to the policy details.
Cost Per GB	The cost per unit that will be charged.
RAID	
Total Capacity	Total capacity of the array groups in the storage arrays for this policy.
Virtual Capacity	Applicable to HDS arrays only - Capacity from a back-end system, derived from daisy-chained external storage (back-end systems). This is array group capacity that does not have PDEVs in the storage array.
Local Capacity	All the capacity from the attached disks.
Used	Total used capacity of the array groups (by LUNs) in the storage array.
Available	Available capacity of the array groups in the storage array. Available Capacity = Total Capacity - Allocated Capacity.
LUNs	
Total	Total LUN capacity for the policy.
Allocated	Allocated LUN capacity for the policy.
UnAllocated	Unallocated LUN capacity for the policy.
Thin Provisioned	
Total	The Hitachi Dynamic Provisioning capacity associated with this policy. Or, thin provisioned capacity for: EMC Symmetrix, EMC CLARiiON, or IBM SVC.
Used	The consumed Hitachi Dynamic Provisioning capacity associated with this policy. Or, consumed thin provisioned capacity for: EMC Symmetrix, EMC CLARiiON, or IBM SVC.
Pool Available	The total unused Hitachi Dynamic Provisioning (HDP), EMC Symmetrix, EMC CLARiiON, or IBM SVC capacity. For HDP, This is the total Journal Pool Available, NOT just for this policy.

Policy Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports

cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Chargeback and Billing > Aggregated Chargeback

Then click a Policy link.

Table 12-7 Policy Detail

Column name	Description
Item Type	Displays the Policy Type, such as Array Family, Array Type, Array Name, Drive Speed, Drive Type, Drive Capacity, RAID Type, Is Local, Is HDP, Is DAS, or Domain.
Item Value	The value associated with the Policy/Item type.
Display Value	The value to be displayed in reports.

Object Store Chargeback

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Object Store Chargeback

View hourly dynamic statistics for a tenant's namespace or a user's volume or file system. Use this report to identify ingested volumes, used capacity, and number of reads/writes. For example, this report is relevant for Hitachi Content Platform (HCP) chargeback.

Table 12-8 Object Store Chargeback

Column name	Description
Cluster/Array	Name of the storage array.

Table 12-8 Object Store Chargeback *(continued)*

Column name	Description
Model	Storage array family.
Tenant/User	Name of the tenant.
# of Namespaces	Number of namespaces owned by the tenant. Click the link to view namespace-level statistics.
Start Date	The start time of the reporting interval for the set of statistics.
End Date	The end time of the reporting interval for the set of statistics.
# of Objects	The total number of objects and versions stored in the namespace.
Ingested Volume	Amount of stored data and custom metadata before it was added to the namespace. The total size of the stored data and custom metadata, in bytes, before it was added to the identified namespace, to any of the namespaces owned by the identified tenant, or to any of the namespaces in the system.
Used Capacity	The amount of storage currently occupied by all data stored in the namespace, including object data, metadata, and any redundant data. Values are stored as KiB in the database and rendered according to your user profile preferences.
Bytes In	Number of KB read from the namespace. The total number of KB read from the identified namespace, from any of the namespaces owned by the identified tenant, or from any of the namespaces in the system during the reporting interval. If data (including XML for directory listings) was compressed before being transmitted, this is the number of bytes before compression.
Bytes Out	Number of KB written to the namespace. The total number of KB successfully written to the identified namespace, to any of the namespaces owned by the identified tenant, or to any of the namespaces in the system during the reporting interval. If data was compressed before being transmitted, this is the number of bytes before compression.
# of Reads	Number of namespace reads. The total number of read operations performed in the identified namespace, in all the namespaces owned by the identified tenant, or in all the namespaces in the system during the reporting interval.

Table 12-8 Object Store Chargeback *(continued)*

Column name	Description
# of Writes	Number of namespace writes. The total number of write operations successfully performed in the identified namespace, in all the namespaces owned by the identified tenant, or in all the namespaces in the system during the reporting interval.
Deleted?	Namespace or tenant delete status.
Valid?	The status of the set of statistics in the line. Possible values: true = All statistics in the set have been successfully collected by the system.. false = The statistics in the set do not reflect all the activity that occurred during the reporting interval. This may be due, for example, to one or more nodes being unavailable during that time, to a network failure, or to other hardware issues.
Last Updated	The date and time the node data was last updated in the NetBackup IT Analytics database.

Object Store Chargeback Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Chargeback and Billing > Object Store Chargeback

Click the # of Namespaces link.

This report provides namespace-level statistics and is accessed via the Object Store Chargeback Summary report. Use this report to identify ingested volumes, used capacity, and number of reads/writes. For example, this report is relevant for Hitachi Content Platform (HCP) chargeback.

Table 12-9 Object Store Chargeback Detail

Column name	Description
Cluster/Array	Name of the storage array.
Model	Storage array family.
Tenant/User	Name of the tenant.
Namespace	Name of the namespace.
Start Date	The start time of the reporting interval for the set of statistics.
End Date	The end time of the reporting interval for the set of statistics.
# of Objects	The total number of objects and versions stored in the namespace.
Ingested Volume	Amount of stored data and custom metadata before it was added to the namespace. The total size of the stored data and custom metadata, in bytes, before it was added to the identified namespace, to any of the namespaces owned by the identified tenant, or to any of the namespaces in the system.
Used Capacity	The amount of storage currently occupied by all data stored in the namespace, including object data, metadata, and any redundant data. Values are stored as KiB in the database and rendered according to your user profile preferences.
Bytes In	Number of KB read from the namespace. The total number of KB read from the identified namespace, from any of the namespaces owned by the identified tenant, or from any of the namespaces in the system during the reporting interval. If data (including XML for directory listings) was compressed before being transmitted, this is the number of bytes before compression.

Table 12-9 Object Store Chargeback Detail *(continued)*

Column name	Description
Bytes Out	Number of KB written to the namespace. The total number of KB successfully written to the identified namespace, to any of the namespaces owned by the identified tenant, or to any of the namespaces in the system during the reporting interval. If data was compressed before being transmitted, this is the number of bytes before compression.
# of Reads	Number of namespace reads. The total number of read operations performed in the identified namespace, in all the namespaces owned by the identified tenant, or in all the namespaces in the system during the reporting interval.
# of Writes	Number of namespace writes. The total number of write operations successfully performed in the identified namespace, in all the namespaces owned by the identified tenant, or in all the namespaces in the system during the reporting interval.
# of Objects Deleted	The number of objects and versions that have been deleted and are waiting to be shredded.
Deleted?	Namespace or tenant delete status.
Valid?	The status of the set of statistics in the line. Possible values: true = All statistics in the set have been successfully collected by the system.. false = The statistics in the set do not reflect all the activity that occurred during the reporting interval. This may be due, for example, to one or more nodes being unavailable during that time, to a network failure, or to other hardware issues.
Last Updated	The date and time the node data was last updated in the NetBackup IT Analytics database.

Chapter 13

Host Capacity Utilization Reports

This chapter includes the following topics:

- Host Capacity & Utilization
- Host Utilization Detail
- Host Utilization Detail (ASM Example)
- Veritas Volume Details
- Host Executive Summary
- Host HBA Summary

Host Capacity & Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

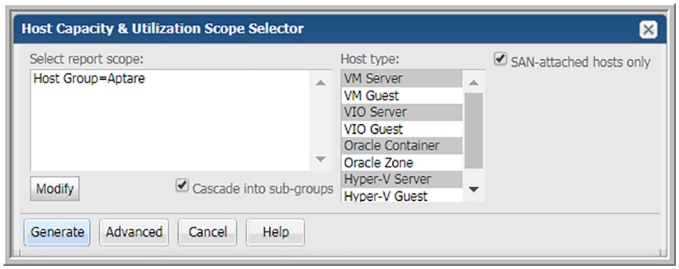
Capacity Manager > Host Capacity & Utilization > Host Capacity & Utilization

Use the Host Capacity and Utilization report to view an overview of host utilization to determine if there's contention for storage. Typically, once you've reviewed this high-level report, you'll want to access the details reported by the Host Utilization Detail.

See " Host Utilization Detail " on page 481.

In the report's scope selector, several widgets enable special filtering of the report's output:

- **Host Type** - Use this selector to include virtual servers in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other.To view a similar report for VM Guests, see the VM Summary report.
- **SAN-attached hosts only**- **SAN-attached hosts only**- Check this box if you want to see data for only SAN-attached hosts. When this box is not checked, the report will include VM Guests and hosts that have been collected by other NetBackup IT Analytics products, such as Backup Manager or Virtualization Manager.



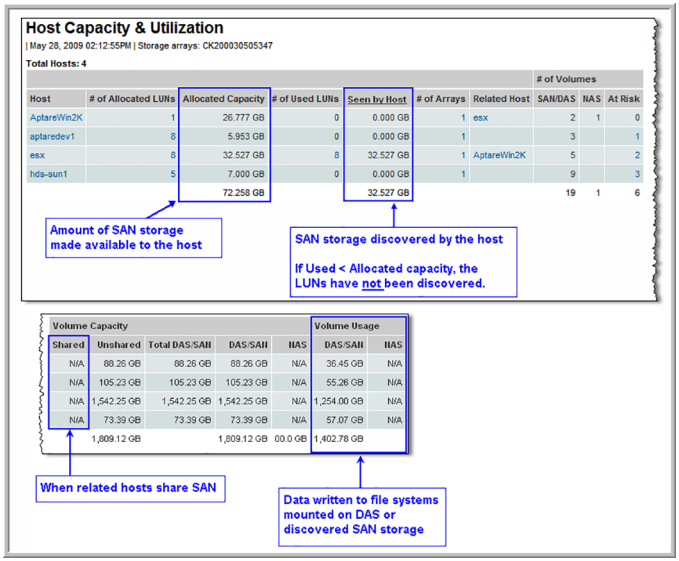


Table 13-1

Column name	Description
Host	The name of the host utilizing the storage. Click on a Host name link to view the details. See " Host Utilization Detail " on page 481.
# of Allocated LUNs	The number of LUNs that have been allocated to this host. Click on this number to launch the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Allocated Capacity	The total capacity of all Volumes both free and used that has been given to the host.
# of Used LUNs	LUNs that are currently used by the host. Click on this number to launch the LUN Utilization Summary. See " LUN Utilization Summary " on page 344.
Seen by Host	The sum of all SAN capacity actually deployed or accepted by the host. This is the storage that the host has been allocated and is using--but this use could be as simple as adding the storage to a volume group. It does not mean that the server has that storage mounted. Mounted storage is reflected in Total DAS/SAN.
# of Arrays	The total arrays associated with this host. Click on the number to launch the Array Capacity and Utilization. See " Array Capacity and Utilization " on page 164.
Related Hosts	Hosts that are related to this host via clustering--they have multi-pathing to the same LUNs for fail over.
# Volumes	
SAN/DAS	The number of SAN/DAS volumes.
NAS	The number of NAS volumes.
At Risk	The number of host volumes that are at risk of reaching capacity, based on configured thresholds.
Volume Capacity	
Shared	Represents the capacity that is shared by related hosts.
Unshared	Capacity that is not shared by related hosts.
Total DAS/SAN	Total of Shared and Unshared volume capacity.

Table 13-1 *(continued)*

Column name	Description
DAS/SAN	DAS/SAN storage capacity for this host.
NAS	Total capacity of external, mounted filesystems (estimated).
Volume Usage	
DAS/SAN	This is the actual amount of SAN/DAS storage used by the host. This may include shared capacity within the cluster.
NAS	The sum of currently allocated storage associated with external, mounted filesystems. The same NAS storage may be accounted for in multiple external filesystem mount points, resulting in double counting.

Host Utilization Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Host Capacity & Utilization > Host Capacity & Utilization

Click a **Host Name** link or **# of Volumes at Risk** link.

In addition, this report can be accessed from several other reports via drill-down links, as shown in the following table:

Table 13-2 Host Utilization Detail

From this Report:	Click this Link:
LUN Utilization Summary	Host
Host Utilization Summary	Host
Array Port Utilization	Host

Table 13-2 Host Utilization Detail *(continued)*

From this Report:	Click this Link:
LUNs At Risk	Host
Hosts At Risk	Name

See "Host Utilization Detail (ASM Example)" on page 490.

Use this Host Utilization Detail report to quickly determine critical capacity issues. In particular, pay attention to the **% Available** column to find specific at-risk volumes. When a host falls below the designated threshold, the data in the % Available column displays in a different color to call your attention to the problem.

Host Details

Table 13-3 Host Details

Column name	Description
Server ID	The internal numeric ID of the server.
Hostname	Host name links to Host Details.
Location	Physical location of the host, if available.
Display Name	The external host/server name.
Description	Description of the server, if available.
Make	Make of the server, if field is populated.
IP Address	Host IP address.
Model	Model of the server, if field is populated.
Type	Type of server, such as TSM Media Server.
Host Type	Type of host.
Virtual Guest	Name of the virtual guest, if relevant.
Operating System	The host's operating system.

The tables that are displayed in this report are dependent on the host data that has been collected in your environment. Therefore, not all tables may be available in a report for a particular host.

- See "Host Details by Source" on page 483.

- See “Host Capacity Detail Report” on page 484.
- See “Network Attached Storage Shares” on page 485.
- See “Allocated and Unused LUNs” on page 486.
- See “Disks Details” on page 487.
- See “Disk Partitions/Volume Group Details” on page 487.
- See “Host Groups” on page 489.

Host Details by Source

Hosts can be collected by any number of sources including vendor products (subsystems), such as VMware, or other sources such as Host Discovery. This table lists the details that have been collected by a particular source so that you can validate the collected details.

Table 13-4 Host Details by Source

Column name	Description
Display Name	Name that is displayed in reports.
Host Name	Name of the host as it was collected by the vendor product.
IP Address	IP address of the host.
Product Group	The group, such as Capacity, for which the host data was collected.
Source	A specific source from which host data was collected; for example, VMware.
Date Created	The date and time the host was created in the Portal database.
Last Updated	The date and time the host data was updated in the Portal database.
OS Platform	Operating system of the host, as collected from this vendor product; for example, Windows Server 2008 R2.
OS Version	The host’s specific operating system version, as collected from this vendor product; for example, 6.01.
Make	The host’s make, as collected from this vendor product; for example,Windows-x64.
Model	The host’s model, as collected from this vendor product; for example, PowerEdge 2950.

Table 13-4 Host Details by Source *(continued)*

Column name	Description
Backup Server Type	The host's role in a backup relationship, client or server.

Host Capacity Detail Report

Table 13-5 Host Capacity Detail Report

Column name	Description
Mount Point	Mount point of the filesystem; for Oracle ASM, the mount point is the name of the ASM instance.
Raw Filesystem	Filesystem name; for Oracle ASM, the raw filesystem is not applicable.
Filesystem Type	When you know the filesystem type, you can determine the host's storage requirements and limitations. Examples of filesystem values: ■ ASM - Oracle ASM ■ NTFS - Windows NT Filesystem ■ ext2, ext3 - Extended Filesystems 2 and 3 ■ ufs - Linux Filesystem ■ vxfs - Veritas Filesystem ■ FAT - File Allocation Table ■ hfs - Hierarchical Filesystem (Apple) ■ ctfs - Cooperative Temporary Filesystem ■ devfs - Device Filesystem ■ hsfs - High Sierra Filesystem ■ jfs - Journaled Filesystem ■ jfs2 - Journaled Filesystem 2 ■ nfs - Network Filesystem ■ objfs - Filesystem Object ■ sfs - Secure Filesystem ■ gpfs - General Parallel Filesystem (IBM) ■ proc - a virtual filesystem ■ procfs - Process Filesystem ■ tmpfs - a virtual memory filesystem ■ smbfs - Server Message Block Filesystem ■ Null - empty ■ NA - not applicable

Table 13-5 Host Capacity Detail Report *(continued)*

Column name	Description
Storage Type	NAS - Network-attached storage; file servers and software dedicated to providing files over a network.SAN - Storage area network; Fibre Channel connected file servers. DAS - direct-attached storage; part of the host computer. MIXED - for Oracle ASM, if heterogenous storage is used.
Volume Group	The volume group and logical volume names. For Oracle ASM, this is the ASM Disk Group name.
Capacity	The host's filesystem storage capacity, as collected by Host Resources data collection. For mounted filesystems, this is the same as `dir` or `df -k`. This capacity value is always from the host's perspective.
Used	The amount of the host's storage that is currently in use.
% Available	The percentage of the capacity that is unused.
90 Days Min Used	Use these three columns in combination to substantiate the need for additional storage. Using 90 days worth of data, this provides a realistic snapshot of how the forecast was derived. These statistics are available after 90 days of data collection.
90 Days Avg Used	The average 90-day usage. Using 90 days worth of data, this provides a realistic snapshot of how the forecast was derived. These statistics are available after 90 days of data collection.
90 Days Max Used	The maximum usage for the past 90-day period. These statistics are available after 90 days of data collection.

Network Attached Storage Shares

This Network Attached Storage (NAS) Shares table provides details similar to the mount point table, without the Volume Group column.

Table 13-6 Network Attached Storage Shares

Column name	Description
Mount Point	Mount point of the filesystem.
Raw Filesystem	Filesystem name
Capacity	The host's filesystem storage capacity.

Table 13-6 Network Attached Storage Shares *(continued)*

Column name	Description
Used	The amount of the host's storage that is currently in use.
% Available	The percentage of the capacity that is unused.
90 Days Min Used	Use these three columns in combination to substantiate the need for additional storage. Using 90 days worth of data, this provides a realistic snapshot of how the forecast was derived. These statistics are available after 90 days of data collection.
90 Days Avg Used	The average 90-day usage. Using 90 days worth of data, this provides a realistic snapshot of how the forecast was derived. These statistics are available after 90 days of data collection.
90 Days Max Used	The maximum usage for the past 90-day period. These statistics are available after 90 days of data collection.

Allocated and Unused LUNs

Use this table to identify LUN storage that may be able to be reclaimed. This table shows any LUNs that have not been partitioned or allocated to a volume group. Therefore, they are not associated with a mount point or file system. This list will be displayed for LUNs that match this "allocated but unused" criteria.

Table 13-7 Allocated and Unused LUNs

Column name	Description
Name	LUN name.
RAID Type	The RAID type of the LUN.
Storage Array	The storage array of the LUN links to the Array Capacity and Utilization.
Array Group	The Array Group for the LUN.
Product	Array product identifier.
Capacity	Capacity of the LUN.
Host	Links to the Host Utilization Detail.

Disks Details

This table lists the disks associated with the host's utilization. Expand the Disk Details (click on the arrow) to view the raw disk paths, storage types, and LUNs associated with disks.

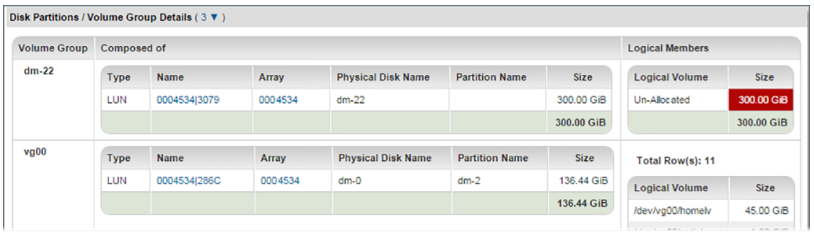
Multi-pathed disks also are represented in this list. Multi-pathing is supported for the following products: EMC PowerPath, Hitachi Dynamic Link Manager (HDLM), and Veritas Dynamic Multi-Pathing (DMP).

Table 13-8 Disk Details

Column name	Description
Disk Name	Disk name.
Disk Make	The vendor information.
Storage Type	DAS, SAN, NAS.
Storage Array	The Array Group for the disk links to the Array Capacity and Utilization. If the disk maps to a storage array, it is a SAN storage type (by WWN).
LUN	Indicates the LUN to which the disk is mapped, as seen by the host; links to the LUN Utilization Summary.
Size	Disk size, as derived from Host Resources data collection.
Raw Disk Path	Disk path. For DAS or SAN, when the disk does not map to a storage array, the raw disk path is used.
Aliases	Aliases for the disk.

Disk Partitions/Volume Group Details

Expand the Volume Group Details (click on the arrow) to identify unallocated space in a Volume Group and logical volumes that are not being used by hosts. These details highlight inefficient use of storage.



This table will be present only when there are one or more volume groups set up on the selected host. The table displays each volume group along with the physical disks that have been allocated to the volume group and the logical volume details. Use this table to find any un-allocated storage--that is, space that is available within the volume group that has not been assigned to a logical volume.

If the Array name is missing from a table entry it is because it is a local disk.

Oracle Containers include storage pools that are represented in the Volume Group column of this table. The corresponding Disk Capacity column displays the virtual disk capacity allocated to the Zone.

Note: Totals in a Volume Group reflect usable capacity--so, for example, RAID 1 mirrored disks will show only the capacity that is usable.

Table 13-9 Disk Partitions/Volume Group Details

Column name	Description
Volume Group	The name of the volume group.
Type	LUN or DSK.
Name	LUN or Disk name.
Array	Storage array; if this field is empty, it means it is a local disk (DAS).
Physical Disk Name	Name of the disk.
Partition Name	The disk partitions that are part of the volume group.
Size	Disk capacity, derived from Host Resources data collection. This is the same as the value returned from `df -k`.
VVM Disk Capacity	Volume manager disk capacity.

Table 13-10 Disk Partitions/Volume Group Details

Column name	Description
RAID Type	Disk or virtual disk RAID type associated with the storage pool.
Disk Name	Name of the disk or virtual disk.
Array	Array where the disk and LUN reside.

Table 13-10 Disk Partitions/Volume Group Details *(continued)*

Column name	Description
Disk Capacity	Capacity that has been allocated. In the case of Oracle containers, this is the capacity allocated to the zones.

Table 13-11 Disk Partitions/Volume Group Details

Column name	Description
Zone	Oracle VM (OS instance)
Volume	If the volume value is blank, the Used value represents the data set used capacity. If there is a value listed for the volume, this means it is for a file system.
Used	Represents the capacity “written to” the file system.

Table 13-12 Disk Partitions/Volume Group Details

Column name	Description
Logical Volume	Logical volume name. If listed as Un-allocated, it is not mounted.
Size	Usable volume space.
Physical Size	Physical size of logical member.

Note: The Capacity and Used columns are aggregated and then converted to the units selected in the Advanced option of the report from which the Host Utilization Detail was launched. Note that in the conversion to the displayed units, the totals may appear to be slightly inaccurate. This is due to the division or multiplication by 1024.

Host Groups

Displays the host groups for which the host is a member.



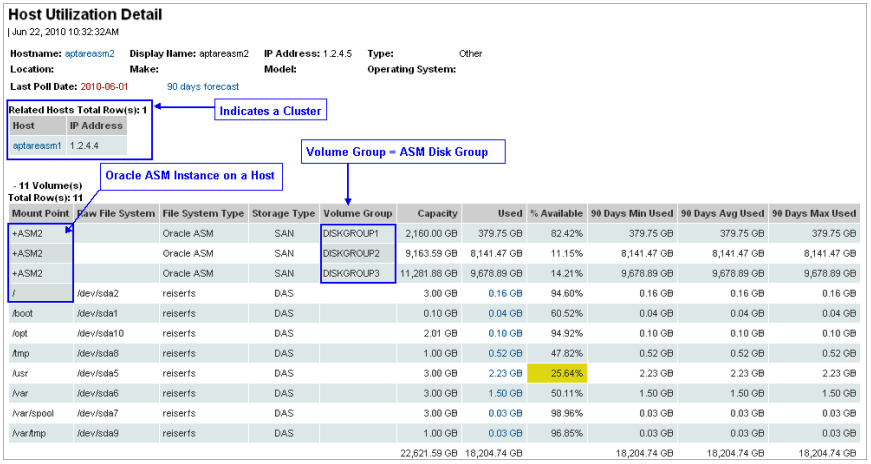
Table 13-13 Host Groups

Column name	Description
R1C1 Member of the Following Groups	R1C2 The full pathname of the host group links to additional host group details.
Member Since	Indicates when this host became a member of the group.
Added By	The user who added the host to the group.

Host Utilization Detail (ASM Example)

NetBackup IT Analytics reveals the following:

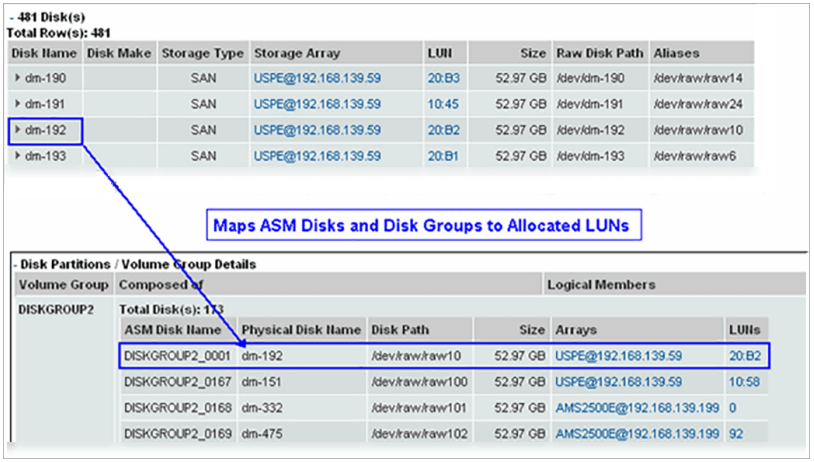
- ASM Storage Usage that the Filesystem Does Not Report: Oracle ASM does not use a mount-point for storage usage. Since this storage is managed by Oracle ASM, it is unknown to the filesystem.
- LUN-to-Host Perspective: The number of LUNs used for the ASM instance on the host provides a complete picture of storage usage. Without the exposure that NetBackup IT Analytics offers, the number of allocated/used LUNs would be smaller--that is, your view of storage usage would be incomplete without the inclusion of ASM usage.
- Mapping of ASM Volume to Storage Tier: By exposing the Array storage used for Oracle ASM instances, capacity and chargeback numbers become more comprehensive and accurate.



Along with traditional filesystem usage, the NetBackup IT Analytics Host Utilization Detail report incorporates ASM Disk Group usage. In the above example, the first three rows of the Volume table display ASM Disk Groups. This Host Utilization Detail report provides additional physical disk details, raw disk paths, and aliases (see example below).

- Mapping of disk paths to storage array LUNs
- View of shared storage details

This host-centric view provides:



Veritas Volume Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Host Capacity & Utilization > Host Capacity & Utilization

Click a Host link to view the Host Utilization Detail report.

Click the **Logical Volume** link in the Disk Partitions/Volume Group Details table.

Table 13-14 Veritas Volume Details

Column name	Description
Partition Name	The Veritas Volume Manager partition name.
Size	Size of the partition.
Raid Type	RAID type, such as mirror.
Plex	The Plex that contains the Sub Disks.
Raid Type	RAID type, such as Striped or Concat.
Size	Capacity of the Plex.
Sub Disk	Sub Disk name.
Sub Disk Size	Capacity of the Sub Disk.
Disk	Disk ID.

Host Executive Summary

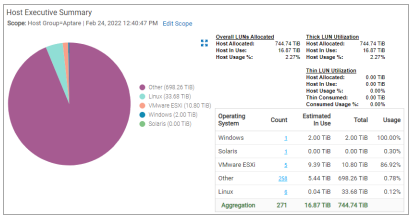
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Host Capacity & Utilization > Host Executive Summary

The Host Executive Summary provides a single view of hosts' thick and thin storage usage by operating system.

In the report's scope selector, the Host Type option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other.



In this report, the following definitions apply:

- Thick - Capacity that has been allocated by a traditional array group LUN
- Thin - Capacity that has been allocated by a thin-provisioned LUN

Table 13-15 Host Executive Summary

Column name	Description
Operating System	Each row represents the hosts, by operating system, and the storage that they are consuming.
Count	The number of hosts for this operating system; click on the link to view the Host Capacity and Utilization report.
Estimated In Use	The amount of used SAN filesystem capacity used by hosts.
Total	The total amount of storage allocated to hosts of this operating system.
Usage	The percentage of the total storage for hosts (by operating system) that is in use, calculated as In-Use divided by Total Capacity.
Overall LUNs Allocated	
Host Allocated	Total amount of recognized capacity allocated to the hosts in this report scope.

Table 13-15 Host Executive Summary *(continued)*

Column name	Description
Host In Use	Total amount of recognized capacity actually in use by the hosts in this report scope.
Host Usage %	Percentage of total capacity used, calculated as Host In-Use divided by Host Allocated.
Thick LUN Utilization	
Host Allocated	The total amount of recognized thick capacity that has been allocated to hosts in this report scope.
Host in Use	The amount of recognized thick storage that has been written to by the hosts in this report scope.
Host Usage %	Thick Host In-Use divided by Thick Host Allocated.
Thin LUN Utilization	
Host Allocated	The amount of thin-provisioned storage that has been presented to hosts.
Host In Use	The total amount of thin-provisioned (dynamically provisioned) storage that has been written to by hosts.
Host Usage %	Thin Host In-Use divided by Thin Host Allocated.
Thin Consumed	The total amount of thin-provisioned storage that has been written to by hosts.
Consumed Usage %	Percentage of consumed thin capacity. Thin Host In-Use divided by Thin Consumed.

Host HBA Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Host Capacity & Utilization > Host HBA Summary

Table 13-16 Host HBA Summary

Column name	Description
Name	Host name.
Device Name	The HBA device name collected from the host.
Node WWN	The node's worldwide name.
Port WWN	The port's worldwide name.
Node Nickname	The node's nickname, if there is one assigned.
Driver	HBA driver.
Driver Version	HBA driver version.
Firmware Version	Firmware version.
HBA Model	Model of the host bus adapter.
Manufacturer	HBA manufacturer.
Vendor	Vendor of the host bus adapter.
Last Updated	Last time data collection updated these values in the portal database.

Chapter 14

SnapMirror Reports

This chapter includes the following topics:

- SnapMirror Summary
- SnapMirror Details
- SnapMirror Schedule
- SnapMirror Status Summary
- SnapMirror Data Transfer
- SnapMirror Volumes/QTrees At Risk
- SnapMirror Data Protection Dashboard

SnapMirror Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Summary

The SnapMirror Summary provides a listing of the status of SnapMirror transactions. It lists the source-destination pairs, along with the most frequently used relevant metrics. Choose to report by asynchronous and/or synchronous schedules. For additional details about a Failed transaction, mouse over the Summary Status to view the message.

Table 14-1 SnapMirror Summary

Column name	Description
Source Location	The source location will be displayed in one of the following formats: NetApp Storage System: Volume Name and NetApp Storage System: Volume Name/QTree Name See " SnapVault Secondary Configuration " on page 508.
Destination Location	The destination location will be displayed in one of the following formats: NetApp Storage System: Volume Name and NetApp and Storage System: Volume Name/QTree Name See " SnapVault Secondary Configuration " on page 508.
Mirror Timestamp	Date and time of last SnapMirror transaction, including unsuccessful mirrors. See " SnapMirror Details " on page 498.
Contents	The active filesystem state of the destination: replica, transitioning, or original
Last Transfer Size	Size of the last SnapMirror transfer
Last Transfer Duration (Secs)	Duration in seconds of the last SnapMirror transfer
Transfer KB/Secs	Rate of data transfer
Base Snapshot	The name of the last snapshot transferred from source to destination
State	The state of the SnapMirror pair: un-initialized, snapmirrored, broken-off, quiesced, source, or unknown
Status	Transfer completion status for the source-destination pair: idle, transferring, pending, aborting, migrating, quiescing, resyncing, waiting, syncing, in-sync, or paused.
Summary Status	Mouse over the Failed status to reveal the actual error message

Table 14-1 SnapMirror Summary *(continued)*

Column name	Description
Lag Time	The length of time since the last SnapMirror transfer
Last Updated	The last time the Data Collector updated the Portal database

SnapMirror Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Summary

Click a **Mirror Timestamp** link.

The SnapMirror Summary provides the details of a SnapMirror transaction.

Table 14-2 SnapMirror Details

Column name	Description
Source Storage Array	The name of the Source storage array links to the NetApp Storage System Detail.
Destination Storage Array	The name of the Destination storage array links to the NetApp Storage System Detail.
Source Volume	The source volume name links to the NetApp Volume Detail.
Destination Volume	The destination volume name links to the NetApp Volume Detail.

Table 14-2 SnapMirror Details *(continued)*

Column name	Description
Source QTree	The source QTree name links to the NetApp QTree Summary.
Destination QTree	The destination QTree name links to the NetApp QTree Summary.
Contents	The active filesystem state of the destination: replica, transitioning, or original.
Mirror Timestamp	Date and time of last SnapMirror transaction, including unsuccessful mirrors.
State	The state of the SnapMirror pair: un-initialized, snapmirrored, broken-off, quiesced, source, or unknown.
Status	Transfer completion status for the source-destination pair: idle, transferring, pending, aborting, migrating, quiescing, resyncing, waiting, syncing, in-sync, or paused.
Transfer Progress	Indicates if a current transfer is in progress.
Current Transfer Type	The current SnapMirror transfer type: initialize, store, schedule, retry, retrieve, resync, or migrate. A value is displayed only if there is a transfer in progress.
Last Transfer Duration (Seconds)	Length of time for the last mirror data transfer.
Last Transfer From	The source location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name.
Last Transfer Size	Amount of data in the last mirror data transfer.
Last Updated	The last time the Data Collector updated the Portal database.

SnapMirror Schedule

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Schedule

List the SnapMirror Schedules, showing error/success status. Use this report to identify replication exposure. Choose to report by asynchronous and/or synchronous schedules.

Note that the SnapMirror schedule is initiated by the Destination system.

Table 14-3 SnapMirror Schedule

Column name	Description
Source Location	The source location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name. See " NetApp Volume Details " on page 398.
Destination Location	The destination location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name. See " NetApp Volume Details " on page 398.
Ops Throttle	For synchronous schedules, this value indicates the number of outstanding operations allowed before blocking occurs on the source. This may be shown as ops (number of operations), s (number of seconds), ms (number of milliseconds).
Visibility Frequency	For synchronous schedules, this frequency (in seconds) controls how often the source snapshot will be visible on the destination mirror.

Table 14-3 SnapMirror Schedule *(continued)*

Column name	Description
Day of the Month	For Asynchronous schedules. The notation is similar to crontab format. 1 means match day 1; 1,3 means match days 1 and 3; 1-3 means match days 1, 2, 3; 1-30/7 means match every 7th day until the 30th day; * means match all possible values
Hours	For Asynchronous schedules. The notation is similar to crontab format: 1 means match hour 1; 1,3 means match hours 1 and 3; 1-3 means match hours 1, 2, 3; 1-24/4 means match every 4th hour until the end of the day; * means match all possible values
Minutes	For Asynchronous schedules. The notation is similar to crontab format: 1 means match minute 1; 1,3 means match minutes 1 and 3; 1-3 means match minutes 1, 2, 3; 0-55/5 means match every 5th minute until the 55th minute; * means match all possible values
Max Transfer Rate	The maximum transfer rate in KB/sec. Typically, this is the fastest possible rate of the storage system.
Error	Error message displayed only if there is a transfer error
Last Successful	Date/timestamp of the last successful SnapMirror transfer See " SnapMirror Details " on page 498.

Table 14-3 SnapMirror Schedule *(continued)*

Column name	Description
Last Updated	The last time the Data Collector updated the Portal database

SnapMirror Status Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

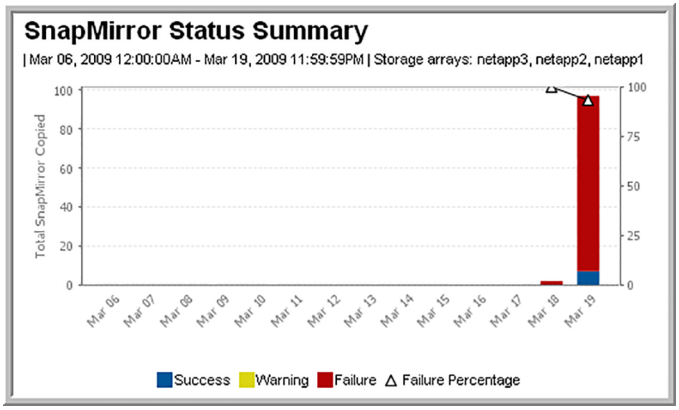
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Status

Note: SnapMirror Status is based on the last sampling of data. NetApp keeps track of the latest SnapMirror schedule, however, the data shown in this report reflects the history of what was captured from the NetApp SnapMirror data. In some cases, a SnapMirror transaction may have occurred between data collection cycles.

This bar chart provides an at-a-glance view of the success or failure for each selected time frame, such as days, weeks, or months. Mouse over the colored section of each bar to view the numeric details. Click on a colored section of a bar to drill down to the summary.

See " SnapMirror Summary " on page 496.



SnapMirror Data Transfer

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

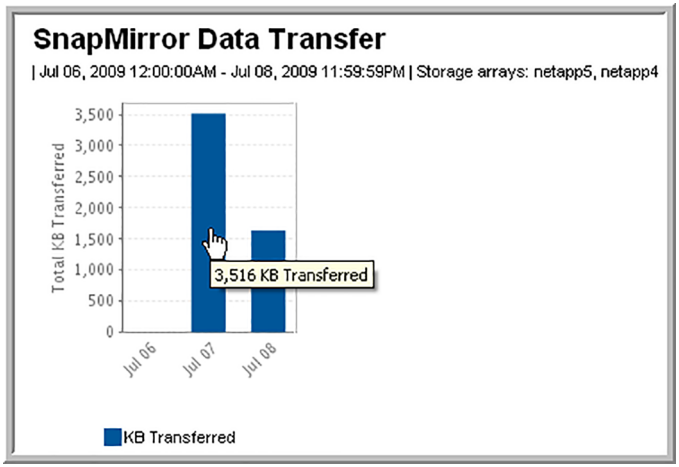
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Data Transfer

This bar chart provides an at-a-glance view of the amount of SnapMirror data transferred from NetApp source storage systems to destination storage systems.

Click a bar to drill down to the summary for the corresponding time period.

See " SnapMirror Summary " on page 496.



SnapMirror Volumes/QTrees At Risk

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Volume/QTrees At Risk

Use this report to pro-actively track the storage being consumed by SnapMirror transactions. Identify SnapMirror Volume and QTree usage that is expected to exceed capacity.

To configure this report to alert you when Volumes and/or QTrees are at risk, generate the report and save an instance. Then, configure an alert to email you when this report is populated with volumes or QTrees at risk.

Table 14-4 SnapMirror Volumes/QTrees At Risk

Column name	Description
Source Capacity	Capacity of the SnapMirror source
Source Capacity Used	Amount of the source that is in use
Risk Date	The point at which SnapMirror storage will be limited
Destination Capacity	Capacity of the SnapMirror destination
Destination Capacity Used	Amount of the destination that is in use
Space Guarantee on Destination	Indicates if the destination volume has a space guarantee enabled
Source Location	The source location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name
Destination Location	The destination location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name
Base Snapshot	The name of the last snapshot transferred from source to destination
Last Updated	The last time the Data Collector updated the Portal database

SnapMirror Data Protection Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapMirror Reports > SnapMirror Data Protection Dashboard

This report provides an at-a-glance view of the status of your SnapMirror data transfers, with drilldowns to details so that you can quickly resolve issues.

Click a status icon to drill down to the summary.

See " SnapMirror Summary " on page 496.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Chapter 15

SnapVault Reports

This chapter includes the following topics:

- SnapVault Relationship Summary
- SnapVault Secondary Configuration
- SnapVault Summary
- SnapVault Detail
- SnapVault Schedule
- SnapVault Status
- SnapVault Data Transfer
- SnapVault Relationship Summary
- SnapVault Data Protection Dashboard
- SnapVault QTrees at Risk

SnapVault Relationship Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapVault Reports > SnapVault Relationship Summary

List status for source-destination pairs for SnapVault transactions. Choose to report by volume or QTree and select the mirror states to examine. For additional details about a Failed transaction, mouse over the Summary Status to view the message.

Table 15-1 SnapVault Relationship Summary

Column name	Description
Source Location	The source location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name
Destination Location	The destination location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name
Mirror Timestamp	Date and time of last SnapMirror transaction, including unsuccessful mirrors
Contents	The active filesystem state of the destination: replica, transitioning, or original.
Last Transfer Size (KB)	Amount of data in the last mirror transfer.
Last Transfer Duration (Secs)	Length of time for the last mirror data transfer.
Transfer KB/Secs	Rate of mirror data transfer.
Base Snapshot	The name of the last snapshot transferred from source to destination
State	uninitialized, snapmirrored, broken-off, quiesced, source, unknown.
Status	Transfer completion status for the source-destination pair: idle, transferring, pending, aborting, migrating, quiescing, resyncing, waiting, syncing, in-sync, or paused.
Summary Status	Success, Failed - Mouse over the Failed status to view the error message that contains the details of why the transfer failed.
Lag Time	The length of time since the last SnapVault transfer.

Table 15-1 SnapVault Relationship Summary *(continued)*

Column name	Description
Last Updated	The last time the Data Collector updated the Portal database.

SnapVault Secondary Configuration

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapVault Reports > SnapVault Sec Configuration

Use this report to get an overview of the state of SnapVault secondary configurations. List SnapVault primary and secondary paths, with success/failure statistics.

Table 15-2 SnapVault Secondary Configuration

Column name	Description
Primary Path	This path is used for the source of the baseline transfer.
Primary System	The primary system of the SnapVault relationship and also the baseline transfer.
Secondary Path	This path is the destination of the baseline transfer and as such, is created during the baseline transfer
Max Transfer Rate	Maximum transfer rate; KB treated as 1024 bytes
Tries Count	Indicates the maximum number of times this transfer will be attempted, as set in the SnapVault configuration. If a transfer isn't successful after this number of tries, the secondary will not try again and all the data that was transferred on previous attempts will be removed.
Last Successful Date	Timestamp of the last successful transfer.

Table 15-2 SnapVault Secondary Configuration *(continued)*

Column name	Description
Last Transfer Size	Amount of data in the last SnapVault transfer.
Lag Time	The amount of time from the start of the most recent successful transfer.
Last Base SnapShot Name	The name of the last SnapShot for this primary system.
Last Updated	The last time the Data Collector updated the Portal database.

SnapVault Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapVault Reports > SnapVault Summary

List SnapVault events with status and performance data. View a list of SnapVault transactions, by destination location.

Table 15-3 SnapVault Summary

Column name	Description
Source Location	Links to the NetApp QTree Summary for the source. The location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name
Destination Location	The destination location will be displayed in one of the following formats: NetApp Storage System: Volume Name or NetApp Storage System: Volume Name/QTree Name
Mirror Timestamp	The snapshot's creation date and time (for the most recent successful transfer); links to the SnapVault Detail.

Table 15-3 SnapVault Summary *(continued)*

Column name	Description
Contents	The active filesystem state of the destination: replica, transitioning, or original
Last Transfer Size	Amount of data in the last SnapVault transfer
Last Transfer Duration (Secs)	Length of time for the last SnapVault data transfer
Transfer KB/Sec	Rate of SnapVault data transfer
Base Snapshot	The name of the last snapshot transferred from source to destination
State	uninitialized, snapvaulted, broken-off, quiesced, source, restoring, or unknown
Status	aborting, idle, pending, quiescing, restoring, resyncing, or transferring
Summary Status	Success, Failed - Mouse over the Failed status to view the error message that contains the details of why the transfer failed
Lag Time	The length of time since the beginning of the last successful transfer from the source. This field is populated only for successful baseline transfers.
Last Updated	The last time the Data Collector updated the Portal database.

SnapVault Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > SnapVault Reports > SnapVault Detail

Click a **Mirror Timestamp** link.

Table 15-4 SnapVault Detail

Column name	Description
Source System	Hostname of the source system.
Source Volume	Volume of the source system.
Source QTree	QTree of the source system.
Contents	The active filesystem state of the destination: replica, transitioning, or original.
State	The state of the source-destination relationship: unitialized, snapvaulted, broken-off, unknown, source, restoring, quiesced.
Transfer Progress	Indicates the amount of data transferred so far, available only for active transfers.
Destination System	Hostname of the destination system.
Destination Volume	Volume of the destination system.
Destination QTree	QTree of the destination system.
Mirror Timestamp	The snapshot's creation date and time (for the most recent successful transfer).
Status	The status of the mirror transfer: aborting, broken-off, idle, pending, quiescing, resynching, or restoring.
Current Transfer Type	This field is populated only for active transfers. initialize, update, retry, or resync
Last Transfer Duration (Secs)	Duration of the last transfer.
Last Transfer Size	Amount of data transferred in the last SnapVault transaction--in kilobytes (1024 bytes).

Table 15-4 SnapVault Detail *(continued)*

Column name	Description
Last Transfer From	The name of the last transfer's source system.
Last Transfer Type	initialize, update, retry, or resync.
Last Updated	The last time the Data Collector updated the Portal database.

Schedule Report (Primary and Secondary)

For a description of the fields in this table refer to the following.

See " SnapVault Schedule " on page 512.

SnapVault Schedule

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Replication Manager > SnapVault Reports > SnapVault Schedule

Capacity Manager > SnapVault Reports > SnapVault Schedule

List the SnapVault schedules with array and volume information. View primary or secondary SnapVault schedules for an overview of the associated volumes and retention count.

Table 15-5 SnapVault Schedule

Column name	Description
Schedule Name	Unique identifier of a SnapVault schedule; this name serves as a prefix in the name of snapshots created by this schedule
Storage Array	Link to the Storage System Details.
Volume Name	The volume related to this schedule.

Table 15-5 SnapVault Schedule *(continued)*

Column name	Description
Day of Week	A comma-separated list or a range of days (3-letter abbreviation) for this schedule.
Hour of Day	A comma-separated list or a range of hours for this schedule, where 0 represents midnight.
Retention Count	The maximum number of snapshots that will be kept for this schedule.
Status	Status of the schedule: Aborting, Active, Idle, Queued, or Saving.
Last Updated	The last time the Data Collector updated the Portal database.

SnapVault Status

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Replication Manager > SnapVault Reports > SnapVault Status

Capacity Manager > SnapVault Reports > SnapVault Status

Displays a bar chart of SnapVault status, with drilldown access to details. Click a colored section of a bar to drill down to the summary. This at-a-glance view lets you quickly identify warnings and failures that need to be investigated.

See “ SnapVault Summary ” on page 509.

SnapVault Data Transfer

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapVault Reports > SnapVault Data Transfer

Displays an at-a-glance view of the amount of SnapVault data transferred. Mouse over a bar to view the amount of data transferred. This bar chart represents the amount of data transferred from NetApp source to destination systems.

Click a bar to view the summary.

See " SnapVault Summary " on page 509.

SnapVault Relationship Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Replication Manager > SnapVault Reports > SnapVault Relationship Summary

Capacity Manager > SnapVault Reports > SnapVault Relationship Summary

Table 15-6 SnapVault Relationship Summary

Column name	Description
Source Location	Links to the NetApp QTree Summary for the source.
Destination Location	Links to the NetApp QTree Summary for the destination.
Mirror Timestamp	The snapshot's creation date and time (for the most recent successful transfer).
Contents	State of the active filesystem on the destination location: original, replica, or transitioning.
Last Transfer Size	Amount of data in the last transfer.
Last Transfer Duration (Secs)	Amount of time taken for the last data transfer.

Table 15-6 SnapVault Relationship Summary *(continued)*

Column name	Description
Transfer KB/Sec	Transfer rate.
Base Snapshot	The name of the last snapshot transferred from source to destination.
State	State of the SnapVault-pair: broken-off, quiesced, restoring, source, snapvaulted, uninitialized, unknown.
Status	Status of the SnapVault relationship: Success, Warning, Failure. Mouse over the Failed error message to view the details.
Summary Status	Mouse over a Failed summary status to view the error details.
Lag Time	The length of time since the beginning of the last successful transfer from the source. This information is available only if there was a successful baseline transfer.
Last Updated	The last time the Data Collector updated the Portal database.

SnapVault Data Protection Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapVault Reports > SnapVault Data Protection Dashboard

Click a colored status indicator to view the SnapVault Summary, which lists details about the snapshot transfers.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

SnapVault QTrees at Risk

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > SnapVault Reports > SnapVault QTrees At Risk

List SnapVault QTrees that are at risk of exceeding available capacity. Identify the risk data; that is, the date and time that the destination location is expected to have a limited amount of storage available.

Table 15-7 SnapVault QTrees at Risk

Column name	Description
Source Location	Path of the location of the SnapVault source.
Source Capacity	Capacity of the SnapVault source.
Source Capacity Used	Amount of the source capacity that is in use.
Destination Location	Path of the SnapVault destination.
Destination Capacity	Capacity of the SnapVault destination.
Destination Capacity Used	Amount of the destination capacity that is in use.
Base Snapshot	The name of the last snapshot transferred from source to destination.
Risk Date	The date and time that the destination location is expected to have a limited amount of storage available.
Last Updated	The last time the Data Collector updated the Portal database.

Chapter 16

Capacity Forecasting Reports

This chapter includes the following topics:

- Array Capacity Forecast
- Data Domain Filesystem Capacity & Forecast
- Data Domain Filesystem Status
- EMC Symmetrix Thin Pool Capacity Forecast
- EMC VNX (CLARiiON) Thin Pool Capacity Forecast
- Object Store Usage Over Time
- Hitachi NAS Storage Pool Capacity Over Time
- Host Capacity Forecast
- Host Volume Growth and Forecast
- Host Storage Usage Over Time
- Storage Capacity Over Time
- IBM COS Storage Pool Growth and Forecast
- IBM COS Vault Growth and Forecast
- IBM COS Vault Summary Growth and Usage
- IBM COS Site Summary
- IBM COS Device Summary

- IBM COS Mirror Summary

Array Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

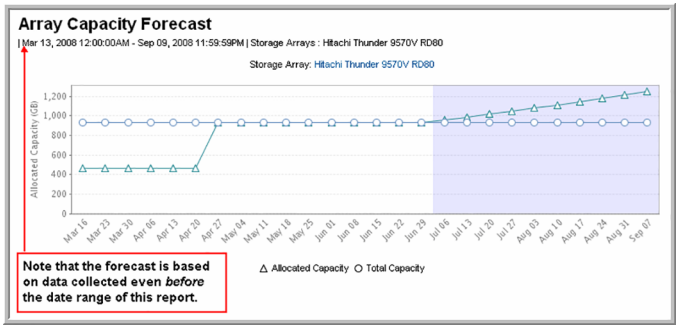
Capacity Manager > Storage Capacity & Forecast > Array Capacity Forecast

Use the Capacity Forecast report to project growth in storage usage. You can generate the report by the following criteria:

- Historical date range
- Forecast horizon (date range)
- Storage Array Group
- Forecast by Array Capacity or LUN Capacity

In addition, you can choose to generate graphs for each, individual array or instead, display a consolidated graph for all the arrays selected in the scope selector.

Note: This report is limited to a maximum of 50 arrays.



Mouse-over the markers on the lines to view:

- Historical Usage: Allocated and Total Capacity
- Forecasted Usage: Allocated and Total Capacity

The Array name links to the Array Capacity and Utilization.

See " Array Capacity and Utilization " on page 164.

Table 16-1 Array Capacity Forecast

Column name	Description
Array	The name of the storage array. See " Array Capacity and Utilization " on page 164.
△ Allocated Capacity	This line tracks the historical and forecasted allocation.
○ Total Capacity	This line tracks the historical and forecasted capacity.

Data Domain Filesystem Capacity & Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Forecasting & Capacity Planning > Data Domain Filesystem Capacity & Forecast

This chart shows the total physical disk space, after de-duplication and compression, available for data storage.

The scope selector provides a drop down Filesystem Names list to enable scope selection from the following:

- Data - This is the amount of total physical disk space available for data, the actual physical space used for compressed data, and physical space still available for data storage.
- /ddvar - The approximate amount of space used by and available to log and core files. Storage often can be reclaimed by removing old logs and core files.
- Index - This data is only collected when SNMP is configured for data collection.
- Metadata - This data is only collected when SNMP is configured for data collection.

Table 16-2 Data Domain Filesystem Capacity & Forecast

Column name	Description
System	The Data Domain system links to the details. See " Data Domain System Details" on page 246.
Used	The percentage of the filesystem capacity that is used; click this bar to view the status. See " Data Domain Filesystem Status " on page 520.
Available	Total disk space available for data storage.
Used Forecast	Forecast based on past usage.
Available Forecast	Forecast based on past usage.

Data Domain Filesystem Status

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Forecasting & Capacity Planning > Data Domain Filesystem Capacity & Forecast

Click a bar in the chart.

Table 16-3 Data Domain Filesystem Status

Column name	Description
Filesystem Name	data, /ddvar, index, or metadata.
Date	Start date of status.
Available	Filesystem size.

Table 16-3 Data Domain Filesystem Status *(continued)*

Column name	Description
Used	Amount of filesystem that is in use.
Cleanable	Potential cleanable space.
Pre-Comp Size	Pre-compression size.

EMC Symmetrix Thin Pool Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

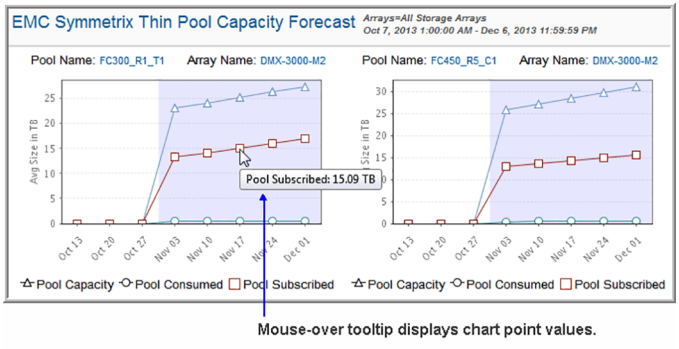
Capacity Manager > > EMC Symmetrix Thin Pool Capacity Forecast

View this report to quickly assess if EMC Symmetrix subscribed storage is close to exceeding physical capacity.

Note that because the rendering of charts consumes memory resources, it is recommended that you manage the scope of this report with the following Report Designer parameters:

- Time Period
- Group by

Reducing the Time Period also enables a more focused view of thin-provisioned storage usage and forecast for multiple thin pools.



EMC VNX (CLARiiON) Thin Pool Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

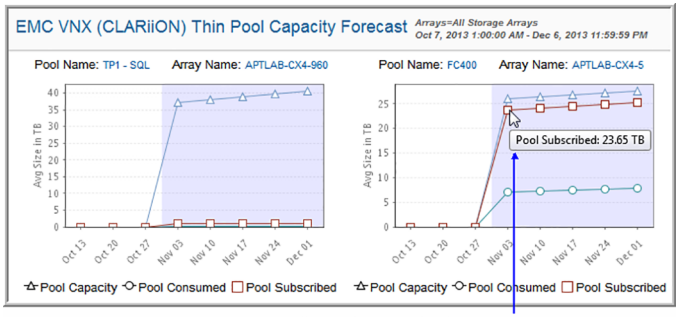
Capacity Manager > EMC VNX (CLARiiON) Thin Pool Capacity Forecast

View this report to quickly assess if EMC VNX CLARiiON subscribed storage is close to exceeding physical capacity.

Note that because the rendering of charts consumes memory resources, it is recommended that you manage the scope of this report with the following Report Designer parameters:

- Time Period
- Group by

Reducing the Time Period also enables a more focused view of thin-provisioned storage usage and forecast for multiple thin pools.



Mouse-over tooltip displays chart point values.

Object Store Usage Over Time

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Capacity & Forecast > Object Store Usage Over Time

Graphically represent object array capacity usage for all nodes, over time. Compare used capacity against total capacity to assess future capacity needs. For example, view Hitachi Content Platform (HCP) used and total capacity.

Table 16-4 Object Store Usage Over Time

Column name	Description
Total Capacity	Total amount of storage space on all the logical volumes managed by the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
Used Capacity	Total amount of used space on all the logical volumes managed by the node. Values are stored as KiB in the database and rendered according to your user profile preferences.
Log Date	The date and time the data was logged in the NetBackup IT Analytics database.

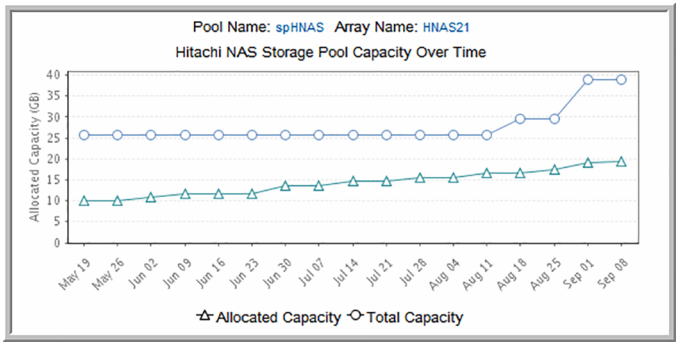
Hitachi NAS Storage Pool Capacity Over Time

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Capacity & Forecast > Hitachi NAS Storage Pool Capacity Over Time

This report provides an at-a-glance history of allocated and total capacity over time.



- Click the **Pool Name** link to view the Hitachi NAS Storage Pool Summary report.
- Click the **Array Name** link to view the Array Capacity & Utilization report.

Host Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Capacity & Forecast > Host Capacity Forecast

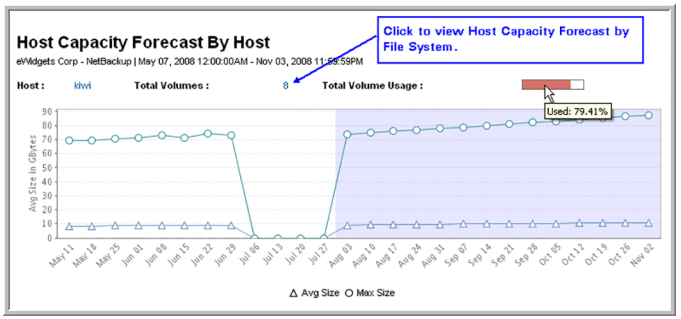
This forecast report provides an at-a-glance overview of historical and predicted usage.

Note: To ensure that your report reflects both historical and forecasted usage, when you generate the report, verify that your date range starts on a past date and ends on a future date. The current day is treated as the future for the purpose of forecasting. Typically, if you specify three months before and after the current day, you'll have sufficient data for a valid forecast. The default display is **Grouped by Months**. However, to see results based on previous weeks, choose **Group by Weeks** in the Report Designer. The Group by Weeks option is especially useful during the first few weeks of Capacity Manager deployment.

In the report's scope selector, the **Host Type** option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other. You also

can narrow the scope of this report by specifying a storage type filter of: SAN, DAS, or NAS.

When you generate this report, you can select the type of output--either **by host** or **by filesystem**. If you select filesystem, the specific path, host name, and mount point will be associated with each graph.



1. Mouse-over the thermometer to view usage.
2. Mouse-over the markers on the lines to view:
   - Historical Usage: Average and Maximum Size
   - Forecasted Usage: Average and Maximum Size
3. Click the **Host** link to launch the detail report.

   See " Host Utilization Detail " on page 481.
4. Click the **Total Volumes** link to view the Host Capacity Forecast by Filesystem.

Table 16-5 Host Capacity Forecast

Column name	Description
Date	The first day of the period for which the usage is projected.
Avg Size	The historical average usage.
Max Size	The historical maximum usage.
Avg Size Future	The projected usage based on the average usage history (or historical average usage for periods in history).
Max Size Future	The projected usage based on the maximum usage history (or historical average usage for periods in history).

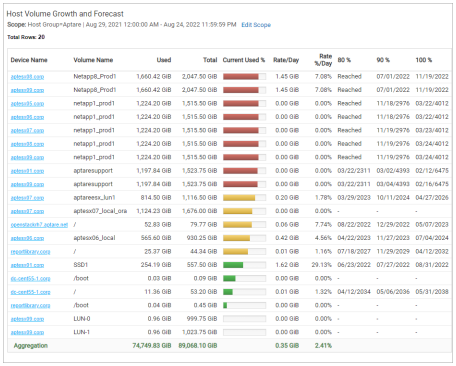
Host Volume Growth and Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Capacity & Forecast > Host Volume Growth and Forecast

This report lists from the host perspective, all discovered file systems, capacities, growth rates and projected out-of-space conditions allowing you to monitor system (physical and virtual) volume usage and avoid service outages. The time frame for the forecast is set in scope selector when you generate a report. In the report's scope selector, the **Host Type** option can be used to include virtual hosts in this host-based report: VM Server, VM Guest, VIO Server, VIO Guest, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest, or Other.


Host Volume Growth and Forecast

Scope: Host Group=Aptare | Aug 29, 2021 12:00:00 AM - Aug 24, 2022 11:59:59 PM Edit Scope

Total Rows: 20

Device Name	Volume Name	Used	Total	Current Used %	Rate/Day	Rate %/Day	80 %	90 %	100 %
aptess98.corp	Netapp8_Prod1	1,660.42 GiB	2,047.50 GiB		1.45 GiB	7.08%	Reached	07/01/2022	11/19/2022
aptess99.corp	Netapp8_Prod1	1,660.42 GiB	2,047.50 GiB		1.45 GiB	7.08%	Reached	07/01/2022	11/19/2022
aptess95.corp	netapp1_prod1	1,224.20 GiB	1,515.50 GiB		0.00 GiB	0.00%	Reached	11/18/2976	03/22/4012
aptess96.corp	netapp1_prod1	1,224.20 GiB	1,515.50 GiB		0.00 GiB	0.00%	Reached	11/18/2976	03/22/4012
aptess97.corp	netapp1_prod1	1,224.20 GiB	1,515.50 GiB		0.00 GiB	0.00%	Reached	11/19/2976	03/23/4012
aptess98.corp	netapp1_prod1	1,224.20 GiB	1,515.50 GiB		0.00 GiB	0.00%	Reached	11/19/2976	03/24/4012
aptess99.corp	netapp1_prod1	1,224.20 GiB	1,515.50 GiB		0.00 GiB	0.00%	Reached	11/19/2976	03/24/4012
aptess91.corp	aptaresupport	1,197.84 GiB	1,523.75 GiB		0.00 GiB	0.00%	03/22/2311	03/02/4393	02/12/6475
aptess93.corp	aptaresupport	1,197.84 GiB	1,523.75 GiB		0.00 GiB	0.00%	03/22/2311	03/04/4393	02/16/6475
aptess97.corp	aptaressx_lun1	814.50 GiB	1,116.50 GiB		0.20 GiB	1.78%	03/29/2023	10/11/2024	04/27/2026
aptess07.corp	aptissx07_local_ora	1,124.23 GiB	1,676.00 GiB		0.00 GiB	0.00%	-	-	-
openstackrh7.aptare.net	/	52.83 GiB	79.77 GiB		0.06 GiB	7.74%	08/22/2022	12/29/2022	05/07/2023
aptess06.corp	aptissx06_local	565.60 GiB	930.25 GiB		0.42 GiB	4.56%	04/22/2023	11/27/2023	07/04/2024
repofilecarv.corp	/	25.37 GiB	44.34 GiB		0.01 GiB	1.16%	07/18/2027	11/29/2029	04/12/2032
aptess91.corp	SSD1	254.19 GiB	557.50 GiB		1.62 GiB	29.13%	06/23/2022	07/27/2022	08/31/2022
ds-cent65-1.corp	/boot	0.09 GiB	0.09 GiB		0.00 GiB	0.00%	-	-	-
ds-cent65-1.corp	/	11.36 GiB	53.20 GiB		0.01 GiB	1.32%	04/12/2034	05/06/2036	05/31/2038
repofilecarv.corp	/boot	0.04 GiB	0.45 GiB		0.00 GiB	0.00%	-	-	-
aptess99.corp	LUN-0	0.96 GiB	999.75 GiB		0.00 GiB	0.00%	-	-	-
aptess99.corp	LUN-1	0.96 GiB	1,023.75 GiB		0.00 GiB	0.00%	-	-	-
Aggregation		74,749.83 GiB	89,068.10 GiB		0.35 GiB	2.41%			

Table 16-6 Host Volume Growth and Forecast

Column name	Description
Device Name	Host name.
Volume Name	Volume name.
Used	Amount of used capacity for the file system from host perspective. The value for the Used column is always the current used, but the projection is based on the time period set in the scope selector.
Total	Total allocated capacity of the file system from host perspective.

Table 16-6 Host Volume Growth and Forecast *(continued)*

Column name	Description
Current Used %	Thermometer 0% to 100%, colored green, yellow, red according to the current utilization of the file system.
Rate/Day	Average daily growth of the file system.
Rate %/Day	Average daily percentage increase or decrease for the file system.
80%	Expected date when the file system will reach 80% capacity, or an indication it has been reached.
90%	Expected date when the file system will reach 90% capacity, or an indication it has been reached.
100%	Expected date when the file system will reach 100% capacity, or an indication it has been reached.

Host Storage Usage Over Time

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

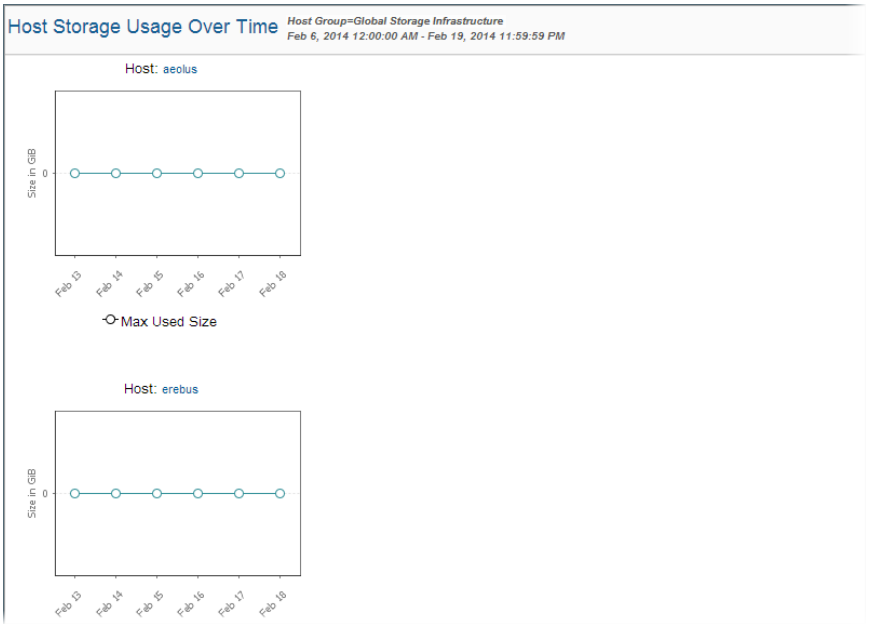
Capacity Manager > Storage Capacity & Forecast > Host Storage Usage Over Time

Line charts, aggregated by Host Groups or Arrays, and grouped by a configurable time period, show allocated capacity over time. You can view a historical perspective on the amount of storage allocated to the ESX Clusters or Server Groups. This helps to illustrate the growth of your environment by ESX clustered or Host Group.

Use the Capacity Forecast report to project growth in storage usage. You can generate the report by the following criteria:

- Historical date range
- Forecast horizon (date range)
- Storage Type (SAN, DAS and NAS)
- Host Type

In addition, you can choose to generate graphs for each host or instead, display a consolidated graph for all the hosts selected in the scope selector.



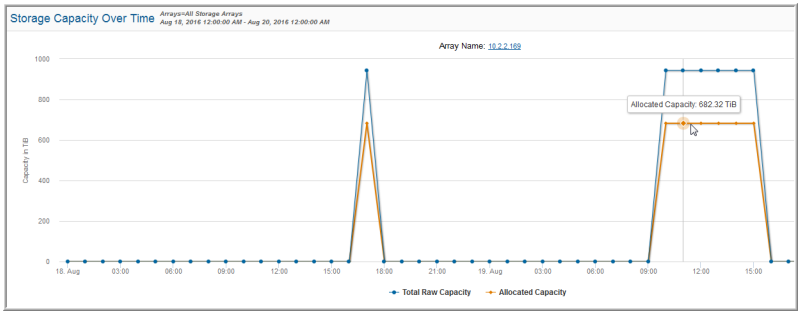
Storage Capacity Over Time

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

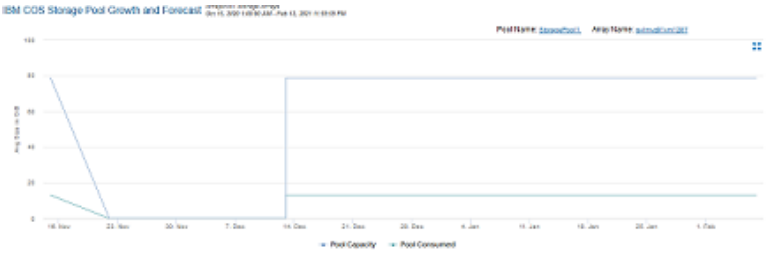
Capacity Manager > Storage Capacity & Forecast > Storage Capacity Over Time

View multi-line charts of capacity metrics for storage arrays, such as Pure Storage FlashArrays. Use this report for an at-a-glance view of raw and allocated storage to determine future capacity requirements.



IBM COS Storage Pool Growth and Forecast

These Line chart reports provide an at-a-glance view of Vault and Storage pool capacity and consumption trends.



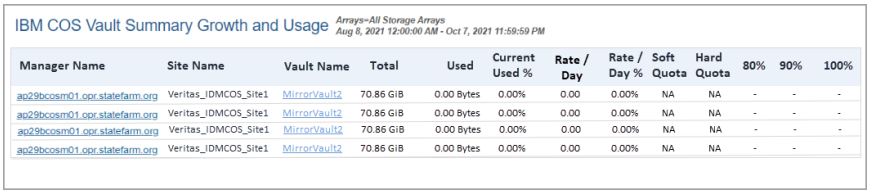
IBM COS Vault Growth and Forecast

These Line chart reports provide an at-a-glance view of Vault and Storage pool capacity and consumption trends.



IBM COS Vault Summary Growth and Usage

This report shows list of Vaults for the COS manager showing the Sites, Capacities, Growth rates, and projected out-of-space conditions.



IBM COS Vault Summary Growth and Usage
Arrays=All Storage Arrays
Aug 8, 2021 12:00:00 AM - Oct 7, 2021 11:59:59 PM

Manager Name	Site Name	Vault Name	Total	Used	Current Used %	Rate / Day	Rate / Day %	Soft Quota	Hard Quota	80%	90%	100%
ap29bcosm01.opr.statefarm.org	Veritas_IDMCOS_Site1	MirrorVault2	70.86 GiB	0.00 Bytes	0.00%	0.00	0.00%	NA	NA	-	-	-
ap29bcosm01.opr.statefarm.org	Veritas_IDMCOS_Site1	MirrorVault2	70.86 GiB	0.00 Bytes	0.00%	0.00	0.00%	NA	NA	-	-	-
ap29bcosm01.opr.statefarm.org	Veritas_IDMCOS_Site1	MirrorVault2	70.86 GiB	0.00 Bytes	0.00%	0.00	0.00%	NA	NA	-	-	-
ap29bcosm01.opr.statefarm.org	Veritas_IDMCOS_Site1	MirrorVault2	70.86 GiB	0.00 Bytes	0.00%	0.00	0.00%	NA	NA	-	-	-

IBM COS Site Summary

This reports shows listing of Site for the COS Manager. We can drill-down to these reports from 'Array Capacity and Utilization Report'.

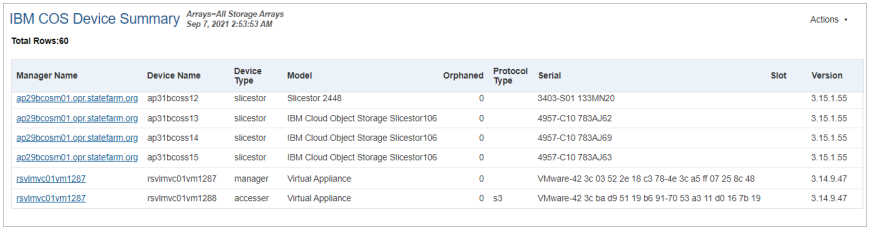


IBM COS Site Summary
Arrays=All Storage Arrays
Sep 7, 2021 2:59:26 AM
Actions

Manager Name	Site Name	Abbreviation	Company	Description	Address	Phone	Latitude	Longitude
ap29bcosm01.opr.statefarm.org	Olathe Data Center	Site31	24400 W. Valley Pkwy	Failover Production Environment	Olathe, KS 66061United States of America	1-877-363-0153, Option 5		
ap29bcosm01.opr.statefarm.org	Site29		State Farm Insurance Company	Primary Production Environment.	Richardson Data Center1402 E. Lookout DriveRichardson, TX 75082	1-877-363-0153, Option 5		
rsvlmvc01vm1287	Veritas_IBMCOS_Site1							

IBM COS Device Summary

This reports shows listing of Device for the COS Manager. We can drill-down to these reports from 'Array Capacity and Utilization Report'.

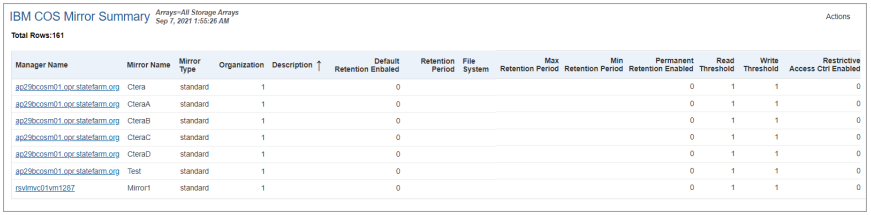


IBM COS Device Summary
Arrays=All Storage Arrays
Sep 7, 2021 2:53:53 AM
Actions
Total Rows:60

Manager Name	Device Name	Device Type	Model	Orphaned	Protocol Type	Serial	Slot	Version
ap29bcosm01.opr.statefarm.org	ap31bcoss12	slicestor	Slicestor 2448	0		3403-S01 133MN20		3.15.1.55
ap29bcosm01.opr.statefarm.org	ap31bcoss13	slicestor	IBM Cloud Object Storage Slicestor106	0		4957-C10 783AJ62		3.15.1.55
ap29bcosm01.opr.statefarm.org	ap31bcoss14	slicestor	IBM Cloud Object Storage Slicestor106	0		4957-C10 783AJ69		3.15.1.55
ap29bcosm01.opr.statefarm.org	ap31bcoss15	slicestor	IBM Cloud Object Storage Slicestor106	0		4957-C10 783AJ63		3.15.1.55
rsvlmvc01vm1287	rsvlmvc01vm1287	manager	Virtual Appliance	0		VMware-42 3c 03 52 2e 18 c3 78-4e 3c a5 ff 07 25 8c 48		3.14.9.47
rsvlmvc01vm1287	rsvlmvc01vm1288	accesser	Virtual Appliance	0	s3	VMware-42 3c ba d9 51 19 b5 91-70 53 a3 11 d0 16 7b 19		3.14.9.47

IBM COS Mirror Summary

This reports shows listing of Mirror for the COS Manager. We can drill-down to these reports from 'Array Capacity and Utilization Report'.



IBM COS Mirror Summary *Arrays=All Storage Arrays*
Sep 7, 2021 1:55:26 AM

Actions

Total Rows:161

Manager Name	Mirror Name	Mirror Type	Organization	Description ↑	Default Retention Enbaled	Retention Period	File System	Max Retention Period	Min Retention Period	Permanent Retention Enabled	Read Threshold	Write Threshold	Restrictive Access Ctrl Enabled
ap29bcosm01.opr.statefarm.org	Ctera	standard	1		0					0	1	1	0
ap29bcosm01.opr.statefarm.org	CteraA	standard	1		0					0	1	1	0
ap29bcosm01.opr.statefarm.org	CteraB	standard	1		0					0	1	1	0
ap29bcosm01.opr.statefarm.org	CteraC	standard	1		0					0	1	1	0
ap29bcosm01.opr.statefarm.org	CteraD	standard	1		0					0	1	1	0
ap29bcosm01.opr.statefarm.org	Test	standard	1		0					0	1	1	0
rsvlmvc01vm1287	Mirror1	standard	1		0					0	1	1	0

Chapter 17

Storage Performance Reports

This chapter includes the following topics:

- Performance Profiles Overview
- EMC Isilon Metrics
- NetApp Cluster-Mode Metrics
- EMC Symmetrix Enhanced Performance Metrics
- Determining Performance States
- Populating Performance States
- Updating Performance Profiles
- Aggregated Profiles
- Array Performance by RAID Group
- Array Performance by LUN
- LUN Details
- Data Domain Disk Performance
- Data Domain System Performance
- EMC Isilon CPU Performance by Cluster
- EMC Isilon CPU Performance by Node
- EMC Isilon Disk Performance by Cluster

- EMC Isilon Disk Performance by Node
- EMC Isilon File System Performance by Protocol
- EMC Symmetrix Array Performance
- EMC Symmetrix Device Group Performance
- EMC Symmetrix Disk Performance
- EMC Symmetrix FE Director Performance
- EMC Symmetrix Storage Group Performance
- EMC Symmetrix Thin Tier Performance
- EMC VNX (Celerra) Data Mover Performance
- Host Load Average
- Host CPU Utilization
- Host Memory Utilization
- Host Memory Swap
- Host Disk reads/writes
- Host Network Packets sent/received
- LUN Performance
- Mission Control for Performance Analysis
- NetApp Cluster-Mode Aggregate Performance by Cluster
- NetApp Cluster-Mode Aggregate Performance by Node
- NetApp Cluster-Mode CIFS Performance by Cluster
- NetApp Cluster-Mode CIFS Performance by Vserver
- NetApp Cluster-Mode Disk Performance by Cluster
- NetApp Cluster-Mode Disk Performance by Node
- NetApp Cluster-Mode NFS Performance by Cluster
- NetApp Cluster-Mode NFS Performance by Vserver
- NetApp Cluster-Mode Volume Performance by Cluster
- NetApp Cluster-Mode Volume Performance by Vserver

- NetApp Cluster-Mode Performance by Cluster
- NetApp Cluster-Mode Performance by Node
- Performance Profiles
- Port Throughput by Array

Performance Profiles Overview

Data Collection of array performance data is supported for certain vendor storage arrays. See the *Certified Configurations Guide* for a complete list of supported arrays.

Performance profiles reflect key performance metrics for groups of LUNs. These profiles are used to assign performance states to LUNs and hosts, as shown in several NetBackup IT Analytics array performance reports. NetBackup IT Analytics maintains aggregated performance profiles for each collection of RAID groups and thin pools that share common physical characteristics. These performance profiles categorize a host's performance state, based on LUN performance: Slow, Normal, or Fast. All LUNs are mapped to either a RAID group or thin pool thereby enabling array performance reports to list the associated drives, drive types, and drive speed. Refer to the following reports:

See " Array Performance by RAID Group " on page 537.

See " Array Performance by LUN " on page 540.

Performance profiles are derived from previously captured array raw performance statistics. Internal procedures examine captured performance metrics for the specified interval and populate profiles with actual ranges of performance and throughput metrics for slow, normal, and fast performance states.

Performance profiles should be based on a typical peak array performance interval. By default, the updating of performance profiles occurs daily and uses an interval of 10:00 a.m. to 11:00 a.m. You should change the default interval to correspond to the peak array performance interval for your environment. When the Data Collector policy is configured, a recurring frequency is selected to set the length of the profile interval. Baseline profiling determines the range of metrics for I/Os, response time, and cache hits.

See "Determining Performance States" on page 535.

See "Populating Performance States" on page 536.

An on-going scheduled daily task profiles collected data. Since certain times of day are more prone to performance issues, the performance profile schedule can be modified for your environment. For Performance Profile Schedule Customization,

see Customize the Array Performance Profile Schedule in the System Administrator's Guide.

Both internal and community profiling can be enabled.

- Internal profiling compares periodic collection cycles within just your environment. Essentially, a library of metrics is maintained for comparisons.
- Community profiling allows contributors securely import/export performance profiles over https, enabling comparisons with a broader community of similar arrays and environments.
  See "Aggregated Profiles" on page 536.

EMC Isilon Metrics

Isilon performance data (raw, hourly, and daily) is collected from SNMP MIB statistics. For example, collected data includes such metrics as cluster, node, protocols (CIFS, SMB, FTP, HTTP), and disk performance. See the Array Performance Statistics Technical Note for an extensive list of collected metrics.

NetApp Cluster-Mode Metrics

A large variety of NetApp Cluster-Mode performance data (raw, hourly, and daily) is collected. For example, collected data includes such metrics as system, protocols (CIFS and NFS), volume, LUN, and target port performance. See the Array Performance Statistics Technical Note for an extensive list of collected metrics.

EMC Symmetrix Enhanced Performance Metrics

In addition to LUN and Port performance metrics that can be collected from EMC Symmetrix arrays, data collection gathers other performance metrics by accessing storage devices via the EMC Unisphere REST API. See the Array Performance Statistics Technical Note for an extensive list of collected metrics.

Determining Performance States

After performance profiles have been created, future reports will use the performance profiles for determining the performance state of LUNs, hosts, and other objects (in future releases).

Populating Performance States

Performance profiling establishes criteria for deriving a LUN's performance state, as part of future NetBackup IT Analytics reporting. The criteria is established based on actual performance metrics captured from each LUN. Reads and writes are treated independently.

Updating Performance Profiles

Performance profiles are used to characterize response times, I/Os per second, MBs per second, and cache hits per second as Poor, Normal, and Fast. Creation of and updating Performance Profiles establishes the criteria used based on performance and throughput metrics captured within a profile interval, using the following logic:

1. For all LUNs with performance data captured during the specified performance interval, group the data based on common physical attributes: array vendor, array family, array model, RAID type, number of drives in RAID Group, drive type, drive form, and drive speed. Each unique grouping of LUNs based on this criteria represents an individual NetBackup IT Analytics Performance reporting profile entry for this Portal.
2. Calculate the average and standard deviation for all captured metrics based on all LUNs with performance data in the specified interval and within a common performance profile. When there is more than one poll to the array during the profile updating interval, averages and standard deviation will be based on only the portion of the interval covered for the polls that return data, ignoring polls that did not capture I/Os or where the response times were zero.
3. To determine boundaries within the range, subtract one standard deviation from the average. If the standard deviation is larger than the average, divide the average by two. This establishes the low end of the Normal boundary and the high end of the Fast boundary.
4. Add one standard deviation to the average. If the standard deviation is larger than the average, add half of the average to the average. This establishes the high end of the Normal boundary and the low end of the Poor boundary.

Aggregated Profiles

An aggregated array performance profile is maintained at Profile Central. This profile takes into account unique configurations: array vendor, array model. RAID configuration, number of data and parity drives, drive size, drive type, and drive speed. In addition, the total MB transferred during a profile interval weights individual

profiles with aggregated equivalents in the community profile. As more published profiles are aggregated, the MB transferred will have less impact on the weighting factor.

Array Performance by RAID Group

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > Array Performance by RAID Group

For a list of Array Families, collected performance statistics, and details on how performance profiles are determined, see *Certified Configurations Guide*.

- For CLARiiON arrays, a LUN that is assigned to a Pool (as is required for dynamic provisioning) will not have any performance statistics collected.
- HP 3PAR arrays are represented uniquely with each Common Provisioning Group (CPG) represented as a RAID Group. Virtual volumes (VV) can be allocated directly, outside of a CPG with each VV being represented as an individual RAID Group.
- Array LUN performance response times (total, read, write) are not provided in reports for the following array families: EMC Symmetrix, Dell Compellent, and HP 3PAR.

Note: Column aggregations are displayed in green on a separate line with the label Aggregation. Rollover the value to display the aggregation type: Avg., Min., Max., or Sum. Column filtering is not supported on reports with aggregated columns.

Table 17-1 Array Performance by RAID Group

Column name	Description
Storage Array	Array name.
RAID Group	RAID group name.

Table 17-1 Array Performance by RAID Group *(continued)*

Column name	Description
Performance	Either Fast, Normal, or Slow, based on the performance index calculation for the LUN using the last baseline profile ranges for Fast and Slow performance; takes into account the worst performance of a LUN in the RAID Group. Each of these performance states are represented by a colored dot as shown in the legend at the top of the report.
RAID Type	RAID type.
Drive Type	Drive type, such as SAS or SATA.
Drive Speed	Drive speed in RPM.
Drive Form	Drive enclosure.
# LUNs	# of LUNs with drilldown to Array Performance by LUN. See " Array Performance by LUN " on page 540. You can also click the LUN name to show the details. See " LUN Details " on page 542.
Response (ms)	I/O response time = the average response time for all read and write operations (average I/Os). Response-time metrics are not available for Dell Compellent, EMC Symmetrix, IBM SVC, or HP 3PAR arrays.
Max Read	Identifies the LUN in the group with the longest read response time. This is the highest response time for a single LUN, from all the LUNs in the group. This can be used to identify the LUN with a performance problem.
Max Write	Identifies the LUN in the group with the longest write response time. This is the highest response time for a single LUN, from all the LUNs in the group. This can be used to identify the LUN with a performance problem.
Avg Total	Average of the average response times.
Avg Read	Average of the average response times (for the LUNs in the RAID Group) during the specified interval in milliseconds. Response-time metrics are not available for Dell Compellent, EMC Symmetrix, IBM SVC, or HP 3PAR arrays.

Table 17-1 Array Performance by RAID Group *(continued)*

Column name	Description
Avg Write	Average of the average response times (for the LUNs in the RAID Group) during the specified interval in milliseconds. Response-time metrics are not available for Dell Compellent, EMC Symmetrix, IBM SVC, or HP 3PAR arrays.
Throughput IO/sec	Average Read IOs/sec + Average Write IOs/sec.
Max Read	Identifies the LUN in the group with the largest read throughput (IOs/sec). This is the largest value for a single LUN, from all the LUNs in the group. This can be used to identify the LUN with a performance problem.
Max Write	Identifies the LUN in the group with the largest write throughput (IOs/sec). This is the largest value for a single LUN, from all the LUNs in the group. This can be used to identify the LUN with a performance problem.
Avg Total	Sum of the throughput (IOs/sec) averages.
Avg Read	Sum of the average Read IOs for the interval/length of interval in seconds.
Avg Write	Sum of the average Write IOs for the interval/length of interval in seconds.
Max Read	Identifies the LUN in the group with the highest read transfer rate. This is the highest read transfer rate for a single LUN, from all the LUNs in the group. This can be used to identify the LUN with a performance problem.
Max Write	Identifies the LUN in the group with the highest write transfer rate. This is the highest write transfer rate for a single LUN, from all the LUNs in the group. This can be used to identify the LUN with a performance problem.
Avg Total	Average of the average total transfer rates for the LUNs in the group.
Avg Read	Average of the average read transfer rate for the LUNs in the group, for the interval/length of interval in ms.
Avg Write	Average of the average write transfer rate for the LUNs in the group, for the interval/length of interval in ms.
Cache Hits/sec Read	Total Read Hits for interval/length of interval in seconds.

Table 17-1 Array Performance by RAID Group *(continued)*

Column name	Description
Cache Hits/sec Write	Total Write Hits for interval/length of interval in seconds.

Array Performance by LUN

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Storage Performance > Array Performance by RAID Group

- Click **# of LUNs**.
- Select a row and click the **Show Chart** button to display the LUN Performance report. This shows performance charts as a time series. Roll-over the chart to show date and time. Select multiple rows to compare performance charts of individual LUNs.
- Click **LUN** to display the LUN Details report to show performance graphs and LUN attributes.
  See " LUN Details " on page 542.

Table 17-2 Array Performance by LUN

Column name	Description
Storage Array	Array name
RAID Group	RAID group name

Table 17-2 Array Performance by LUN *(continued)*

Column name	Description
Performance	Either Fast, Normal, or Slow, based on the performance index calculation for the LUN using the last baseline profile ranges for Fast and Slow performance; takes into account the worst performance of a LUN in the RAID Group. Each of these performance states are represented by an icon as shown in the legend at the top of the report.
LUN	LUN name. Click the **LUN** name to view the details. See " LUN Details " on page 542.
RAID Type	RAID type
Drive Type	Drive type, such as SAS or SATA
Drive Speed	Drive speed in RPM
Hosts	The name of the host to which storage from the LUN was allocated.
Max Read	Maximum read response time.
Max Write	Maximum write response time.
Avg Total	Average total response time.
Avg Read	Average response time (for the LUNs in the RAID Group) during the specified interval in milliseconds. Response-time metrics are not available for Dell Compellent, EMC Symmetrix, IBM SVC, or HP 3PAR arrays.
Avg Write	Average response time (for the LUNs in the RAID Group) during the specified interval in milliseconds. Response-time metrics are not available for Dell Compellent, EMC Symmetrix, IBM SVC, or HP 3PAR arrays.
Throughput IO/sec	Average Read IOs/sec + Average Write IOs/sec = Throughput Average.
Max Read	Maximum read throughput (IOs/sec)
Max Write	Maximum write throughput (IOs/sec)
Avg Total	Average total throughput (IOs/sec)
Avg Read	Average Read IOs for interval/length of interval in seconds
Avg Write	Average Write IOs for interval/length of interval in seconds
Max Read	Maximum read transfer rate.

Table 17-2 Array Performance by LUN *(continued)*

Column name	Description
Max Write	Maximum write transfer rate.
Avg Total	Average total transfer rate.
Avg Read	Average megabytes transferred for the interval/length of interval in milliseconds
Avg Write	Average megabytes transferred for the interval/length of interval in milliseconds
Cache Hits/sec Read	Total Read Hits for interval/length of interval in seconds
Cache Hits/sec Write	Total Write Hits for interval/length of interval in seconds
Show Chart	After selecting a row, click the **Show Chart** button to display a time series chart. Roll-over the chart to show date and time. You can select multiple rows to show the LUN Performance of multiple LUNs and compare them side-by-side.

LUN Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

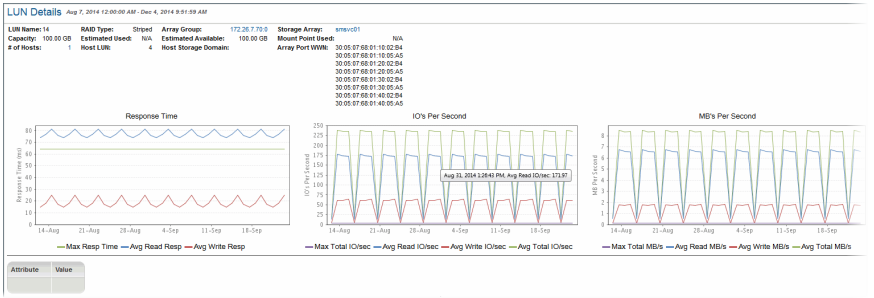
Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Storage Performance > Array Performance by RAID Group

1. Click **# of LUNs**.
2. Click **LUN** to display the **LUN Details** report to show performance graphs and LUN attributes for an individual LUN.

The **LUN Details** report is divided into three sections:

- LUN Summary
- Performance Graphs
- LUN Attributes



Data Domain Disk Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > Data Domain Disk Performance

This report display the average throughput for all disks in the Data Domain system along with the average disk busy statistics. These statistics are averaged over time, since the last “disk re-set performance” command or since the last system power-up cycle. The Data Domain system links to details.

See “ Data Domain System Details” on page 246.

Data Domain System Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > Data Domain System Performance

Graphs in this report provide a view of performance for the following metrics:

Table 17-3 Data Domain System Performance

Column name	Description
Peak CPU Busy %	The percentage of time that the Data Domain system's CPUs are busy
Peak FS Proc %	The number of NFS operations per second
Peak Disk Busy %	The percentage of time that all disks in the Data Domain system are busy

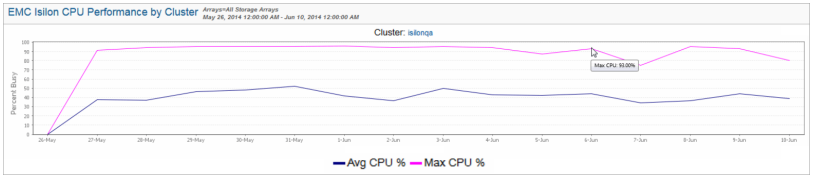
EMC Isilon CPU Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Isilon CPU Performance by Cluster

Use this report to determine a cluster's CPU utilization over a historical timeline. For the report scope, the average and maximum CPU percentages (%) are shown for each cluster.



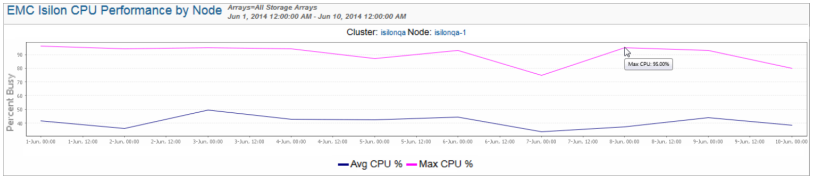
EMC Isilon CPU Performance by Node

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Isilon CPU Performance by Node

This report displays a node's CPU utilization over a historical timeline. For the report scope, the average and maximum CPU percentages (%) are shown for each node.



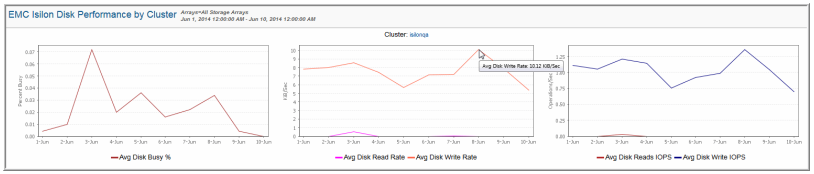
EMC Isilon Disk Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Isilon Disk Performance by Cluster

Displays disk performance by cluster metrics over a historical timeline. Line charts represent disk performance: average response times, busy %, read/write rates, and IOPS.



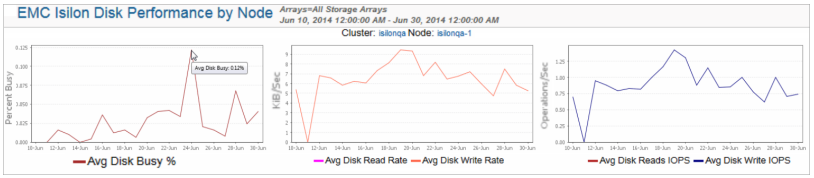
EMC Isilon Disk Performance by Node

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Isilon Disk Performance by Node

Use this report to display disk performance by node metrics over a historical timeline. Line charts represent disk performance: average response times, busy %, read/write rates, and IOPS.



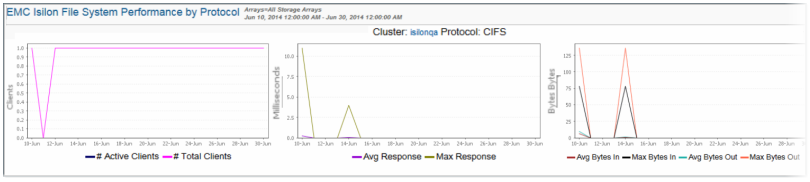
EMC Isilon File System Performance by Protocol

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Isilon File System Performance by Protocol

This report illustrates specific performance metrics for Isilon file systems. Displays line charts representing active clients, along with average and maximum transfer rates and response times.



EMC Symmetrix Array Performance

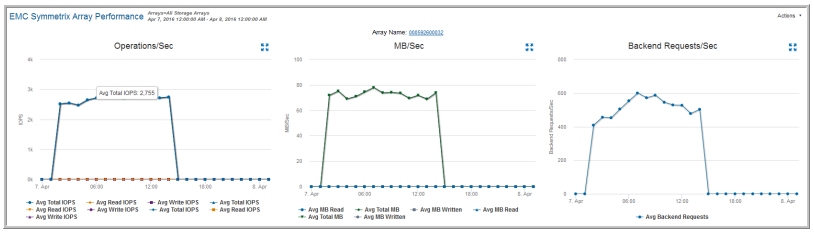
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Symmetrix Array Performance

This report illustrates specific performance metrics for EMC Symmetrix arrays. It displays line charts representing operations per second, data read and written, and backend requests. Mouse-over the data points to view specific values. View line charts of Symmetrix array performance metrics over time.

Metrics that can be selected for the charts include Total IOPS (Input/Output Operations Per Second), Read and Write IOPS, backend requests (requests from the director to the attached backend devices), and data transfer rates: Read, Write, and Total MBps. Values represented in this report are averages for the time period.



EMC Symmetrix Device Group Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Symmetrix Device Group Performance

View a table of Symmetrix device group performance metrics.

Metrics that can be selected for the charts include % Busy, Host IO per second, Host MBps, Requests per second, and Read and Write response times. Values represented in this report are averages for the time period.

Table 17-4 EMC Symmetrix Device Group Performance

Column name	Description
Array Name	Name of the storage array. This name is derived from the Symmetrix ID.
Avg Read Latency	The average time that it took the Symmetrix to serve one read IO for this group.
Max Read Latency	The MAX of the average time that it took the Symmetrix to serve one read IO for this group.
Read Latency Std Dev	The Std Dev of average time that it took the Symmetrix to serve one read IO for this group.
Avg Write Latency	The average time that it took the Symmetrix to serve one write IO for this group.
Max Write Latency	The MAX of the average time that it took the Symmetrix to serve one write IO for this group.
Write Latency Std Dev	The Standard Deviation of average time that it took the Symmetrix to serve one write IO for this group.
Avg Read Hits	The number of host read operations performed each second by the group that were immediately satisfied from cache.
Max Read Hits	The MAX of the number of host read operations performed each second by the group that were immediately satisfied from cache.
Avg Write Hits	The number of host write operations performed each second by the group that were immediately satisfied from cache.
Max Write Hits	The MAX of the number of host write operations performed each second by the group that were immediately satisfied from cache.
Avg Write MB/Sec	The cumulative number of host MBs written per second by the group.
Max Write MB/Sec	The MAX of the cumulative number of host MBs written per second by the group.
Avg Read MB/Sec	The cumulative number of host MBs read per second by the group.
Max Read MB/Sec	The cumulative number of host MBs read per second by the group.
Avg Write IOPS	The number of host write operations performed each second by the group.
Max Write IOPS	The MAX of the number of host write operations performed each second by the group.

Table 17-4 EMC Symmetrix Device Group Performance *(continued)*

Column name	Description
Avg Read IOPS	The number of host read operations performed each second by the group that were immediately satisfied from cache.
Max Read IOPS	The MAX of the number of host read operations performed each second by the group.
Avg Total IOPS	Average total input/output operations per second, including both read and write operations.

EMC Symmetrix Disk Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Symmetrix Disk Performance

View a table of Symmetrix disk performance metrics.

Metrics listed in the report, by array and device, include capacity, Average Busy %, Average Queue Depth, Response Times (Average Read and Write), Input/Output Operations Per Second - IOPS (Average Read and Write), and data transfer rates (Average Read and Write MBps). Values represented in this report are averages for the time period.

Table 17-5 EMC Symmetrix Disk Performance

Column name	Description
Array Name	Name of the storage array. This name is derived from the Symmetrix ID.
Vendor	Vendor/Manufacturer of the disk.
Serial #	Serial number of the disk. Example: 3LN3216B.
Avg Busy%	The percent of time that the disk is busy serving IOs.
Max Busy%	The MAX of the percent of time that the disk is busy serving IOs.

Table 17-5 EMC Symmetrix Disk Performance *(continued)*

Column name	Description
Avg Queue Depth	Calculated value: Accumulated queue depth/total SCSI command per sec.
Max Queue Depth	Calculated value: Max of the Accumulated queue depth/total SCSI command per sec.
Avg Read Latency	The average time it took the disk to serve one read command.
Max Read Latency	Max of the average time it took the disk to serve one read command.
Read Latency Std Dev	The Std Dev of the average time it took the disk to serve one read command.
Avg Write Latency	The average time it took the disk to serve one write command.
Max Write Latency	The Max of the average time it took the disk to serve one write command.
Write Latency Std Dev	The Std Dev of the average time it took the disk to serve one write command.
Avg IOPS	The number of host read and write requests for the disk.
Max IOPS	The Max of the number of host read and write requests for the disk.
Avg Read IOPS	The number of host reads per second for the disk.
Max Read IOPS	The Max of the number of host reads per second for the disk.
Avg Write IOPS	The number of host writes per second for the disk.
Max Write IOPS	The Max of the number of host writes per second for the disk.
Avg MB Read	The read throughput (MBs) of the disk per second.
Max MB Read	The Max of the read throughput (MBs) of the disk per second.
Avg MB Written	The write throughput (MBs) of the disk per second.
Max MB Written	The max of the write throughput (MBs) of the disk per second.

EMC Symmetrix FE Director Performance

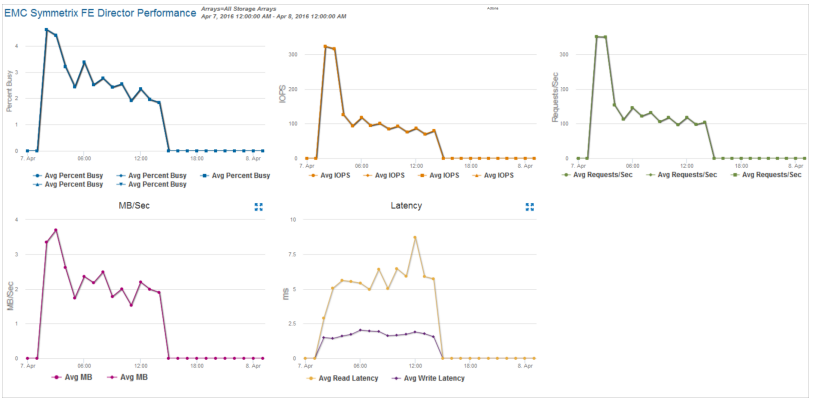
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Symmetrix FE Director Performance

View line charts of performance metrics for select Symmetrix front-end directors.

Metrics that can be selected for the charts include % Busy, Host IO per second, Host MBps, Requests per second, and Read and Write response times. Values represented in this report are averages for the time period.



EMC Symmetrix Storage Group Performance

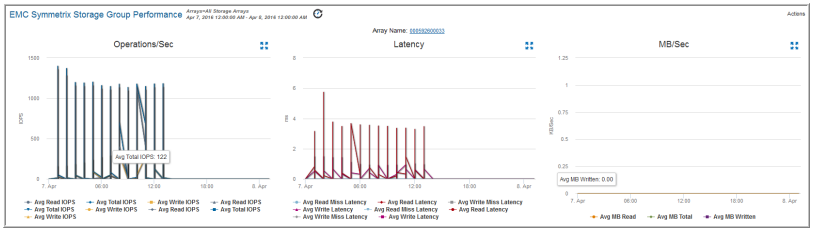
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Symmetrix Storage Group Performance

View line charts of performance metrics for select Symmetrix storage groups.

Metrics that can be selected for the charts include Total IOPS (Input/Output Operations Per Second), Read and Write IOPS, data transfer rates (Read, Write, and Total MBps), Read and Write response times, and Read and Write Missed (not found in buffer) response times.



EMC Symmetrix Thin Tier Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC Symmetrix Thin Tier Performance

View a table of Symmetrix thin pool performance metrics.

Metrics listed in the report, by array, pool, and tier, include capacity, Response Times (Average Read and Write), Average Input/Output Operations Per Second - IOPS, and Average MBps. Values represented in this report are averages for the time period.

Table 17-6 EMC Symmetrix Thin Tier Performance

Column name	Description
Array Name	Name of the storage array. This name is derived from the Symmetrix ID.
Tier Name	The thin tier ID, as returned by Unisphere.
Avg Capacity(MB)	Average Capacity in MB.
Max Capacity(MB)	Max of the capacity in MB.
Avg Used Capacity(MB)	Average Used Capacity in MB.
Max Used Capacity(MB)	Max of the Used Capacity in MB.

Table 17-6 EMC Symmetrix Thin Tier Performance *(continued)*

Column name	Description
Avg Read Latency	The average time that it took the Symmetrix to serve one read IO for this group.
Max Read Latency	The Max of the average time that it took the Symmetrix to serve one read IO for this group.
Read Latency Std Dev	Std Dev of the average time that it took the Symmetrix to serve one read IO for this group.
Avg Write Latency	The average time that it took the Symmetrix to serve one write IO for this group.
Max Write Latency	Max of the average time that it took the Symmetrix to serve one write IO for this group.
Write Latency Std Dev	Std Dev of the average time that it took the Symmetrix to serve one write IO for this group.
Avg MB/Sec	Cumulative number of host MBs read/writes per second by the group.
Max MB/Sec	Cumulative number of Max of the host MBs read/writes per second by the group.
Avg Read IOPS	The average time it took the disk to serve one read command.
Max Read IOPS	Max of the average time it took the disk to serve one read command.
Avg Write IOPS	The number of host writes per second for the disk.
Max Write IOPS	Max of the number of host writes per second for the disk.

EMC VNX (Celerra) Data Mover Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

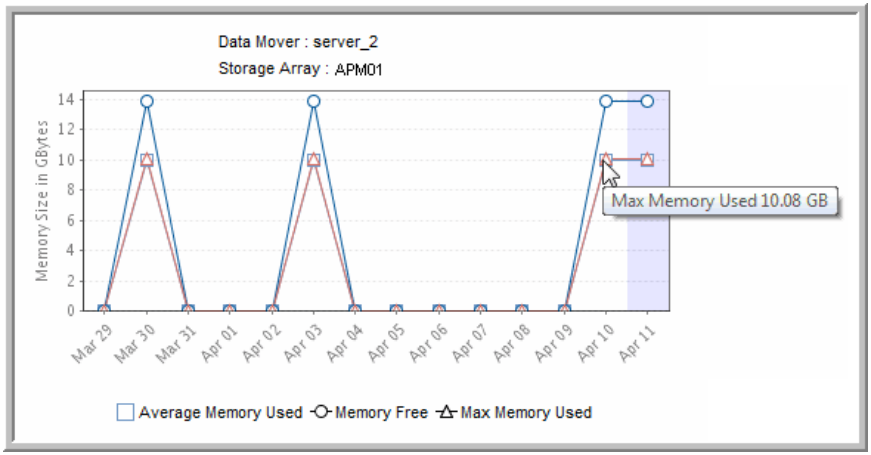
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > EMC VNX (Celerra) Data Mover Performance

Graphs for the selected report time period represent the following collected statistics:

- CPU Utilization: Average and Maximum
- Memory Use: Average, Free, and Maximum

To view the actual values plotted on the graphs, mouse over the symbols.



Host Load Average

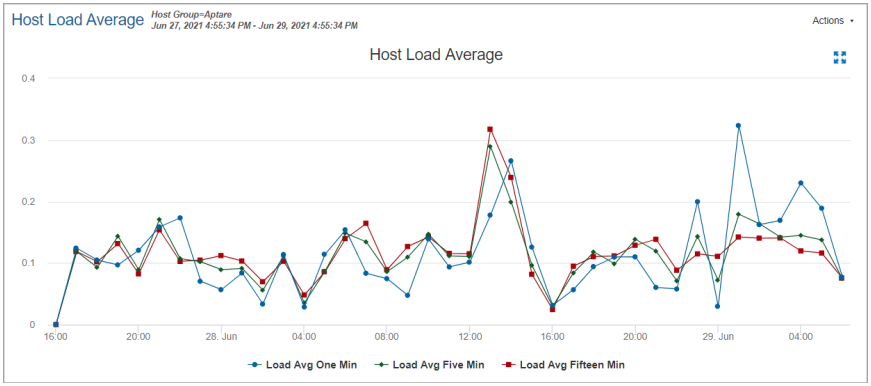
This report displays the chart based on values derived from hosts running only on UNIX operating systems.

This chart displays historical load averages derived from the uptime command. The uptime command displays the current time, how long the system is running, the number of active logged in users, and the host load averages in the past 1, 5, and 15 minutes.

Host load average is the average of the processes in either pre-run or uninterruptible state. A process in pre-run state either already using the CPU or waiting to use it. However, a process in uninterruptible state is awaiting the use of CPU, which means it is waiting for an I/O access, such as waiting for disk. The averages displayed in the report are taken over the three time intervals mentioned before, but are not normalized for the number of CPUs in a system.

This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:**Capacity Manager** > **Storage Performance** > **Host Load Average**.

See "Backup Server Performance Dashboard" on page 728.



Host CPU Utilization

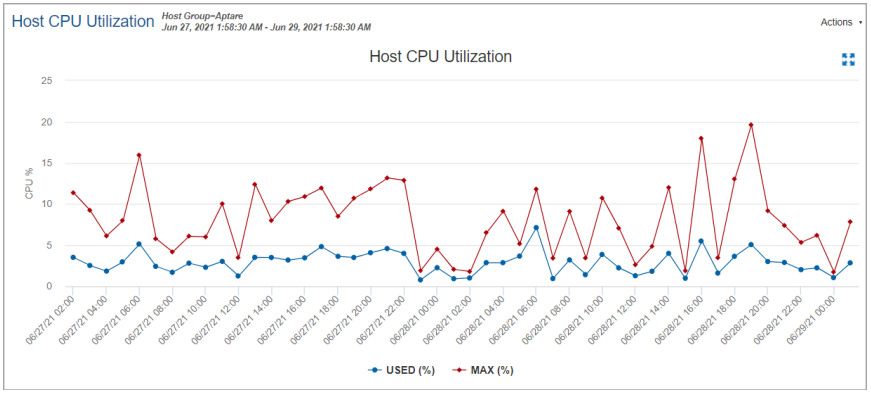
The report displays a chart of the average CPU usage of the backup host over the time duration specified for the report. In addition, it also plots an additional line indicating the maximum CPU usage for the specified time interval.

Even though the report is designed to evaluate the CPU utilization of a single host, you may view a chart for multiple hosts. In that case, the chart will display values based on the average utilization of all the hosts.

This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:**Capacity Manager** > **Storage Performance** > **Host CPU Utilization**.

This report also populates the data collected by Compute Resource and Flex-Appliances.

See “Backup Server Performance Dashboard” on page 728.



Host Memory Utilization

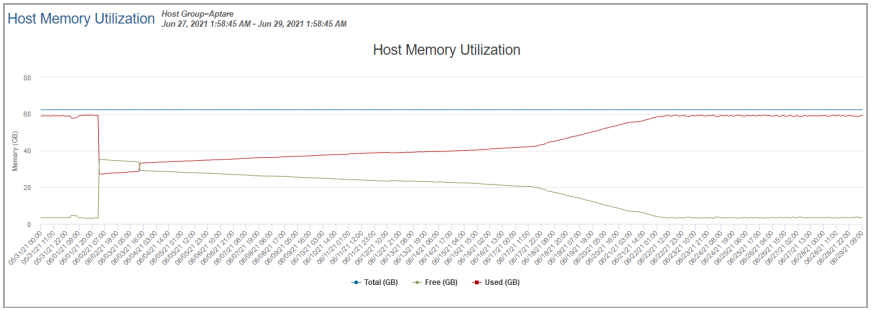
The report displays a chart of the average memory consumption of the backup host over the time duration specified for the report. In addition, it also plots an additional line indicating the total memory usage for the specified time interval. Although the report is designed to evaluate the memory consumption of a single host, you may view a chart for multiple hosts. In that case, the chart will display values based on the average memory consumption of all the hosts.

This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:**Capacity Manager** > **Storage Performance** > **Host Memory Utilization**.

This report also populates the data collected by Compute Resource and Flex-Appliances.

Note: This report does not populate data for Windows exporter.

See “Backup Server Performance Dashboard” on page 728.

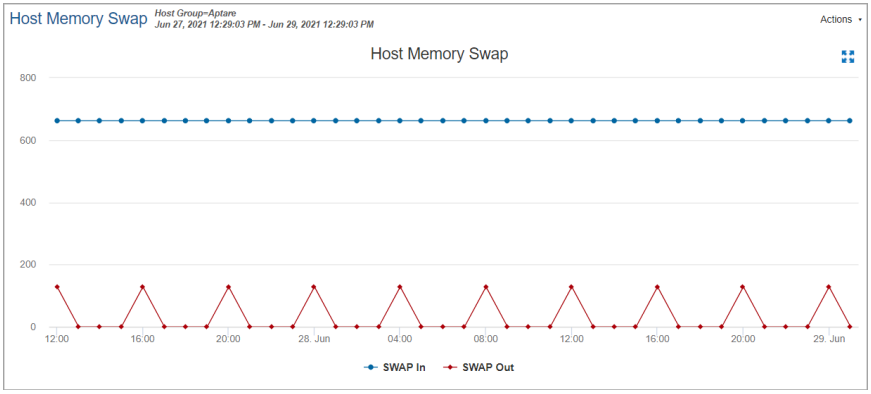


Host Memory Swap

The report displays the swap in and swap out of memory per minute for each host within the duration specified while generating the report. This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:**Capacity Manager** > **Storage Performance** > **Host Memory Swap**.

See “Backup Server Performance Dashboard” on page 728.



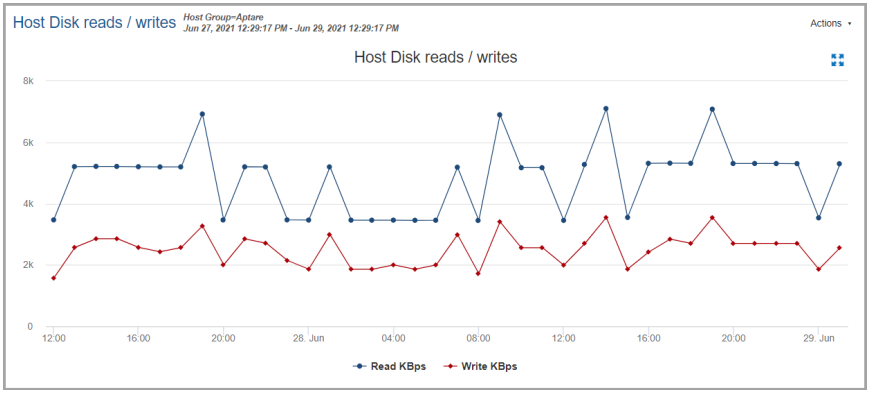
Host Disk reads/writes

This report summarizes the reads and writes of all the disks connected to the host. The values are calculated in Kilobits/second and displayed on the chart for the time duration specified while generating the report. You can evaluate the disk utilization using this report and resolve slowness of applications running on the host.

This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:**Capacity Manager** > **Storage Performance** > **Host Disk reads/writes**.

This report also populates the data collected by Compute Resource and Flex-Appliances.

See "Backup Server Performance Dashboard" on page 728.



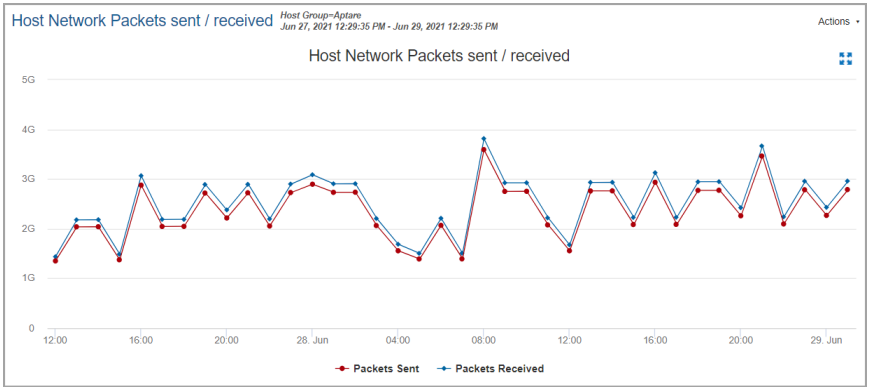
Host Network Packets sent/received

This report displays a chart of summarized values of the packets sent and received over the host network. The chart is a plot of summarized values obtained from all the NICs in the host network. This chart can help you minimize performance issues arising from network congestions.

This report also populates the data collected by Compute Resource and Flex-Appliances.

This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances. This report is located here:**Capacity Manager** > **Storage Performance** > **Host network Packets sent/received**.

See "Backup Server Performance Dashboard" on page 728.



LUN Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

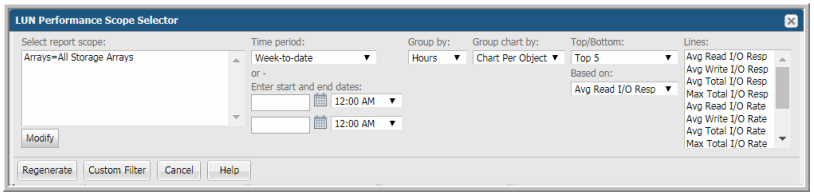
Capacity Manager > Storage Performance > LUN Performance

Graphs for the selected report time period represent the following LUN performance statistics:

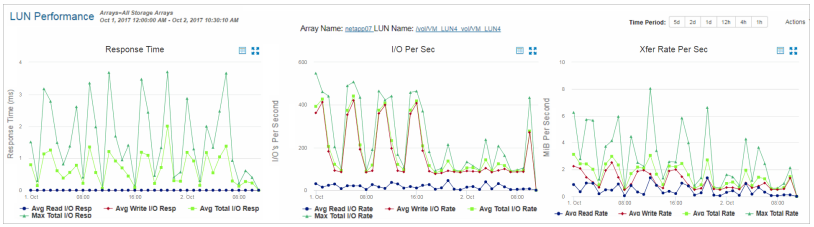
- Response Time: Average Read, Average Write, Average Total I/O, and Maximum Total I/O
- I/O's Per Second
- Read/Write Rates

Note: Array LUN performance response times (total, read, write) are not provided in reports for the following array families: EMC Symmetrix, Dell Compellent, and HP 3PAR.

In the LUN Performance Scope Selector, you can choose a number of filters to create the LUN Performance report:



- Click the **Array Name** link to view the Array Capacity & Utilization report.
- Click the **LUN Name** link to view the LUN Utilization Summary.
- Click the **Time Period** quick filters to change the time period represented in the chart. Choose from the last **5 days, 2 days, 1 day, 12 hours, 4 hours or 1 hour**.



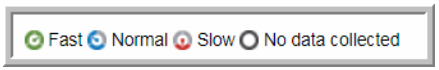
Mission Control for Performance Analysis

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > Mission Control - Performance

The following array performance indicators are displayed for each day in the report's time period:



- Click any **host name** to view the Host Utilization Details.
- Click the arrow in front of a host to expand the list of LUNs and their associated performance statistics for all intervals in the report's scope.

Vendor-Specific Considerations for Array Performance Metrics

- For CLARiiON arrays, a LUN that is assigned to a Pool (as is required for dynamic provisioning) will not have any performance statistics collected.
- LUN Performance (response time) metrics are not available in this report for the following vendor arrays: EMC Symmetrix, Dell Compellent, IBM SVC, and HP 3PAR.
- Array LUN performance response times (total, read, write) are not provided in reports for the following array families: EMC Symmetrix, Dell Compellent, and HP 3PAR.

NetApp Cluster-Mode Aggregate Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Aggregate Performance by Cluster

Aggregate performance rates by cluster are represented by lines: transfers, user reads, and user writes. Select the metrics to be represented as lines. Access the related cluster's details via a link to the Array Capacity & Utilization report.

Click the **Cluster Name** link to view the Array Capacity & Utilization report.

Table 17-7 NetApp Cluster-Mode Aggregate Performance by Cluster

Column name	Description
Cluster Name	Name of the storage array.
Total Transfers	Total number of transfers per second serviced by the aggregate.
# User Reads	Number of user reads per second to the aggregate.
# User Writes	Number of user writes per second to the aggregate.
Log Date	Timestamp at which the performance data was collected.

NetApp Cluster-Mode Aggregate Performance by Node

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Aggregate Performance by Node

Aggregate performance rates by node are represented by lines: transfers, user reads, and user writes. Select the metrics to be represented as lines.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Node** link to view the NetApp Cluster-Mode Node Summary report.

Table 17-8 NetApp Cluster-Mode Aggregate Performance by Node

Column name	Description
Cluster Name	Name of the storage array.
Node Name	The textual name of a node.
Total Transfers	Total number of transfers per second serviced by the aggregate.
# User Reads	Number of user reads per second to the aggregate.
# User Writes	Number of user writes per second to the aggregate.
Log Date	Timestamp at which the performance data was collected.

NetApp Cluster-Mode CIFS Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode CIFS Performance by Cluster

Three line charts represent CIFS performance by cluster metrics: CIFS Operations per Second (IOPS), CIFS latency, and client performance (number of connected shares and open files). Select the IOPS metrics that you want represented by lines.

Click the **Cluster Name** link to view the Array Capacity & Utilization report.

Table 17-9 NetApp Cluster-Mode CIFS Performance by Cluster

Column name	Description
CIFS IOPS	Total number of CIFS operations.
Cluster Name	Name of the storage array.
Log Date	Timestamp at which the performance data was collected.
Avg CIFS Latency	Average latency for CIFS operations.
# CIFS Connected Shares	Number of SMB and SMB2 share connections.
# CIFS Connections	Number of connections.
# Open Files	Number of open files over SMB and SMB2.

NetApp Cluster-Mode CIFS Performance by Vserver

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode CIFS Performance by Vserver

Three line charts represent CIFS performance by Vserver metrics: CIFS Operations per Second (IOPS), CIFS latency, and client performance (number of connected shares and open files). Select the IOPS metrics that you want represented by lines.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Vserver** link to view the NetApp Cluster-Mode Vserver Summary report.

Table 17-10 NetApp Cluster-Mode CIFS Performance by Vserver

Column name	Description
CIFS IOPS	Total number of CIFS operations.
Cluster Name	Name of the storage array.
Vserver Name	Name of the Vserver. When specified as part of a Vserver-create, this is the name of the Vserver.
Log Date	Timestamp at which the performance data was collected.
Avg CIFS Latency	Average latency for CIFS operations.
# CIFS Connected Shares	Number of SMB and SMB2 share connections.
# CIFS Connections	Number of connections.
# Open Files	Number of open files over SMB and SMB2.

NetApp Cluster-Mode Disk Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Disk Performance by Cluster

Three line charts display disk performance by cluster: busy, read/write rates, and read/write latency. Select the lines to be represented in the charts.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click a point in the line chart to view the NetApp Cluster-Mode Disk Performance by Node report.

Table 17-11 NetApp Cluster-Mode Disk Performance by Cluster

Column name	Description
Cluster Name	Name of the storage array.
Disk Busy %	Percentage of time there was at least one outstanding request to the disk.
Log Date	Timestamp at which the performance data was collected.
User Read Blocks	Number of blocks transferred per second for user read operations.
User Write Blocks	Number of blocks transferred per second for user write operations.
User Read Latency	Average latency per block in microseconds for user read operations.
User Write Latency	Average latency per block in microseconds for user write operations.

NetApp Cluster-Mode Disk Performance by Node

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Disk Performance by Node

Three line charts display disk performance by node: busy, read/write rates, and read/write latency. Select the lines to be represented in the charts.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Node Name**link to view the NetApp Cluster-Mode Node Summary report.

Table 17-12 NetApp Cluster-Mode Disk Performance by Node

Column name	Description
Cluster Name	Name of the storage array.
Disk Busy %	Percentage of time there was at least one outstanding request to the disk.
Log Date	Timestamp at which the performance data was collected.
Node Name	Name of the node.

Table 17-12 NetApp Cluster-Mode Disk Performance by Node *(continued)*

Column name	Description
User Read Blocks	Number of blocks transferred per second for user read operations.
User Write Blocks	Number of blocks transferred per second for user write operations.
User Read Latency	Average latency per block in microseconds for user read operations.
User Write Latency	Average latency per block in microseconds for user write operations.

NetApp Cluster-Mode NFS Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode NFS Performance by Cluster

Two line charts represent NFS performance by cluster metrics: NFS read/write success and NFS read/write latency. Select the lines that you want represented for the specific performance metrics.

Click the **Cluster Name** link to view the Array Capacity and Utilization report.

Table 17-13 NetApp Cluster-Mode NFS Performance by Cluster

Column name	Description
# Read Success	Number of successful read procedure requests.
# Write Success	Number of successful write procedure requests.
Avg Read Latency	Average latency of read procedure requests.
Avg Write Latency	Average latency of write procedure requests.
Cluster Name	Name of the storage array.
Log Date	Timestamp at which the performance data was collected.

NetApp Cluster-Mode NFS Performance by Vserver

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode NFS Performance by Vserver

Two line charts represent NFS performance by Vserver metrics: NFS read/write success and NFS read/write latency. Select the lines that you want represented for the specific performance metrics.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Vserver** link to view the NetApp Cluster-Mode Vserver Summary report.

Table 17-14 NetApp Cluster-Mode NFS Performance by Vserver

Column name	Description
# Read Success	Number of successful read procedure requests.
# Write Success	Number of successful write procedure requests.
Avg Read Latency	Average latency of read procedure requests.
Avg Write Latency	Average latency of write procedure requests.
Cluster Name	Name of the storage array.
Log Date	Timestamp at which the performance data was collected.
Vserver Name	Name of the Vserver. When specified as part of a Vserver-create, this is the name of the Vserver.

NetApp Cluster-Mode Volume Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Volume Performance by Cluster

Three line charts represents Cluster-Mode volume performance data by cluster: IOPS, read/write rate, and read/write latency. Select the lines to be represented in the charts.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Vserver** link to view the NetApp Cluster-Mode Vserver Summary report. Click a point in the line chart to view the NetApp Cluster Mode Volume Performance by Vserver report.

Table 17-15 NetApp Cluster-Mode Volume Performance by Cluster

Column name	Description
Cluster Name	Name of the storage array.
Total IOPS	Number of operations per second serviced by the volume.
Read IOPS	Number of reads per second to the volume
Write IOPS	Number of writes per second to the volume.
Log Date	Timestamp at which the performance data was collected.
Read Rate	Bytes read per second from the volume.
Write Rate	Bytes written per second to the volume.
Read Latency	Average latency in microseconds for the WAFL (Write Anywhere File Layout) file system to process read requests.
Write Latency	Average latency in microseconds for the WAFL (Write Anywhere File Layout) file system to process write request to the volume; not including request processing or network communication time.

NetApp Cluster-Mode Volume Performance by Vserver

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Volume Performance by Vserver

Three line charts represents Cluster-Mode volume performance data by Vserver: IOPS, read/write rate, and read/write latency. Select the lines to be represented in the charts.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Vserver** link to view the NetApp Cluster-Mode Vserver Summary report. Click a point in the line chart to view the NetApp Cluster Mode Volume Performance by Vserver report.

Table 17-16 NetApp Cluster-Mode Volume Performance by Vserver

Column name	Description
Cluster Name	Name of the storage array.
Total IOPS	Number of operations per second serviced by the volume.
Read IOPS	Number of reads per second to the volume
Write IOPS	Number of writes per second to the volume.
Log Date	Timestamp at which the performance data was collected.
Vserver Name	Name of the Vserver. When specified as part of a Vserver-create, this is the name of the Vserver.
Read Rate	Bytes read per second from the volume.
Write Rate	Bytes written per second to the volume.
Read Latency	Average latency in microseconds for the WAFL (Write Anywhere File Layout) file system to process read requests.
Write Latency	Average latency in microseconds for the WAFL (Write Anywhere File Layout) file system to process write request to the volume; not including request processing or network communication time.

NetApp Cluster-Mode Performance by Cluster

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Performance by Cluster

Four line charts represent cluster performance metrics: Operations per Second (IOPS), CPU performance, system latency, and disk read/write performance. Select the lines that you want represented for IOPS metrics: FCP IOPS, HTTP IOPS, iSCSI IOPS, NFS IOPS, NFS IOPS, CIFS IOPS.

Click the **Cluster Name** link to view the NetApp Cluster-Mode Performance by Node report.

See "NetApp Cluster-Mode Performance by Node" on page 571.

Table 17-17 NetApp Cluster-Mode Performance by Cluster

Column name	Description
Cluster Name	Name of the storage array.
CIFS IOPS	CIFS operations per second.
FCP IOPS	Fibre Channel Protocol (FCP) operations per second.
HTTP IOPS	HTTP operations per second.
iSCSI IOPS	iSCSI operations per second.
NFS IOPS	NFS operations per second.
Log Date	Timestamp at which the performance data was collected.
CPU Busy %	System CPU resource utilization.
System Read Latency	Average latency for all read operations in the system in milliseconds.
System Write Latency	Average latency for all write operations in the system in milliseconds.
Disk Read	Disk read rate in MB/Sec.
Disk Write	Disk write rate in MB/Sec.

NetApp Cluster-Mode Performance by Node

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > NetApp Cluster-Mode Performance by Node

Four line charts represent node performance metrics: Operations per Second (IOPS), CPU performance, system latency, and disk read/write performance. Select the lines that you want represented for IOPS metrics: FCP IOPS, HTTP IOPS, iSCSI IOPS, NFS IOPS, NFS IOPS, CIFS IOPS.

Click the **Cluster Name** link to view the Array Capacity and Utilization report. Click the **Node Name**link to view the NetApp Cluster-Mode Node Summary report.

Table 17-18 NetApp Cluster-Mode Performance by Node

Column name	Description
Cluster Name	Name of the storage array.
CIFS IOPS	CIFS operations per second.
FCP IOPS	Fibre Channel Protocol (FCP) operations per second.
HTTP IOPS	HTTP operations per second.
iSCSI IOPS	iSCSI operations per second.
NFS IOPS	NFS operations per second.
Node Name	The textual name of a node.
Log Date	Timestamp at which the performance data was collected.
CPU Busy %	System CPU resource utilization.
System Read Latency	Average latency for all read operations in the system in milliseconds.
System Write Latency	Average latency for all write operations in the system in milliseconds.
Disk Read	Disk read rate in MB/Sec.

Table 17-18 NetApp Cluster-Mode Performance by Node *(continued)*

Column name	Description
Disk Write	Disk write rate in MB/Sec.

Performance Profiles

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > Performance Profiles

Table 17-19 Performance Profiles

Column name	Description
Array Vendor	Array vendor
Family	Array family
Product	Array model
RAID Type	RAID type
Drive Type	Drive Type
Drive Form	Drive enclosure
Speed	Drive speed
Read Response Time	
Slow	The low end of the range of the read response time in milliseconds
Normal	Normal read response time in milliseconds
Fast	The high end of the range of the read response time in milliseconds
Write Response Time	
Slow	The low of the range of the write response time in milliseconds
Normal	Normal write response time in milliseconds

Table 17-19 Performance Profiles *(continued)*

Column name	Description
Fast	The high of the range of the write response time in milliseconds
Total Response Time	
Slow	The low of the range of the total response time in milliseconds
Normal	Normal total response time in milliseconds
Fast	The high of the range of the total response time in milliseconds
Read IO/sec	
Slow	The low end of the range of the read I/O performance metrics
Normal	Normal read I/O performance
Fast	The high end of the range of the read I/O performance metrics
Write IO/sec	
Slow	The low end of the range of the write I/O performance metrics
Normal	Normal write I/O performance
Fast	The high end of the range of the write I/O performance metrics
Read Cache Hits	
Slow	The low of the range of read cache hits
Normal	Normal read cache hits
Fast	The high of the range of read cache hits
Write Cache Hits	
Slow	The low of the range of write cache hits
Normal	Normal write cache hits
Fast	The high end of the range of write cache hits
Last Profile	Timestamp of the last time performance was profiled
Data Transferred	This data is used to weight each contributor's metrics within the community pool of performance metrics

Port Throughput by Array

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Storage Performance > Port Throughput by Array

Note: Column filtering is not supported for this tabular report.

In addition to the tabular list of performance metrics, a checkbox for each port enables graphing of the collected throughput statistics and comparisons among ports.

Table 17-20 Port Throughput by Array

Column name	Description
Storage Array	Array name
Port #	CIM port identifier
Role	Port role, such as Initiator, RCU, Target
Port Name	Port name For HP 3PAR: Ports refer to fibre channel ports.
IO/sec	For every record in the reporting interval, this is the sum of every data collection interval's IOs divided by the reporting interval in seconds.
IO %	Percentage of the total IOs for the array, across the report's interval, relative to the number of ports on the array. If there is only one port on the array, the percentage will be 100%; otherwise, the total for all ports on the array will equal 100%. This is a measure of how much of the total traffic is carried on each port.
MB Transferred	For every record in the reporting interval, this is the sum of every data collection interval's MB transferred.

Table 17-20 Port Throughput by Array *(continued)*

Column name	Description
Transferred %	Percentage of the MB transferred for this array, across the report's interval, relative to the number of ports on the array. If there is only one port on the array, the percentage will be 100%; otherwise, the total for all ports on the array will equal 100%. This is a measure of how much of the total traffic is carried on each port.
Show Chart	Check the Show Chart checkbox, then click **Show Chart** at the top of the report to graph the performance values. Multiple checks enable comparisons.

Chapter 18

Thin Provisioning Reports

This chapter includes the following topics:

- Thin Provisioning Overview
- EMC Symmetrix Compression Ratio Summary
- EMC Symmetrix SRP Summary
- EMC Symmetrix Thin Pool Summary
- EMC Symmetrix Thin Pool Detail
- EMC VNX (CLARiiON) Thin Pool Summary
- EMC VNX (CLARiiON) Thin Pool Details
- Hitachi Disk Summary
- Hitachi Dynamic Provisioning Pool Summary
- Hitachi Dynamic Provisioning Pool Utilization
- HP 3PAR Common Provisioning Group (CPG) Summary
- HP 3PAR Common Provisioning Group (CPG) Detail
- Storage Pool Summary
- Hitachi Vantara DP Pool Summary
- Hitachi Vantara Capacity by Storage

Thin Provisioning Overview

By using an over-subscription strategy, thin provisioning technology addresses the issue of inefficient storage utilization. Applications are given physical space from a

storage pool when data is written to the application, avoiding the pre-allocation commitment required by traditional storage provisioning approaches. In this way, an application can appear to have more storage than is actually reserved for it. As the application's data grows, it can be given more storage from the storage pool, with minimal management intervention. With thin provisioning, existing capacity can be more efficiently used, eliminating the need for over-provisioning and often delaying the acquisition of additional storage.

Often storage is over-promised and then actually dealt out as needed, making it incumbent upon the Storage Administrator to monitor storage usage so that an application's write requests can be fulfilled without failures.

NetBackup IT Analytics provides the following key capacity-at-risk reports to facilitate this monitoring process:

- See " Exception Summary " on page 450.
- See " Hosts at Risk " on page 452.
- See " LUNs at Risk " on page 453.
- See " NetApp Aggregates at Risk " on page 455.
- See " NetApp Volumes at Risk " on page 458.
- See " Over-Provisioned Hosts " on page 461.

EMC Symmetrix Compression Ratio Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > EMC Symmetrix Compression Ratio Summary

View compression ratios and savings for VMAX All Flash Adaptive (AFA) storage resource pools (SRP). Evaluate storage group ratios to better understand how savings can be achieved.

Table 18-1 EMC Symmetrix Compression Ratio Summary

Column name	Description
Array Name	Name of the storage array. This name is derived from the Symmetrix ID.
SRP Name	Name of the Storage Resource Pool.
Last Updated	The date and time the EMC Symmetrix SRP (Storage Resource Pool) data was last updated in the Portal database.
Virtual Provisioning	
Saved %	Percentage capacity saved due to thin provisioning.
Shared Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity, due to dedup of thin provisioning.
Overall Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity, due to thin provisioning.
Snapshot	
Saved %	Percentage capacity saved due to snap (versus full copy).
Shared Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity, due to dedup of snap.
Overall Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity, due to snap.
Compression	
VP Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity due to thin provisioning.
Snapshot Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity, due to snap.
Overall Ratio	Current ratio of logical storage (thick LUN capacity) to overall compressed capacity.
Storage Resource Pool (SRP)	
Overall Ratio	Current ratio of logical storage (thick LUN capacity) to compression and dedup.

EMC Symmetrix SRP Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > EMC Symmetrix SRP Summary

This report displays capacity usage and over-subscription data for Symmetrix thin-provisioned SRP (storage resource pools). Use this report to stay on top of SRP capacity subscription and usage for all storage resource pools.

Table 18-2 EMC Symmetrix SRP Summary

Column name	Description
SRP Name	Name of the Storage Resource Pool.
Array Name	Name of the storage array. This name is derived from the Symmetrix ID.
Usable Capacity	Indicates the usable capacity of all the disk groups in the SRP (Storage Resource Pool).
Allocated Capacity	Indicates the sum of the device allocations, snapshot allocations, and SRDF/A DSE allocations on the SRP (Storage Resource Pool).
Free Capacity	This is the difference between the usable and the allocated capacity.
Subscribed Capacity	Subscribed capacity is the sum of the sizes of all the thin devices subscribed against the SRP (Storage Resource Pool).
% Subscribed	Subscribed Capacity percentage is the percentage of the usable capacity used by thin devices.
User Subscribed Capacity	Subscribed Capacity for a user.
# Thin Devices	Number of LUNs associated with the EMC Symmetrix SRP (Storage Resource Pool).
# Disks	Number of Physical Disks associated with the EMC Symmetrix SRP (Storage Resource Pool).

Table 18-2 EMC Symmetrix SRP Summary *(continued)*

Column name	Description
Effective Used Capacity %	Percentage of effective used capacity.

EMC Symmetrix Thin Pool Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > EMC Symmetrix Thin Pool Summary

Table 18-3 EMC Symmetrix Thin Pool Summary

Column name	Description
Pool Name	The name of the storage pool links to EMC Symmetrix Thin Pool Detail.
Array Name	The array associated with the storage pool
Dev Emulation	Symmetrix drives can emulate standard disk drives, such as IBM, CKD (count-key-data), DASD, or FBA (Fixed Block Architecture) volumes.
RAID Type	RAID type
State	Pool state
Pool Capacity	The total capacity of the pool
Pool Used	The physical capacity that is assigned to all LUNs in the storage pool. From the array's perspective, consumed capacity is the same as allocated capacity.
% Used	Pool Used divided by Pool Capacity.
Pool Subscribed	The physical capacity that has been given to hosts/applications; this is from the host's perspective. This value is derived from total number of tracks in the storage pool times the pool subscribed percentage.
% Subscribed	The % of the pool capacity that has been presented to hosts/applications.

Table 18-3 EMC Symmetrix Thin Pool Summary *(continued)*

Column name	Description
# of Devices	This number is the total of both thin and data devices.
# of Hosts	The number of hosts using this pool.
# of Enabled Devices	The number of the devices that are enabled.

EMC Symmetrix Thin Pool Detail

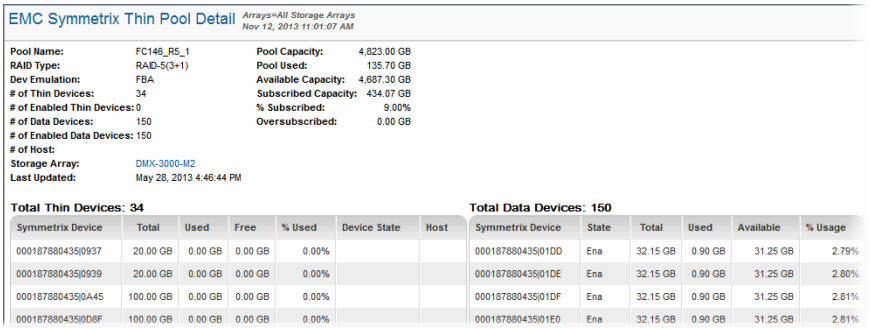
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization >EMC Symmetrix Thin Pool Summary

Click a **Pool Name** to view the Thin Pool Detail report.



EMC Symmetrix Thin Pool Detail
Arrays=All Storage Arrays
Nov 12, 2013 11:01:07 AM

Pool Name:	FC146_R5_1	**Pool Capacity:**	4,823.00 GB
RAID Type:	RAID-5(3+1)	**Pool Used:**	135.70 GB
Dev Emulation:	FBA	**Available Capacity:**	4,687.30 GB
# of Thin Devices:	34	**Subscribed Capacity:**	434.07 GB
# of Enabled Thin Devices:	0	**% Subscribed:**	9.00%
# of Data Devices:	150	**Oversubscribed:**	0.00 GB
# of Enabled Data Devices:	150		
# of Host:			
Storage Array:	DMX-3000-M2		
Last Updated:	May 28, 2013 4:46:44 PM		

Total Thin Devices: 34

Symmetrix Device	Total	Used	Free	% Used	Device State	Host
000187880435|0937	20.00 GB	0.00 GB	0.00 GB	0.00%		
000187880435|0939	20.00 GB	0.00 GB	0.00 GB	0.00%		
000187880435|0A45	100.00 GB	0.00 GB	0.00 GB	0.00%		
000187880435|0D8F	100.00 GB	0.00 GB	0.00 GB	0.00%		

Total Data Devices: 150

Symmetrix Device	State	Total	Used	Available	% Usage
000187880435|01DD	Ena	32.15 GB	0.90 GB	31.25 GB	2.79%
000187880435|01DE	Ena	32.15 GB	0.90 GB	31.25 GB	2.80%
000187880435|01DF	Ena	32.15 GB	0.90 GB	31.25 GB	2.81%
000187880435|01E0	Ena	32.15 GB	0.90 GB	31.25 GB	2.81%

Table 18-4 EMC Symmetrix Thin Pool Detail

Column name	Description
Pool Name	The name of the storage pool links to EMC Symmetrix Thin Pool Detail.
Storage Array	The array associated with the storage pool.
Dev Emulation	Symmetrix drives can emulate standard disk drives, such as IBM, CKD, DASD, or FBA (Fixed Block Architecture) volumes.
RAID Type	RAID type.
# of Thin Devices	The number of thin devices that are used to provision storage to a host as the storage is needed.
# of Enabled Thin Devices	The number of thin devices in this pool that are enabled.
# of Data Devices	The number of data devices that are providing the actual physical storage to the thin devices.
# of Enabled Data Devices	The number of the data devices that are enabled.
# of Hosts	The number of hosts using this pool.
Pool Capacity	The total capacity of the pool.
Pool Used	The physical capacity that is assigned to all LUNs in the storage pool. From the array's perspective, consumed capacity is the same as allocated capacity.
Available Capacity	Pool Capacity minus Pool Used.
Subscribed Capacity	The physical capacity that has been given to hosts/applications; this is from the host's perspective.
% Subscribed	Percentage of the pool capacity that has been presented to hosts/applications.
Oversubscribed	The amount of storage that has been given to hosts, but that exceeds the amount of physical storage.
Total Thin Devices	
Symmetrix Device	The name of the array and drive.
Total	Total amount of thin-provisioned storage on this device.
Used	Amount of provisioned storage that is in use.

Table 18-4 EMC Symmetrix Thin Pool Detail *(continued)*

Column name	Description
Free	Amount of provisioned storage that is available.
% Used	Percentage of provisioned storage that is in use.
Device State	Device state at last data collection.
Host	Host that is using the storage.
Total Data Devices	
Symmetrix Device	The name of the array and drive.
State	State of the drive.
Total	Total amount of physical storage on the drive.
Used	Total amount of storage that is in use.
Available	Amount of storage that is available.
% Usage	Percentage of storage on this device that is in use.

EMC VNX (CLARiiON) Thin Pool Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity and Utilization > EMC VNX (CLARiiON) Thin Pool Summary

Table 18-5 EMC VNX (CLARiiON) Thin Pool Summary

Column name	Description
Pool Name	The name of the storage pool links to EMC CLARiiON Thin Pool Details.
Array Name	The array associated with the storage pool
RAID Type	RAID type
Raw Capacity	The raw capacity of the pool

Table 18-5 EMC VNX (CLARiiON) Thin Pool Summary *(continued)*

Column name	Description
User Capacity	The physical capacity that's available to all LUNs in the storage pool; also known as usable capacity or host-visible capacity.
Consumed Capacity	The physical capacity that is assigned to all LUNs in the storage pool. From the array's perspective, consumed capacity is the same as allocated capacity.
% Used	Consumed Capacity divided by User Capacity
Available Capacity	The physical pool storage that is not allocated to pool LUNs
Pool Subscribed	The physical capacity that has been given to hosts/applications; this is from the host's perspective
% Oversubscribed	Oversubscribed capacity refers to Subscribed Capacity that exceeds User Capacity.
# of Thin LUNs	Thin LUNs contain the capacity that is visible to hosts (applications); this capacity typically is greater than the actual because it is from the host's perspective and is the subscribed capacity total. Drill down to the LUN Utilization Summary for details.
# of Hosts	The number of hosts using this pool. Drill down to the Host Capacity & Utilization report.
Oversubscribed	Oversubscribed capacity refers to User Capacity that exceeds the physical capacity of the pool.

EMC VNX (CLARiiON) Thin Pool Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > EMC VNX (CLARiion) Thin Pool Details

Click **Pool Name** to view the Thin Pool Details.

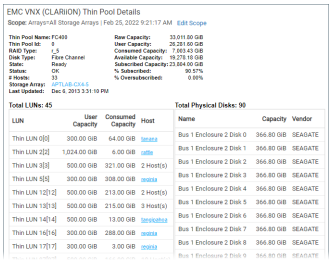


Table 18-6 EMC VNX (CLARiiON) Thin Pool Details

Column name	Description
Thin Pool Name	The name of the storage pool links to EMC CLARiiON Thin Pool Detail.
Thin Pool ID	The unique pool ID
RAID Type	RAID type
Disk Type	Type of disk, such as Fibre Channel
State	State of the pool
Status	Pool status
# Hosts	The number of hosts that have been allocated storage from this pool
Storage Array	The array associated with the storage pool
Raw Capacity	The raw capacity of the pool
User Capacity	The physical capacity that's available to all LUNs in the storage pool; also known as usable capacity or host-visible capacity.
Consumed Capacity	The physical capacity that is assigned to all LUNs in the storage pool. From the array's perspective, consumed capacity is the same as allocated capacity.
Available Capacity	The physical pool storage that is not allocated to pool LUNs
Subscribed Capacity	The physical capacity that has been given to hosts/applications; this is from the host's perspective
% Subscribed	The percentage of the raw capacity that has been given to hosts

Table 18-6 EMC VNX (CLARiiON) Thin Pool Details *(continued)*

Column name	Description
% Oversubscribed	Oversubscribed capacity refers to User Capacity that exceeds the pool's physical capacity.
Pool LUN Capacity and Usage	
LUN	Name of the LUN associated with this Thin Pool
User Capacity	Amount of capacity presented to the host for use
Consumed Capacity	Amount of capacity actually in use by the host(s)
Host	Links to the Host Utilization Details.
Total Physical Disks	
Name	Name of the physical disk that contains the storage for this LUN
Capacity	Capacity of the disk
Vendor	Manufacturer of the disk

Hitachi Disk Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager >Thin Provisioning Capacity & Utilization > Hitachi Disk Summary

Table 18-7 Hitachi Disk Summary

Column name	Description
Device Name	The name of the disk
Storage Array	The storage array in which the disk is located; links to the Array Capacity & Utilization report

Table 18-7 Hitachi Disk Summary *(continued)*

Column name	Description
Group Name	The group to which the disk belongs; links to the Array Group Details report
Capacity	Capacity of the disk
Vendor	Vendor of the disk
Serial #	Serial number of the disk
Model	Disk model
Firmware Version	The disk's firmware version
RPM	Disk revolutions per minute
Disk Type	FC, SATA, BD, SAS, or Unknown
Chassis #	The chassis number
Role	unknown, data, unallocated, offline, outofservice, spare, spareunitialized, n/a

Hitachi Dynamic Provisioning Pool Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > Hitachi DP Pool Summary

The Dynamic Provisioning Pool uses virtual volumes for data storage, with actual capacity available upon request to a host writing to LUN storage. Use this report to ensure that your pool always has sufficient free capacity, enabling on-demand provisioning to hosts.

Table 18-8 Hitachi Dynamic Provisioning Pool Summary

Column name	Description
Pool Capacity	

Table 18-8 Hitachi Dynamic Provisioning Pool Summary *(continued)*

Column name	Description
Pool ID	The name of the Dynamic Provisioning Pool.
Storage Array	The storage array on which the pool resides.
Used%	Mouse over the thermometer to view the percentage of the pool that is used. Used capacity is the total capacity of all the pages that have data written to them; that is, consumed or touched storage. The percentage is the used size divided by the total physical capacity.
# of DP Pool VOLs	Number of Hitachi journal pool volumes. For AMS (AMS200/500/1000) and WMS arrays: ■ This is the number of the QuickShadow volumes. For AMS (AMS2010/2100/2300/2500, MSS) arrays: ■ If the pool function is 3, then this is the number of the QuickShadow volumes or TrueCopyAsync volumes. For USP arrays: ■ If the pool function is 3, then this is the number of QuickShadow volumes. ■ If the pool function is 4, then this is the number of the Universal Replicator Journal volumes (for Open Volume). For USP_V arrays: ■ If the pool function is 3, this is the number of the QuickShadow volumes. ■ If the pool function is 4, this is the number of the Universal Replicator Journal volumes (for Open Volume). ■ If the pool function is 5, this is the number of the Hitachi Dynamic Provisioning Pool (HDP) volumes. For VSP arrays: ■ If the pool function is 3, then this is the number of the QuickShadow volumes. ■ If the pool function is 4, this is the number of the Universal Replicator volumes. ■ If the pool function is 5, this is the number of the HDP/HDT volumes.
# Array Groups	Number of Array Groups. Links to List Array Groups report.
Capacity of DP Pool	Virtual volumes can have a capacity that is greater than the pool capacity; however, it is imperative that you monitor the pool capacity to ensure that space is available on demand.

Table 18-8 Hitachi Dynamic Provisioning Pool Summary *(continued)*

Column name	Description
Consumed	This is consumed or touched capacity; that is, the sum of the 42 MB pages assigned from the pool that have data written to them.
Free In Pool	Capacity - Used Size = Available Size. Monitor this space to ensure that you do not deplete the available space in a pool, ensuring that host requests for storage are not denied
Physical Capacity	
Total	Physical capacity of pool (KB). This item is output when the value of poolFunction is 5 and the storage system is VSP G1000 (with a microcode version of 80-04-0X-XX/XX or later), VSP Gx00 models (with a microcode version of 83-03-0X-XX/XX or later), or VSP Fx00 models (with a microcode version of 83-03-0XXX/XX or later).
Used	Physical capacity usage of the pool (KB). This item is output when the value of poolFunction is 5 and the storage system is VSP G1000 (with a microcode version of 80-04-0X-XX/XX or later), VSP Gx00 models (with a microcode version of 83-03-0X-XX/XX or later), or VSP Fx00 models (with a microcode version of 83-03-0XXX/XX or later).
Used %	The percentage is the used size divided by the total physical capacity.
Physical FMC Pool Volumes	
Total	Total physical capacity of the pool volume whose drive type is FMC (KB). This item is output when the value of poolFunction is 5 and the storage system is VSP G1000 (with a microcode version of 80-04-0X-XX/XX or later), VSP Gx00 models (with a microcode version of 83-03-0X-XX/XX or later), or VSP Fx00 models (with a microcode version of 83-03-0XXX/XX or later).
Used	Total physical capacity usage of the pool volume whose drive type is FMC (KB). This item is output when the value of poolFunction is 5 and the storage system is VSP G1000 (with a microcode version of 80-04-0X-XX/XX or later), VSP Gx00 models (with a microcode version of 83-03-0X-XX/XX or later), or VSP Fx00 models (with a microcode version of 83-03-0XXX/XX or later).

Table 18-8 Hitachi Dynamic Provisioning Pool Summary *(continued)*

Column name	Description
FMC Pool Volume Capacity	
Total	Total capacity of the pool volume whose drive type is FMC (KB). This item is output when the value of poolFunction is 5 and the storage system is VSP G1000 (with a microcode version of 80-04-0X-XX/XX or later), VSP Gx00 models (with a microcode version of 83-03-0X-XX/XX or later), or VSP Fx00 models (with a microcode version of 83-03-0XXX/XX or later).
Used	Total capacity usage of the pool volume whose drive type is FMC (KB). This item is output when the value of poolFunction is 5 and the storage system is VSP G1000 (with a microcode version of 80-04-0X-XX/XX or later), VSP Gx00 models (with a microcode version of 83-03-0X-XX/XX or later), or VSP Fx00 models (with a microcode version of 83-03-0XXX/XX or later).
Saving	Saved Capacity for FMC (KB).
Virtual	
# of DP VOLs	The number of virtual volumes in this Dynamic Provisioning pool; links to the LUN Utilization Summary. This is capacity that is available to be allocated to hosts; associated capacity.
Capacity of DP VOLs	This is the capacity of all DP Volumes that have been allocated to hosts. It also includes overhead for dynamic provisioning. This capacity could be over-subscribed, which means that the capacity shown here could exceed the capacity shown in the Capacity column.
Overprovision	Percentage = Capacity of DP Vols/Capacity of DP Pool
Allocated	Allocated Capacity (KB).
Threshold1	Journal Pool threshold used to monitor consumed capacity.
Threshold2	This threshold is hard-coded in the firmware of the array.
Status	Reported by Device Manager: Normal, Blocked, Over Threshold

Hitachi Dynamic Provisioning Pool Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > Hitachi DP Pool Utilization

This dashboard provides an at-a-glance view of dynamically provisioned storage, enabling insight into potential exposure, as well as, storage that may be reclaimable.

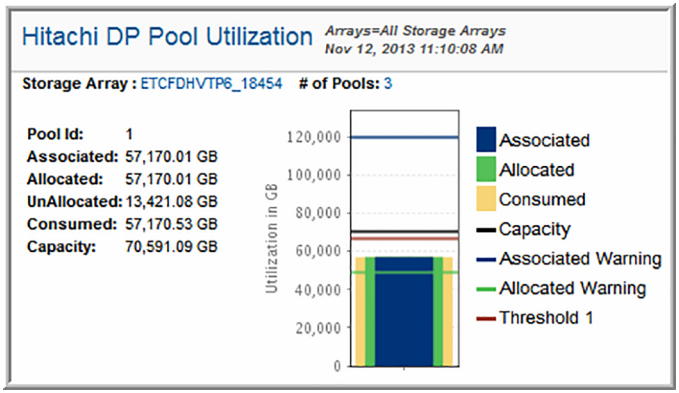


Table 18-9 Hitachi Dynamic Provisioning Pool Utilization

Column name	Description
Pool ID	The name of the Dynamic Provisioning Pool
Array	The storage array on which the pool resides
Associated	Of the DP Vols that have been created in the pool, Associated storage represents potential promised capacity; that is, the amount of storage that can be given to a host. The Associated Warning line defaults to 170%.
Allocated	The green bar in the chart indicates exposure; this is storage that can be written to by a host. When the green bar crosses the Capacity line, there is potential that hosts may attempt to write to storage that is not physically available. The Allocated Warning line defaults to 70%.
Unallocated	Unallocated = Capacity - Allocated storage

Table 18-9 Hitachi Dynamic Provisioning Pool Utilization *(continued)*

Column name	Description
Consumed	This is Consumed, or touched capacity; that is, the sum of the 42 MB pages assigned from the pool that have data written to them. Consumed could be greater than Allocated, in cases where de-provisioned hosts have not had storage reclaimed. Consumed can never exceed capacity, but as it nears capacity, this chart provides an indication of a potential issue.
Capacity	This is the total physical size of all DP pool volumes that have been created within the pool. When Allocated (green bar) crosses this line, there may not be storage available for future writes. Note that the Threshold 1 line represents a % of capacity that is user-defined. If not customized, Threshold 1 defaults to 70%.
Threshold1	This is a user-definable threshold to monitor pool usage. If not customized by the user, the default value is 70%.

Reclaimable Storage

There are times when storage is not in use, yet it's difficult to make that determination. For example, a host can be de-provisioned and therefore the storage is un-masked, but data remains on the disk. These pages could be freed up and storage reclaimed. In the Hitachi DP Pool Utilization report, storage may be reclaimable when the Consumed value is greater than the Allocated value.

HP 3PAR Common Provisioning Group (CPG) Summary

A common provisioning group (CPG) is a virtual pool of logical disks that allows up to 64 volumes.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal. Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > HP 3PAR Common Provisioning Group (CPG) Summary

The following illustration represents the view available for HP 3PAR storage arrays.

Table 18-10 HP 3PAR Common Provisioning Group (CPG) Summary

Column name	Description
CPG Name	Name assigned to the Common Provisioning Group (CPG). See "HP 3PAR Common Provisioning Group (CPG) Detail" on page 594.
Storage Array	The name of the storage array.
Allocated	Sum of the allocated size for all associated logical drives.
User	Sum of the space allocated for users on all associated logical drives.
Snap	Sum of the snap size for all associated logical drives.
Admin	Sum of the admin size for all associated logical drives.
Logical Drives	Number of allocated logical drives.
Allocation Warning	User-defined threshold that can be set for Thinly Provisioned Virtual Volumes (TPVVs) and Commonly Provisioned Virtual Volumes (CPVVs) to alert users when they reach a certain size.
Growth Incr	Unit of storage by which additional logical disks are created and allocated to a Common Provisioning Group (CPG).
Growth Warning	User-defined threshold that can be set for Common Provisioning Groups (CPGs) to alert users when they reach a certain size. Growth warning is defined as: Growth Limit - (Allocation Warning * TotalManagedSpace)/100.
Growth Limit	User-defined threshold that can be set for Common Provisioning Groups (CPGs) to cap their potential size.
% Allocated	Calculated as Allocated/Growth Limit.
Device Type (Speed)	Indicates the specified device type for physical disks: FC for Fibre Channel, NL for Nearline, and SSD for Solid State Drive. Indicates a rough performance estimation between the drive and the other drives in the system. This does not represent a rotational speed for the drives without spinning media (SSD). For FC and NL drives, the number corresponds to both a performance measurement and an actual rotational speed. For SSD drives, the number is a relative performance benchmark that takes into account in I/O per second, bandwidth and the access time.
RAID Type	Indicates the RAID type for the logical disk: r0 = RAID-0, r1 = RAID-1, r5 = RAID-5 and r6 = RAID-6. RAID10 is displayed as RAID1+0.

Table 18-10 HP 3PAR Common Provisioning Group (CPG) Summary *(continued)*

Column name	Description
Set Size	Displays the set size in the number of chunklets. The default relies on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.
Availability	Normal or degraded
Step Size	Value derived directly from the Virtual Volume (VV) or from the Common Provisioning Group (for Thinly Provisioned VVs and Commonly provisioned VVs).
Row Size	Indicates the number of sets in a row. If the field is blank, no row limit is imposed.
Disk Filter	Displays the specific logical disks for the Common Provisioning Group (CPG).

HP 3PAR Common Provisioning Group (CPG) Detail

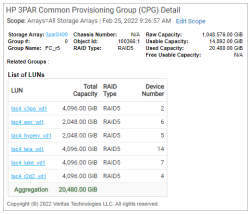
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > HP 3PAR Common Provisioning Group (CPG) Summary

Click the CPG name.


HP 3PAR Common Provisioning Group (CPG) Detail

Scope: Arrays>All Storage Arrays | Feb 25, 2022 9:26:57 AM Edit Scope

Storage Array: 3par6400 Chassis Number: N/A Raw Capacity: 1,048,574.66 GiB
Group #: 0 Object Id: 100366:1 Usable Capacity: 14,892.66 GiB
Group Name: FC_r5 RAID Type: RAID5 Used Capacity: 20,480.00 GiB
Free Usable Capacity: N/A

Related Groups

List of LUNs

LUN	Total Capacity	RAID Type	Device Number
lace4_c2p0_vd1	4,096.00 GiB	RAID5	2
lace4_was_vd1	2,048.00 GiB	RAID5	6
lace4_hyperv_vd1	2,048.00 GiB	RAID5	5
lace4_lace_vd1	4,096.00 GiB	RAID5	14
lace4_kate_vd1	4,096.00 GiB	RAID5	7
lace4_r2k2_vd1	4,096.00 GiB	RAID5	4
Aggregation	20,480.00 GiB		

Copyright (©) 2022 Veritas Technologies LLC. All rights reserved.

Table 18-11 HP 3PAR Common Provisioning Group (CPG) Detail

Column name	Description
Storage Array	The name of the storage array. See "Array Capacity & Utilization (HP 3PAR View)" on page 197.
Group #	The number of the Array Group.
Group Name	The array group name.
Chassis Number	The chassis to which the group belongs.
Object ID	The object ID, for object-based storage.
RAID Type	Indicates the RAID type for the logical disk.
Raw Capacity	Total capacity for this array group.
Usable Capacity	Total storage capacity of the group.
Used Capacity	Amount of space that has already been used.
Free Usable Capacity	Available storage.

Storage Pool Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Capacity Manager > Thin Provisioning Capacity & Utilization > Storage Pool Summary

This report displays capacity usage and subscription data for storage pools.

Use this report to stay on top of storage pool capacity subscription and usage for all storage pools. The hyperlinks enable drilldown to the LUN Utilization Summary for all associated LUNs in a storage pool and the Host Capacity & Utilization report, for the hosts associated with a storage pool.

Table 18-12 Storage Pool Summary

Column name	Description
Pool Name	Name of the storage pool.
Storage Array	The name of the storage array. See “ Array Capacity and Utilization ” on page 164.
RAID Type	RAID type of the storage pool.
Pool Type	Type of storage pool.
Status	Status of the storage pool.
Raw Capacity	Raw capacity the storage pool. (# of drives * drive size).
User Capacity	User capacity (The pool capacity available to LUNs from this pool after definition).
Consumed Capacity	Amount of data written to the pool.
% Used	Consumed capacity divided by User Capacity.
Available Capacity	Amount of space available in the pool.
Subscribed Capacity	The physical capacity that has been given to hosts from the host perspective.
% Oversubscribed	Oversubscribed percentage refers to Subscribed Capacity that exceeds User Capacity.
# LUNs	Number of logical units in the storage pool. See “ LUN Utilization Summary ” on page 344.

Table 18-12 Storage Pool Summary *(continued)*

Column name	Description
# Hosts	Number of hosts that allocated LUNs. See " Host Capacity & Utilization " on page 478.

Hitachi Vantara DP Pool Summary

NetBackup IT Analytics allows to explore the data center using customizable report templates or by using parts of infrastructure as entry points.

Use **Search** to find reports, templates, and dashboards across the portal.

Use the **Reports** tab to examine the catalog of templates, dashboards, and reports which are organized by the products along with user-created, and system folders.

Navigate to *Reports* > *Capacity Manager* > *Thin Provisioning Capacity & Utilization* and then click *Hitachi Vantara DP Pool Summary.* Hitachi Vantara DP Pool Summary report is displayed.

The report displays the virtual volumes for data storage with actual capacity available. Use this report to ensure that the pool has sufficient free capacity, with percentage and ratio values.

Table 18-13 Hitachi Vantara DP Pool Summary report details

Column name	Description
Model	Storage array type
Serial Number	Serial number of the storage array
Pool Type	Type of the pool. Example: Dynamic Provisioning,Dynamic Provisioning (Tiering),Data Direct Mapping, Thin Image
Collection Time	Time (in GMT) when the data was collected from the storage system
Logical Capacity	Capacity of the pool
Used Logical Capacity	Amount of pool capacity actually being used (KB)
Used Logical Capacity in Percentage	Usage rate of the pool capacity
Available Physical Volume Capacity	Total amount of the physical capacity of the parity groups that make up the pool (KB)

Table 18-13 Hitachi Vantara DP Pool Summary report details *(continued)*

Column name	Description
Available Physical FMC Pool Volume Capacity	Total amount of used capacity in the physical capacity of the parity groups that make up the pool (KB)
Subscribed Capacity	Virtual capacity of all Dynamic Provisioning V- VOLs mapped to the pool (KB)
Subscription Percentage	Virtual capacity of all Dynamic Provisioning V- VOLs mapped to the pool (KB)
Threshold Rate	System threshold or depletion threshold set for the percentage of used pool space
Auto Expansion	Indicates whether to automatically add pool volumes according to the compression ratio of a parity group for which accelerated compression is enabled
Data Volumes Used Capacity	Capacity used by a virtual volume (KB)
Data Reduction Capacity	The capacity saved by capacity saving or accelerating the compression of the parity group (KB)
Data Reduction Ratio	Saving effect achieved by capacity saving or accelerating the compression of the parity group. This is the value of N when the ratio of capacity before and after the reduction is N:1
Deduplication Capacity	Capacity saved by deduplication of capacity saving (KB)
Deduplication Ratio	Ratio of disk space that is saved through deduplication when using capacity saving. This is the value of N when the ratio of capacity before and after the reduction is N:1
Compression Capacity	Capacity compressed by capacity saving (KB)
Compression Ratio	Ratio of disk space that is saved through compression when using capacity saving. This is the value of N when the ratio of capacity before and after the reduction is N:1
Deduplication and Compression Capacity	Capacity reduced when deduplication and compression are used in a pool
Deduplication and Compression Ratio	Saving effect achieved by capacity saving. This is the value of N when the ratio of capacity before and after the reduction is N:1

Table 18-13 Hitachi Vantara DP Pool Summary report details *(continued)*

Column name	Description
Pattern Matching	Capacity saved by reclaiming the specified data pattern using capacity saving (KB)1
System Data	The amount of consumed system data (KB)
Preprocessed Data	Capacity of data (before reduction) to be reduced by capacity saving (KB)
FMC Logical Total	The total logical capacity of the pool volume for FMC or the pool volume that supports accelerated compression used in the pool (KB)
FMC Logical Used	The logically used capacity of the pool volume for FMC or the pool volume that supports accelerated compression used in the pool (KB)
FMC Physical Total	The total physical capacity of the pool volume for FMC or the pool volume that supports accelerated compression used in the pool (KB)2
FMC Physical Used	The actually used capacity of the parity group for FMC or the parity group that supports accelerated compression (KB)
FMC Compression Capacity	The compressed capacity with the accelerated compression function (KB)
Demand	Displays the sum of free capacity of allocated DP volumes in a pool.
FMC Compression Ratio	Saving effect achieved by accelerated compression. This is the value of N when the ratio of capacity before and after the reduction is N:1
FMC Pattern Matching	The reclaimed capacity with the accelerated compression function (KB)
Creation Date	Creation Date
Last Updated	Last updated date

Hitachi Vantara Capacity by Storage

NetBackup IT Analytics allows to explore the data center using customizable report templates or by using parts of your IT infrastructure as entry points.

Use Search to find reports, templates and dashboards across the portal.

Access the report to examine the catalog of templates, dashboard, and reports organized by the products along with the user created and system folders.

Navigate to **Reports** > **Capacity Manager** > **Thin Provisioning Capacity & Utilization** > and then click **Hitachi Vantara Capacity By Storage**.

The **Scope Selector** dialog box is displayed. Select or specify the required values and then click **Generate**.

The **Hitachi Vantara Capacity by Storage** report is displayed with the specfied criteria in the scope.

The report displays the storage has sufficient capacity.

The following table contains report columns which are displayed.

Table 18-14 Report columns details

Column name	Description
Storage Array Name	Name of storage array.
Model	Storage array type such as *VSP F1500*.
Serial #	Serial number of the storage array.
Thin Total	Thin total capacity.
Thin Used	Thin utilized capacity.
Thin Free	Thin free capacity.
Subscribed total	Total capacity subscribed.
FMC Capacity Saving	FMC capacity savings.
Deduplication/ Compression Saving Capacity	Deduplication/ Compression Saving capacity.
Total Parity Group Free	Total Parity Group Free
Creation Date	Creation date of storage array.
Last Updated	Storage array updated date.

Chapter 19

File Analytics Reports

This chapter includes the following topics:

- File Analytics Overview
- Capacity Report
- Capacity Forecast
- Capacity Forecast by Device
- Capacity Forecast by Volume/Share
- Host Summary
- Top Candidates for Duplicates
- Volume Summary Report
- Share Summary Report
- New Files Added
- Usage by Owner
- Usage by Owner Forecast
- Usage by Owner Forecast Details
- Duplicate Files (top 50 in each device)
- Deduplication Candidates
- File Categories
- Non-organizational Files
- Ransomware File Types

- File Sizes
- File Types
- Largest Files
- File List Preview
- Access Dates
- Create Dates
- Modified Dates

File Analytics Overview

File Analytics offers a highly optimized, lightweight data collector that uncovers files that are needlessly consuming storage. The bulk of wasted space often can be attributed to unstructured data--that is, files that contain data such as email, word processing documents, presentations, and digital media. Unlike structured data, which is managed by database administrators, unstructured data tends to proliferate with reckless abandon, resulting in stagnant and duplicated data. To assess the value of data files, whether they are critical to operations or are merely forgotten files, File Analytics provides data profiling. This data classification enables aggregation of file types and data growth forecasting. In addition, storage administrators can make archiving decisions based on the age of the data and its relevance to business objectives.

Once data collection completes, a number of reports provide insight into file profiles in your enterprise. Using built-in categorization and also file categories that you can customize, File Analytics becomes part of your arsenal of storage management tools. File Analytics aggregates the data, making it easy for you to identify offenders, such as:

- Deduplication candidates
- File ownership (or lack thereof)
- Largest files
- File Categories (Top file types by volume)

Capacity Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Capacity Management & Chargeback > Capacity

The File Analytics Capacity report can be used as an indicator of storage efficiency. This report lists current capacity and utilization, which aids in maintaining volumes at optimum utilization.

See "File Analytics Overview" on page 602.

Table 19-1 Capacity Report

Column name	Description
Host	Host name for which volume/share data has been collected
Volume/Share	The name of the volume or share
Type	Volume or Share
Capacity	The available capacity of the volume or NetApp device
% of Total	The % of the total capacity for volume, share, or NetApp filer capacity, relative to the total capacity
Used	Used capacity from the OS perspective for volumes (Windows and Linux) and NetApp devices; null for CIFS shares
% of Capacity	Indicates the % of the total volume/share capacity that is in use, calculated as Used/Capacity

Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Capacity Management & Chargeback > Capacity Forecast

See "File Analytics Overview" on page 602.

Using historical data, you can view graphs representing anticipated capacity usage by device, and volume/share.

Table 19-2 Capacity Forecast

Column name	Description
Device	Links to Capacity Forecast by Device.
Volume/Share	Links to Capacity Forecast by Volume/Share.
Current	
Unique Owners	Number of unique users with files consuming capacity.
# Files	Number of files consuming capacity.
Used (6 months)	Amount of storage used by the files (from an OS perspective: ONTAP, Windows, or Linux) in the last six months displayed in an area chart. Note, six months is a default value.
# Files (6 months)	Numbers of files consuming capacity in the last six months displayed in an area chart. Note, six months is a default value.
Used	Amount of storage used by the files (from an OS perspective: ONTAP, Windows, or Linux). For a share, this value may represent the total storage used on the device across all the shares on the device.
Presented	The sum of all the member file sizes on the Volume/Share; this sum could be different from the Volume/Share Used value when de-duplication or compression is in play.
Util %	Used/Capacity.
Capacity	Allocated size of the Volume/Share. For a share, this value may represent the total storage capacity allocated to the device for all the shares on the device.
Estimated Out of Space	
Days	Extrapolated days left prior to 100% used capacity (volumes).
Date	Extrapolated date of 100% used capacity (volumes).
Description	This is the description that was entered when the host was created in the Portal database.

Capacity Forecast by Device

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports

cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

File Analytics > Capacity Management & Chargeback > Capacity Forecast

Click a Device link.

Using historical data, this graph represents anticipated capacity usage by device.

See "File Analytics Overview" on page 602.

Capacity Forecast by Volume/Share

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

File Analytics > Capacity Management & Chargeback >Capacity Forecast

Click a Volume/Share link.

Using historical data, this graph represents anticipated capacity usage by volume or share.

See "File Analytics Overview" on page 602.

Host Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Capacity Management & Chargeback > Host Summary

This report lists a capacity summary for the hosts from which file data has been collected with drilldowns into Volumes and Shares. Use this report to get an overview of how devices and volumes are being used, including an indicator of how much is duplicated.

See "File Analytics Overview" on page 602.

Table 19-3 Host Summary

Column name	Description
Device	The name of the device from which file data has been collected.
Type	The type: NetApp, Windows, Unix, CIFS.
Last Collected	The date of the last data collection.
# Volumes	Number of volumes collected, blank for Linux and Shares, links to the Volume Summary report.
# Shares	Number of shares, links to the Share Summary report.
Total Files	The total number of files collected.
Presented	The sum of the sizes of the files collected for a volume, share, or device.
Unique Owners	The total number of unique owners encountered.
Duplication	The percentage of the files that have been detected as duplicates. Links to the Top Candidates for Duplicates report.

Top Candidates for Duplicates

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

File Analytics > Capacity Management & Chargeback > Host Summary

Click a Duplication link.

The Top Candidates report summarizes the names of files that have been found to be duplicated throughout your storage environment.

See "File Analytics Overview" on page 602.

Table 19-4 Top Candidates for Duplicates

Column name	Description
Name	Name of the duplicate file
# of Files	Number of files found for this name
Last Modified	Last modified date for the files with this name
Size	Total size of the duplicates

Volume Summary Report

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

File Analytics > Capacity Management & Chargeback > Host Summary

Click a link in the # Volumes column.

See "File Analytics Overview" on page 602.

Table 19-5 Volume Summary Report

Column name	Description
Volume	The name of the drive or volume
Total Files	The total number of files collected for this volume
Total File Size	The sum of the sizes of the files collected for this volume
Unique Owners	The total number of unique owners encountered for this volume
Duplication	The percentage of files that are detected as duplicates

Share Summary Report

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

File Analytics > Capacity Management & Chargeback > Host Summary

Click a link in the # Shares column.

See “File Analytics Overview” on page 602.

Table 19-6 Share Summary Report

Column name	Description
Share	The name of the share
Total Files	The total number of files collected for this share
Total File Size	The sum of the sizes of the files collected for this share
Unique Owners	The total number of unique owners encountered for this share
Duplication	The percentage of files that are detected as duplicates

New Files Added

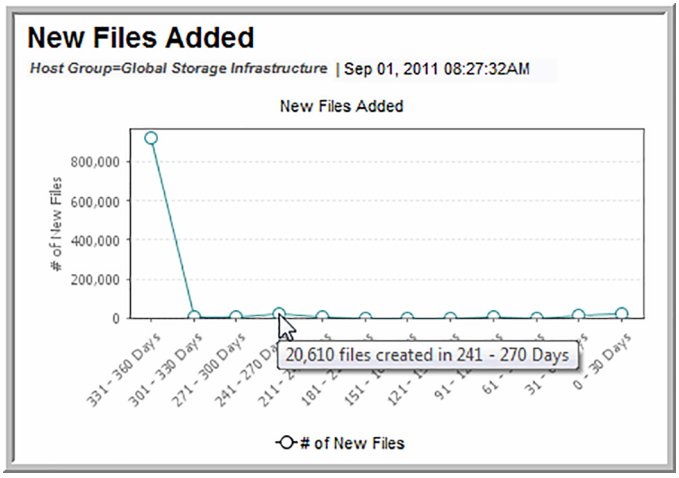
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Capacity Management & Chargeback > New Files Added

The New Files Added report provides a graphical representation of the trend of the number of files that have been added, over time, in your environment. File Analytics data collection groups the files by historic create dates and stores the metadata for these files in the NetBackup IT Analytics database for reports such as this.

Mouse over the data points in the graph to view the number of files added for the associated time period.



See "File Analytics Overview" on page 602.

Usage by Owner

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Capacity Management & Chargeback > Usage by Owner

This report lists, by user, essential storage usage information for files collected. The highest users are listed first. This data could be used to support chargeback reporting.

See "File Analytics Overview" on page 602.

Table 19-7 Usage by Owner

Column name	Description
Owner	The owner of the file data that has been collected
Total Files	The total number of files collected for this owner
Presented	The sum of the sizes of the files collected for this owner
Host	The host on which the files reside. This column is NOT displayed if Group by Owner was selected in the Report Designer.
Volume/Share	The specific location of the files. This column is NOT displayed if Group by Owner was selected in the Report Designer.
Export List	Click this checkbox to launch a window where you can specify the criteria for the file types to be exported.

Usage by Owner Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Capacity Management & Chargeback > Usage by Owner Forecast

See "File Analytics Overview" on page 602.

Table 19-8 Usage by Owner Forecast

Column name	Description
Owner	File owner; click this link to view a graph: Usage by Owner Forecast Details
Host	IP address or host name where the files are located
# Shares	Number of shares on which files are located
Current	
# Files	Number of files for this owner
Presented	Amount of storage consumed by this owner

Usage by Owner Forecast Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

File Analytics > Capacity Management & Chargeback > Usage by Owner Forecast

Click a link in the Owner column.

Use this graph to plan storage usage by individual owners. Using historical usage data, this report forecasts anticipated usage over the next six months.

See “File Analytics Overview” on page 602.

Duplicate Files (top 50 in each device)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > Duplicate Files

Duplicate files often go undiscovered, consuming more than their fair share of storage. This report reveals the duplicate files found in your enterprise. These duplicate files represent an opportunity to reclaim wasted space on NAS volumes.

See "File Analytics Overview" on page 602.

Table 19-9 Duplicate Files (top 50 in each device)

Column name	Description
Device	Device on which the duplicates were found
Volume/Share	The Volume/Share on which the duplicates were found
Name	Name of the file for which duplicates have been found
# Files	Number of duplicates found for this file
Last Modified	The most recent modification timestamp among the duplicates
Size	Total size of all of this file's duplicates
Export List	Check this box to launch the File List Export utility. In the Add Export Request window, the default file size range represents the Size divided by the # of Files divided by 1024 (to convert from MB to KB).

Deduplication Candidates

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > Deduplication Candidates

See "File Analytics Overview" on page 602.

Table 19-10 Deduplication Candidates

Column name	Description
Host	Host name
Volume/Share	Where the duplicates reside
Type	Volume or share
Total Files	The total number of files collected
Size	Sum of the sizes of all files collected
Total Duplicate Files	The total of all duplicate files on the device
Total Duplicate Size	The total of the size of all the duplicate files on the device
Duplicate %	The percentage of the size of the files that are duplicates

File Categories

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > File Categories

This report aggregates file types into application categories. Bar charts provide an at-a-glance view of storage consumption by application.

See “File Analytics Overview” on page 602.



Table 19-11 File Categories

Column name	Description
Category Name	Application category
# Files	Number of files in this category
Average Size	Average size of the files in this category
Total Size	Total size of the files in this category
Size %	Size % of the files in this category
Export List	Click this checkbox to launch a window where you can specify the criteria for the categories to be exported.

Non-organizational Files

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > Non-organizational Files

See “File Analytics Overview” on page 602.

This report aggregates non-organizational files types for every user. The files comprise of audio and video file categories with extensions such as aiff, au, mp3, mp4, wav, wma, mpg, mpeg, avi, mov, m4v.

Table 19-12 Non-organizational Files

Column name	Description
Owner	User name of the file owner
Total Files	File count
File size	Total size of all the files

The report provides a drilldown view aggregated by file category for each owner.

Table 19-13 Non-organizational Files

Column name	Description
Category	Type of file.
Total files	Type-specific file count
Total Size	Total size of all the files of the respective category.

Ransomware File Types

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > **Storage Efficiencies** > **Ransomware File Types**

This report aggregates ransomware files by their respective type. Bar charts provide an at-a-glance view of the number of files and their size.

See "File Analytics Overview" on page 602.

Table 19-14 Ransomware File Types

Column name	Description
File Type	Ransomware file type
# Files	Number of files in this type
Average Size	Average size of the files in this type
Total Size	Total size of the files in this type
Size %	Size in % of the files in this type
Export List	Click this checkbox to launch a window where you can specify the criteria for the file types to be exported.

File Sizes

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > File Sizes

View a breakdown of the file sizes across the profiled devices. Each file size group can be exported to a CSV file, which includes details for each of the files, including owners, file type, sizes, and access information.

See “File Analytics Overview” on page 602.



Table 19-15 File Sizes

Column name	Description
File Size	This column lists ranges of file sizes. Click the link to launch the Add Export Request to specify the criteria for the file list to export.
# Files	The number of files found for this range of sizes
Files %	% of all files found that fall into this range
Total Size	Total size of all files found that fall into this range
Size %	Size % of all files found that fall into this range
Export List	Click this checkbox to launch a window where you can specify the criteria for the file size data to be exported.

Note: A discrepancy may be noted in the # of Files shown in the File List Export output and certain reports, such as the File Sizes report. This is due to rounding used in the reports. File List Export includes all filenames profiled for the Portal, while the File Size report categories are at the KB level and are impacted by precision differences caused by rounding and division by 1024.

File Types

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > File Types

The file types represented in the bar charts enable an at-a-glance insight into storage usage.

See "File Analytics Overview" on page 602.

Table 19-16 File Types

Column name	Description
File Type	Type of the discovered files, based on the filename extension
# Files	Number of files in this file type group
Average Size	Average size of the files in this category
Total Size	Total size of the files in this category
Size %	Size % of the files in this category
Export List	Click this checkbox to launch a window where you can specify the criteria for the file types to be exported.

Largest Files

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > Storage Efficiencies > Largest Files

The largest files discovered are listed in descending order to help you identify the storage consumption offenders.

See "File Analytics Overview" on page 602.

Table 19-17 Largest Files

Column name	Description
File Name	Name of the large file
File Type	Category of the file, based on the filename extension
File Size	Size of the file
Owner	File owner
Last Access	Last file access timestamp

File List Preview

The File List Preview report retrieves data by performing search on the metadata information collected from the following File Analytics policies.

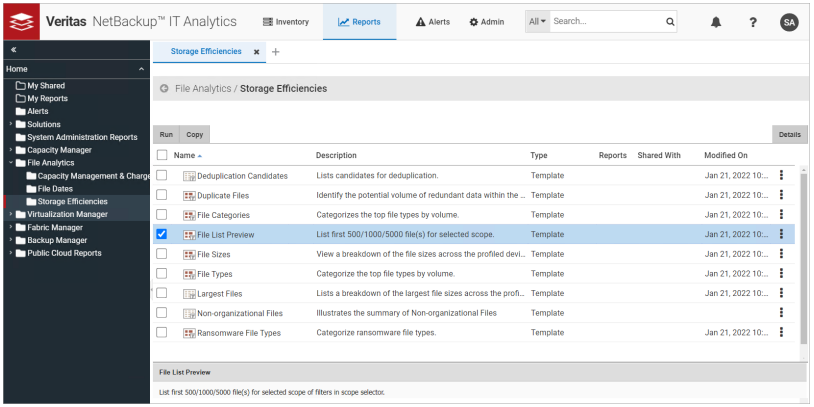
- For more information, refer **Data Collector Installation Guide for Backup Manager**, section **Configuring file analytics in NetBackup Data Collector policy**.
- For more information, refer **Data Collector Installation Guide for File Analytics**, sections **Pre-Installation setup for File Analytics** and **File Analytics Export folder size and folder depth**.

The report displays files data by executing background search.

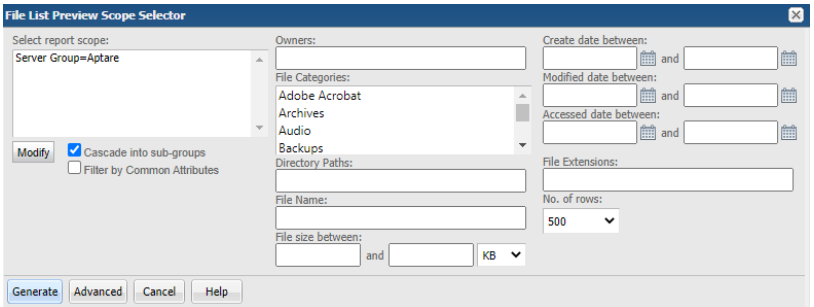
The user can configure various filter parameters based on which the data is displayed. This report is located here:

Reports > File Analytics > Storage Efficiencies > File List Preview.

Note: This report is also available in Inventory. Navigate to **Inventory**, click **File Shares & Volumes**. Click "+" icon to view the report. The report is available under **Management** section.



Click **File List Preview** to specify the report scope.



Specify the filtering criteria for the files to be investigated.

Table 19-18 File List Preview

Column name	Description
Name	User-defined name for the export file
Owners	Search only for files only for certain owners
Create date between	Select only files within the creation date range
Modified date between	Select only files within the modified date range
Accessed date between	Select only files within the accessed date range

Table 19-18 File List Preview *(continued)*

Column name	Description
Scope	Select specific host groups, servers, or Devices (hosts, shares, and volumes). When you click **Modify**, a Device search window facilitates device identification and selection. **Note:** The **Modify** button is disabled when you view the report in **Inventory**.
File Categories	Select the File Categories to be exported. Use Shift-click and Ctrl-click to select multiple categories. To add or modify the default list of file categories, go to Admin > Report > File Categories.
Directory Paths	Enter a comma-separated list of Directory Paths, where the files are located
File Extensions	In addition to the file extensions represented by the File Categories, you can augment this selection with additional file extensions (rather than creating another file category)
No. of Rows	Using this option, you can specify the number of rows displayed in the report. The available options are: **1** 500 **2** 1000 **3** 5000 **4** Custom **Note:** Using the **Custom** option, you can specify the number of rows to be displayed on the report. The allowed range is 1 to 5000.
File size between	Specify the file sizes of interest to you. Select a unit of measure: KB, MB, GB or TB
File Name	Specify the name of the file to export.

The following columns are displayed in the report.

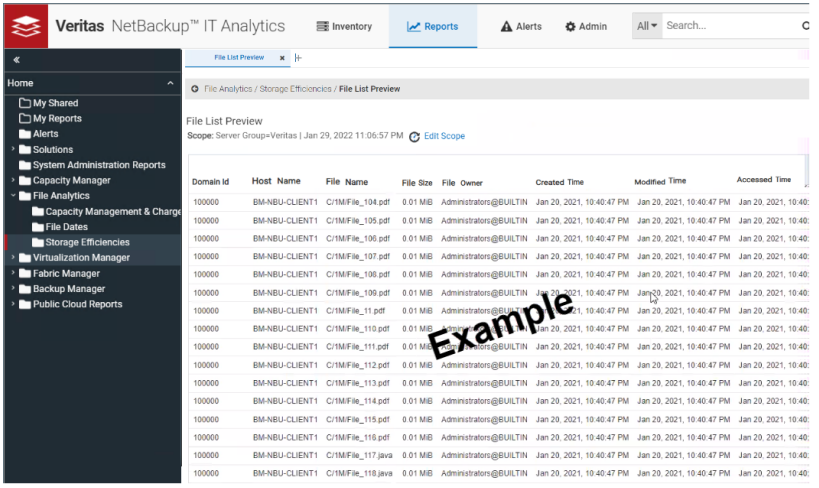
Table 19-19 File List Preview

Column name	Description
Domain ID	The domain in which the host is located.
Host Name	IP address or host name on which the file is located.
File	Name of the file

Table 19-19 File List Preview *(continued)*

Column name	Description
File Size	The size of the file as specified in the filter.
File Owner	Owner of the file.
Created Time	File creation date and time
Modified Time	File modified date and time.
Access Time	The last accessed date and time of the file.
Backup Policies	The Backup Policy applied during extraction of the data. **Note:** The data in Backup Policy column will be displayed only if data collection is done from NetBackup File Analytics feature.
Backup Time	The date and time of the back of the file conducted. **Note:** The data in Backup Time column will be displayed only if data collection is done from NetBackup File Analytics feature.

The following is the example of the report.



Once the report is generated and displayed, various actions can be performed on the displayed report. Click **Actions** on the right upper corner of the report.

The following are various actions associated with the report.

Table 19-20 File List Preview

Column name	Description
Refresh	Click **Refresh** to re-generate the report based on the filter criteria specified.
Edit Scope	Click **Edit Scope** to re-enter the new scope for the report. The File List Preview Scope Selector dialog box is displayed.
Save As	Click **Save As** to save the report as Template for future user. The Save Template as Report dialog box is displayed.
Email	If you want to send the report as an attachment using email, click **Email**. The Email Report dialog box is displayed.
Export	Click **Export** if you want to export the report in CSV file format.
Export All	If the row count exceeds 5000, click **Export All** option.
Filter	Click **Filter** if you want to filter the report data. The Report Filtering dialog box is displayed.

Access Dates

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > File Dates> Access Dates

This report provides an overview of the age of files since they were created.

See "File Analytics Overview" on page 602.

Table 19-21 Access Dates

Column name	Description
Period	Time frame, in months, for the displayed data. Future Dated/Before Create refers to files that have a time stamp that is outside of the normal time frame, typically due to programmatic manipulation; for example, 01 Jan 2030.
# Files	Number of files accessed/created/modified during this time period

Table 19-21 Access Dates *(continued)*

Column name	Description
Files %	% of all files in this report that fall within this time period
Total Size	Sum of the number of files in this time period
Size %	Size, relative to all files represented in this table

Create Dates

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > File Dates > Create Dates

See "File Analytics Overview" on page 602.

View bar charts that represent the age of groups of files since they were first allocated. Use this report to get an understanding of the trend of file creations to plan for storage requirements.

Table 19-22 Create Dates

Column name	Description
Period	Time frame, in months, for the displayed data. Future Dated/Before Create refers to files that have a time stamp that is outside of the normal time frame, typically due to programmatic manipulation; for example, 01 Jan 2030.
# Files	Number of files accessed/created/modified during this time period
Files %	% of all files in this report that fall within this time period
Total Size	Sum of the number of files in this time period
Size %	Size, relative to all files represented in this table

Modified Dates

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

File Analytics > File Dates > Modified Dates

View bar charts that illustrate the period since files were last opened for editing or modification. Use this report to determine the historical activity for files consuming space.

See “File Analytics Overview” on page 602.

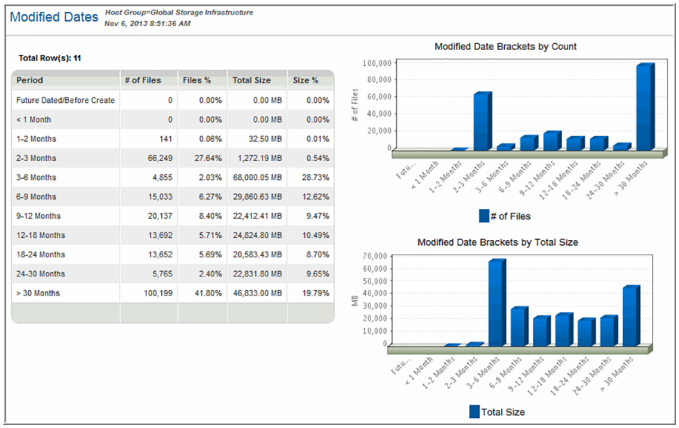


Table 19-23 Modified Dates

Column name	Description
Period	The time frame, in months, for the displayed data. Future Dated/Before Create refers to files that have a time stamp that is outside of the normal time frame, typically due to programmatic manipulation; for example, 01 Jan 2030.
# Files	Number of files accessed/created/modified during this time period
Files %	% of all files in this report that fall within this time period
Total Size	Sum of the number of files in this time period

Table 19-23 Modified Dates *(continued)*

Column name	Description
Size %	Size, relative to all files represented in this table

Chapter 20

Virtualization Manager Reports

This chapter includes the following topics:

- Virtualization Manager Reports
- Virtualization Terminology
- Virtualization Manager Overview
- Understanding the Datastore
- Using Virtualization Manager for Planning and Monitoring
- VM Cluster Summary
- VM Files Summary
- VM Guest On-Prem Right-Sizing Detail by Host
- VM Guest Disk R/W Profiles for Cloud Migration
- Data Collection Options and VM Files Summary
- VM Server Summary
- VM Server Detail
- Microsoft Hyper-V Server Summary
- VM Snapshot Summary
- VM Storage Mapping Summary
- VM Summary

- VM Detail
- Microsoft Hyper-V VM Summary
- VM Detail (Hyper-V)
- VMs Over Time
- Logical Disk Utilization
- Datastore Usage Breakdown
- VM Datastore Overview
- Datastore Usage History and Performance
- Datastore Usage History
- Datastore Disk Read/Write Performance
- Datastore Disk Latency Performance
- Datastore Utilization Summary
- Datastore Detail
- Physical Disk Utilization
- Physical Disk Detail
- Datastore Performance Summary
- Physical Disk Performance Over Time
- Physical Disk Performance Summary
- VM Guest to LUN Performance
- VM Performance Over Time
- VM Performance Summary
- Datastore Capacity Forecast
- VM Size Forecast

Virtualization Manager Reports

NetBackup IT Analytics provides insight into Virtual Machine performance and the capacity utilization of mapped datastores. Virtualization Manager offers reports on storage resource utilization at the file system and virtual disk levels for each VM.

In additional, forecasting reports predict future storage capacity needs in virtualized environments.

The data collected on Virtual Machine storage is displayed in the reports described in this section. Using these reports, you can identify unused, reclaimable storage.

Virtualization Terminology

The following terms are defined in the context of Virtualization Manager.

Note: The usage of VM refers to a Virtual Machine, not Virtualization Manager.

Table 20-1 Virtualization Terminology

Column name	Description
Capacity: Datastore	The capacity of all the datastores hosting a Virtual Machine File System (VMFS) or, for NAS storage, a Network File System (NFS).
Capacity: Logical	Also known as the Virtual Disks or Volume, this is the capacity of the virtual machine, derived from the virtual machine's VMDK file(s).
Datastore	The datastore is the container for the VMs and their configuration files, and other files, such as the ISO files that are used for installing a virtual machine. A datastore provides the virtual storage resources via mapping to physical storage on DAS SCSI, FC SAN arrays, iSCSI, or NAS drives. The datastore is the storage provisioning source for one or more VMs on one or more hosts.
VM Server	The physical server (bare metal box) that is hosting the VMs via the virtualization software, such as ESX. Virtual machines reside on a VM Server, sometimes referred to as a Virtual Host or Virtual Machine Server.
Physical (Raw) Disk	A VM can use either a virtual disk or it can access the host machine's physical disk drive: Raw Device Mapping or Raw Disk Map (RDM) file. Raw disk mapping to a LUN may have been chosen in order to optimize performance.
Virtual Disks	Also known as a Virtual Hard Drive or Virtual Machine Images (.VMDK file), this is the storage that is available to a guest operating system. It appears as a physical disk to the guest operating system, but it actually is a file that encapsulates the OS, applications, and data files of the VM. The VMDK file is accessed as if it were a physical hard disk. VMDK files can be either on the host or stored remotely.

Table 20-1 Virtualization Terminology *(continued)*

Column name	Description
VM Guest	Also known as a VM Instance; the logical server running on a VM Server. **Note:** All host names within a domain must be unique, especially when VM cloning occurs.
Virtual Host	The physical server (bare metal box) that is hosting the VMs via the virtualization software, such as ESX. Virtual machines reside on a VM Server, sometimes referred to as a Virtual Host or Virtual Machine Server.
Virtual Machine (VM)	The container for the guest operating system and applications; each VM is separate from all other VMs, even though they may share physical resources, such as memory and storage devices.
VM Server	VM Server is used throughout the reports to represent the physical server that is hosting the Virtual Machines; for example, the physical server running ESX.
VMDK	Note that for Raw Mapping (RDM), the VMDK file simply contains the metadata to map to a LUN.
VMFS	The Virtual Machine File System, with a hierarchical directory, provides the structure for managing access to shared, clustered storage. Each VM has a single subdirectory in the VMFS volume for the VMDK files (virtual disks). For raw disk mapping, the VMFS is mounted on the Virtual Machine.
Volume	The Volume is the storage that is exposed to the OS (filesystems mounted it on it). A Volume maps to logical disks, such as C:\ and D:\, as seen by the guest OS. These logical disks can be: ■ Virtual Disks, which are VMDK files that contain the actual data that is part of a datastore. ■ Raw Mapping, which is a small VMDK file that contains the metadata that maps to a LUN. **Note:** Volume Usage shown in NetBackup IT Analytics represents logical disk usage.

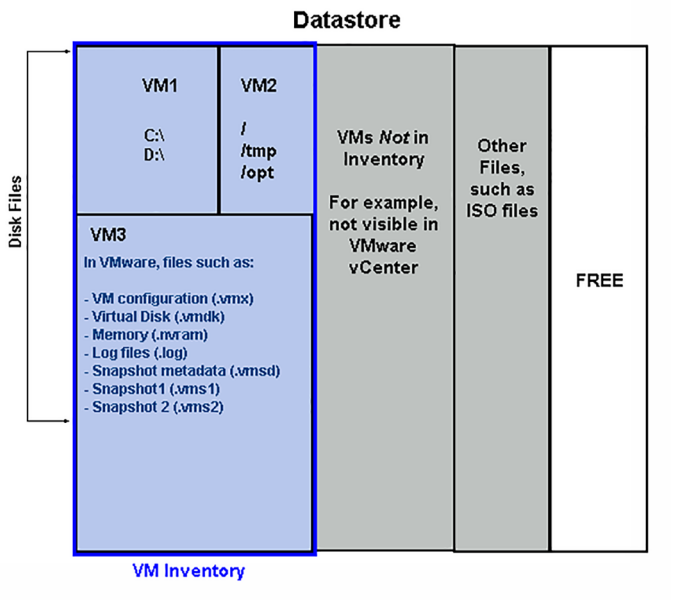
Virtualization Manager Overview

Using Virtualization Manager, you can optimize storage allocation to virtual systems to improve efficiency and reduce costs. Virtualization Manager complements your

virtual management console, with reports that can help you determine if storage resources are over- or under-utilized and if applications running in the virtual environment are contributing to storage access bottlenecks. With Virtualization Manager, you can easily view the virtual-to-physical relationships to understand what virtual systems are consuming physical storage, and more importantly, how much of that storage is consumed by out-dated, unnecessary snapshot files.

Understanding the Datastore

Since Virtualization Manager maps the VMs to the actual physical storage, it is important to understand the Datastore. The Datastore is the container for the VMs and their configuration files, VM disks, and other files, such as the ISO files that are used for installing a virtual machine. A Datastore can be located on DAS SCSI, FC SAN arrays, iSCSI, or NAS drives. The Datastore is the storage provisioning source for one or more VMs, on one or more hosts, mapping to one or more physical storage devices.



All of the objects managed by VMware vCenter comprise the VM inventory: hosts, clusters, VMs, and Networks.

Datastore Characteristics

- Configured as a clustered file system
- Configured per host, but multiple hosts can point to the same datastore
- Managed at the host level and the data center level
- Datastore path format:

  ```
  [<datastore>] <path>
  ```

  where <datastore> is the datastore name and <path> is a slash-delimited path from the root of the datastore.
  Example: **[storage] win2k/testing/bue11dtesting.vmx**

Using Virtualization Manager for Planning and Monitoring

Using Virtualization Manager, you have the tools to effectively plan the deployment of virtual machines in your enterprise and then follow through with monitoring storage allocation and usage. In concert with Virtualization Manager and Backup Manager, you have the ability to monitor all aspects of your storage environment and you can ensure that your data is sufficiently protected.

Virtualization Manager offers the following solutions:

Table 20-2 Using Virtualization Manager for Planning and Monitoring

Virtual Environment Administration	Reports for Planning & Monitoring
■ Determine optimal deployment of Virtual Machines	■ See “ VM Size Forecast ” on page 690.
■ Identify storage access bottlenecks	■ See “ VM Server Summary ” on page 639. ■ See “ Datastore Performance Summary ” on page 676.

Table 20-2 Using Virtualization Manager for Planning and Monitoring *(continued)*

Virtual Environment Administration	Reports for Planning & Monitoring
■ Reclaim space ■ Identify VM files that are taking up storage, but are not in the VM Inventory	■ See " VM Summary " on page 650. - Choose a Datastore and drill down to the details to determine file system allocations.

VM Cluster Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Cluster Summary

Lists all ESX Clusters configured in the scope, one line item per cluster with aggregated capacity utilization and High Availability (HA) metrics illustrated. Includes drilldowns to the VM Server Summary for member servers. Viewing ESX servers in clusters and logical groups allows you to monitor and manage for load balancing and fail over.

Table 20-3 VM Cluster Summary

Column name	Description
Name	Cluster name.
Status	Overall status of the cluster. White circle = Status Unknown, Green check mark = Normal, Yellow triangle = Warning, Red X= Alert, definite problem. When filtering the report on this field, use the color.
Swap Policy	VM swap policy name.
VM Servers	
# Total	Total number of VM servers.

Table 20-3 VM Cluster Summary *(continued)*

Column name	Description
# Active	Total number of active VM servers.
# Inactive	Total number of inactive VM servers.
VM Guest	
# Total	Total number of VM guests in the cluster. Drill down to the VM Summary for these VM guests.
# Active	Number of VM guests in the cluster that have successfully been polled in the last 24 hours. Drill down to the VM Summary for these guests.
# Inactive	Number of inactive VM guests in the cluster (not successfully collected from the last 24 hours). Drill down to the VM Summary for these Guests.
Datastore	
Total Capacity	Total capacity of the datastore.
Used Capacity	Used capacity of the datastore.
Snapshots	Sum of all snapshots' reserved capacity for all VM servers in the cluster.
Committed %	Thermometer (green/yellow/red) of (used+snap)/total.
High Availability (HA)	
Enabled	Yes or blank for high availability enabled.
AC Enabled	Yes or blank for high availability admission control enabled.
Level	High availability fail over level.
Priority	High availability restart priority.
Isolation Response	High availability isolation response.
# Slots	Total number of high availability slots.
# Used Slots	Number of used high availability slots
# Free Slots	Number of free high availability slots.
Slot CPU Speed (MHz)	Speed of CPU slot.
Slot Memory	Size of memory slot.
# VCPU Slots	Number of virtual CPU slots.

Table 20-3 VM Cluster Summary *(continued)*

Column name	Description
DRS (Distributed Resource Scheduler)	
Enabled	Yes or blank.
Mode	Mode for cluster.
Automation Level	Automation level for cluster.

VM Files Summary

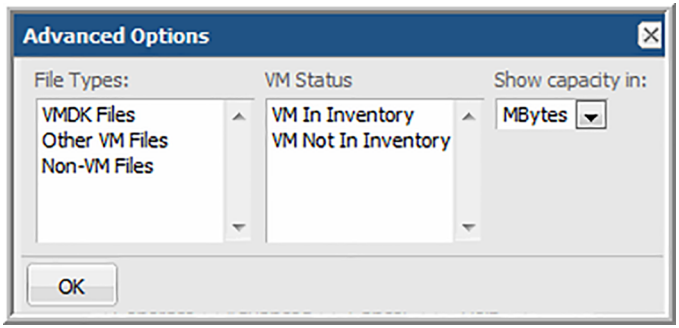
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Files Summary

List all VM files, with associated datastores and LUNs. Use this report to identify space that could be reclaimed.

1. In the Scope Selector, click **Modify** to select the objects for report filtering: Hosts, VM Servers, or Datastores. Note that these are mutually exclusive. For example, if you choose to report by Datastores, you cannot include Hosts in the scope.
2. In the Scope Selector, click **Advanced** to select from the Advanced Options, which enable further filtering.



Note that for performance reasons, the NetBackup IT Analytics System Administrator may have configured an Advanced Parameter (VMWARE_LIST_DATASTORES). If this parameter is set to N, the inventory collection will not list all files in datastores, but only record information about files that are directly associated with virtual machines that are collected (e.g., virtual disk files, log files, configuration files). In this case, an entry (row) in the report output may appear without a Last Modified date. This is because these files are not listed in the datastore, however, the data collector finds them when processing the VM files.

Table 20-4 VM Files Summary

Column name	Description
File Name	Name of the file.
File Path	The file path in datastore format; for example: `[storage] win2k/testing/bue11dtesting.vmx`
Datastore	The datastore in which this file resides.
VM Server	The host that is running the VM software (ESX).
VM Name	Name of the virtual machine.
VM Host Name	The server that is hosting the VMs.
VM Status	Circle = Status Unknown, Green check mark = Normal, Yellow triangle = Warning, Red X = Alert, definite problem
Provision Type	Indicates if the file is thin or thick provisioned.
File Type	File type identifier: VmLog - log for VMware workstation activity (.log) VmNvram - VM BIOS state (.nvrm) VmDisk - contents of the VM's hard drive (.vmdk) VmSnapshot - metadata about snapshots (.vmsd) VmConfigFile - primary configuration file (.vmx)
File Used	Size of the Virtual Machine image (.vmdk file) containing data.
Maximum File Size	Capacity of the virtual disk from the point of view of a virtual machine.
LUN Name	The LUN on which the file resides.
Array	The array on which the file resides.

Table 20-4 VM Files Summary *(continued)*

Column name	Description
Last Modified	Date and time stamp of last file modification. See "Data Collection Options and VM Files Summary" on page 639. for data collection details that impact the Last Modified date.

VM Guest On-Prem Right-Sizing Detail by Host

Use this report to analyze the historical usage patterns of both memory and CPU over a period of time. From this data, determine the maximum amount of memory and CPU used over that time frame. Compare the maximum used with the actual configured memory and apply the delta to determine the right size of the VM.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Guest On-Prem Right-Sizing Detail by Host

Table 20-5 VM Guest On-Prem Right-Sizing Detail by Host

Column name	Description
VM Guest	It is the logical server running on a VM server, also known as a VM instance; **Note:** All host names within a domain must be unique, especially when the VM cloning occurs
#Hours Running	The amount of time the VM was powered on during the specified time frame.
Configured Memory	The amount of virtual memory that is configured on the VM.
Configured #vCPU	The number of vCPUs configured on this VM.
Max Mem Used%	The highest percentage of memory used on a single day over the given time frame.

Table 20-5 VM Guest On-Prem Right-Sizing Detail by Host *(continued)*

Column name	Description
Avg Mem Usage%	The average percentage of memory used per day, over the given time frame.
Max CPU %	The highest percentage of CPU used on a single day used over the given time frame.
Avg CPU %	The average percentage of CPU used per day used over the given time frame.
Adjusted Required Mem	The expected ideal memory usage based on the highest single day memory usage.
Adjusted Required vCPU	The expected CPU usage based on the highest single-day CPU usage.
Over Mem	The amount of over allocated memory (Configured - Max memory), based on the highest single-day memory usage.
Over CPU	The amount of CPU over-allocation (configured #vCPU - Max vCPU) based on the highest single day CPU usage.
Guest Type Profile	A profile name is derived by the #vCPUs and memory. The Type Profile can be used to group and filter VM's of similar characteristics. Example: VMWare_2C_2G will be a VM with 2 vCPU and 2 GB of RAM. Similar to naming conventions used by public cloud providers that is Standard_A3, m4.xlarge, and so on.

VM Guest Disk R/W Profiles for Cloud Migration

This report uses historical disk read/write patterns in order to categorize VMs into Low, Medium, and High in relation to their peers. This will help determine the right size and type of cloud storage that will be required for migration.

Use the **Reports** tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Guest Disk R/W Profiles for Cloud Migration

Table 20-6 VM Guest Disk R/W Profiles for Cloud Migration

Column name	Description
VM Guest	Also known as a VM Instance, it is the logical server running on a VM Server. **Note:** All host names within a domain must be unique, especially when the VM cloning occurs.
Power State	Current power state of the VM as of last collection. Possible value is either PoweredOn or PoweredOff.
Sum R/W	The sum total of all the reads and writes combined over the given time frame.
Avg R/W	The average of all the reads and writes combined over the given time frame.
Peak R/W	The most reads and writes in a single day over the given time frame.
R/W Percentile	To what percentile do the reads and writes rank when compared to their peers.
R/W Usage	The ranking of the reads and writes based on their percentile. Possible values are Low, Medium, or High.
R/W % of Total	What percentage of reads and writes is the VM doing compared to the total of all read and writes.
Capacity	Size of storage that is allocated to this VM.
Arrays	The name of the storage array or arrays attached to the VM.

Table 20-6 VM Guest Disk R/W Profiles for Cloud Migration *(continued)*

Column name	Description
#LUNs	The number of storage logical units allocated to this VM from the storage array(s).
LUNs	The name(s) of the logical units allocated to this VM from the storage array(s).

Data Collection Options and VM Files Summary

If this report is accessed via a drilldown from another report, the values shown represent the type of data collection that was configured to collect this data.

A Virtualization Manager data collection policy can be configured to gather data as follows:

- Inventory (default option): Collect only the data associated with the VMs in the inventory.
- Collect all files in datastore: Perform a full scan of the datastore to discover files that are resident on the datastore, but not associated with a VM in the inventory. This option finds files that may be invisible to VMware, but that are consuming space on the datastore.

When the Inventory option is configured in the data collector policy, it has the following impact on fields shown in this report:

- The Last Modified date will be null.
- Total # Files collected usually will be less than the all files collection.

When the Collect all files in datastore option is configured in the data collector policy, it has the following impact on fields shown in this report:

- The Last Modified date will be missing for files not found in the datastore scan.

VM Server Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Server Summary

Use this report to view the list of Virtual Machine Servers, with access to reports that provide details to help you determine what is using the space on a datastore. This report lists ESX servers, one line item for each physical server, with usage and status information.

The VM Server Summary lists the VM Servers within your selected scope; that is, the physical servers that are hosting the VMs.

When running this report, the following mutually exclusive scope selections are available:

- Host Group
- Datastore

Note: When using advanced filtering for tabular reports to filter rows by the Indicator Lights (for example, Status or Hardware Health) use the color: Red, Yellow, Green, or White.

Table 20-7 VM Server Summary

Column name	Description
Cluster Name	Name of the cluster with which the VM server is associated.
Name	Name of the VM Server that is hosting the virtual server; for example, the ESX host
VM Tags	Displays VMWare tags and custom attributes associated with ESX Hosts and VMs. Tags refer to metadata, such as location or OS, that are assigned to objects to categorize them.
IP Address	IP address of the VM Server
Data Center	The data center in which the server resides.
State	VM connection state: Connected, Not Responding, Disconnected, None
Status	White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem. When filtering the report on this field, use the color.

Table 20-7 VM Server Summary *(continued)*

Column name	Description
Hardware Health	Indicates the overall health of the host's hardware: None, Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert. When filtering the report on this field, use the color.
# CPUs	Total number of processors in the virtual machine
Avg CPU Usage Last 24 Hrs	The host's average CPU usage for the past 24 hours for the number of samples taken within that time frame
Additional VMs	Based on CPU resources, this value helps to determine how many more VMs can be supported on the VM server.
Total Memory	Total memory capacity of the VM Server.
Avg Mem Usage Last 24 Hrs	The VM Server's average memory usage for the past 24 hours for the number of samples taken within that time frame
Additional VMs	Based on memory resources, this value helps to determine how many more VMs can be supported on the VM server.
# Disks	The number of disks that the VM Server can access, either internal disks or LUNs. See " Datastore Utilization Summary " on page 666. Note that the disks are specific to the VM server, whereas the Datastore utilization may be using both VM Server disks and disks apart from the VM Server disks.
Unallocated	Amount of disk space that is unallocated. Mouse over this thermometer to view the total, unallocated, and percentage unallocated. Click on this thermometer to view the utilization summary. See " Datastore Utilization Summary " on page 666.
# Datastores	The number of datastores provisioned to the VM Server. See " Datastore Usage Breakdown " on page 662.
Datastore Usage	Mouse over the thermometer bar to view the total datastore capacity, the amount used, and the % of the total used. See " Datastore Capacity Forecast " on page 689.
# Active VMs	A VM is active when it is powered on. An inactive VM is in the inventory, but is powered off. This column shows the number of active VMs. See " VM Summary " on page 650.

Table 20-7 VM Server Summary *(continued)*

Column name	Description
# Total VMs	The total number of VMs in the inventory, both active and inactive. See " VM Summary " on page 650.

VM Server Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Virtualization Manager > Administration Reports > VM Server Summary

Click the **Name** link.

Note: VMware Tools (VM Tools) must be installed to enable collection of the IP address, host name, mount points, disk path, available space on VM guest volumes, and guest operating system of the VM.

VM Server Detail Data

Table 20-8 VM Server Detail

Column name	Description
Name	Name of the VM Server that is hosting the virtual server; for example, the ESX host
Host Name	Hostname of the guest; value is displayed only if VMware Tools is installed; links to Host Details
IP Address	IP address of the Virtual Machine; value is displayed only if VMware Tools is installed

Table 20-8 VM Server Detail *(continued)*

Column name	Description
Data Center	The data center in which the server resides.
State	Network adapter connection state: Connected, Not Responding, Disconnected, None
Status	White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, a definite problem
Hardware Health	Indicates the overall health of the host's hardware: None, Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert
System Vendor	Host vendor
System Model	Model of the host system
Total Memory	Total memory available as a resource to the VMs on this host
Total CPU	Total CPU available as a resource to the VMs on this host
CPU Vendor	Processor's vendor
CPU Description	Type of processor
CPU Speed	Processor speed
# HBAs/Ports	Number of host bus adapters and ports on the VM Server
# NICs	Number of physical network interface cards
NIC Speed	Speed of the interface cards.
Cluster Name	Name of the cluster associated with the VM Server.
# Total VMs	Number of VMs configured on this host and in the inventory.
# Active VMs	A VM is active when it is in the inventory and powered on; it is inactive if it is in the inventory, but powered off.
# Datastores	Number of datastores that this VM Server is using.
Datastore Usage	Amount of storage in the datastore that this VM Server is using.
# Disks	The number of disks that the VM Server can access, either internal disks or LUNs. See " Datastore Utilization Summary " on page 666.

Table 20-8 VM Server Detail *(continued)*

Column name	Description
Unallocated Disk	Mouse over the thermometer bar to view the total disk capacity, the amount used, and the % of the total used. See " Datastore Utilization Summary " on page 666.
Version	Version of the virtual server software, such as ESX.
Last Updated	Timestamp of last data collection.

VM Tags and Custom Attribute Details

This section displays the VM Tags/Custom Attributes and their values.

Table 20-9 VM Server Detail

Column name	Description
VM Name	Display name of the virtual machine.
Tag Name	Displays VMWare tags and custom attributes associated with ESX Hosts and VMs. Tags refer to metadata, such as location or OS, that are assigned to objects to categorize them.
Tag Type	Identifies the type associated with the tag. Tag is displayed if it is VM Tag, or Custom is displayed to identify a custom attribute.
Tag Value	Displays VMWare tags and custom attribute values associated with ESX Hosts and VMs. Tag values may appear as blank if no value is available.

Performance & Forecast

Table 20-10 VM Server Detail

Column name	Description
CPU	Average and Maximum usage graphed over time for the number of samples taken within that time frame.

Largest VMs

Click the link to access the full list of the largest VMs, listed in descending order, with the top storage-consuming VMs at the top of the list.

Note: VMware Tools (VM Tools) must be installed to enable collection of the IP address, host name, mount points, disk path, available space on VM guest volumes, and guest operating system of the VM.

See " VM Summary " on page 650.

Datastores

For a description of the fields in this table refer to the following.

See " Datastore Usage Breakdown " on page 662.

Disks

These are the physical disks that are available to the VM Server.

For a description of the fields in this table, refer to the following.

See " Datastore Utilization Summary " on page 666.

Host Details by Source

Hosts can be collected by any number of sources such as vendor products. This table lists the details that have been collected by a vendor product so that you can validate the collected details.

Table 20-11 VM Server Detail

Column name	Description
Display Name	Name that is displayed in reports.
Host Name	Name of the host as it was collected by the vendor product.
IP Address	IP address of the host.
Product Group	The product group, such as Capacity, for which the host data was collected.
Source	A specific source such as vendor product (subsystem) from which host data was collected; for example, VMware.
Date Created	The date and time the host was created in the Portal database.
Last Updated	The date and time the host data was updated in the Portal database.
OS Platform	Operating system of the host, as collected from this vendor product; for example, Windows Server 2008 R2.

Table 20-11 VM Server Detail *(continued)*

Column name	Description
OS Version	The host's specific operating system version, as collected from this vendor product; for example, 6.01.
Make	The host's make, as collected from this vendor product; for example,Windows-x64.
Model	The host's model, as collected from this vendor product; for example, PowerEdge 2950.
Backup Server Type	The host's role in a backup relationship, client or server.

Data Collection Options and VM Server Detail

If this report is accessed via a drilldown from another report, the values shown represent the type of data collection that was configured to collect this data.

A Virtualization Manager data collection policy can be configured to gather data as follows:

- **Inventory** (default option): Collect only the data associated with the VMs in the inventory.
- **Collect all files in datastore**: Perform a full scan of the datastore to discover files that are resident on the datastore, but not associated with a VM in the inventory. This option finds files that may be invisible to VMware, but that are consuming space on the datastore.

When the **Inventory** option is configured in the data collector policy, it has the following impact on fields shown in this report:

- VM Not In Inventory will be zero.
- Total VM Used will be reduced by the amount that would have been collected as VM Not In Inventory if an all files collection had been done.

When the **Collect all files in datastore**option is configured in the data collector policy, it has the following impact on fields shown in this report:

- The Last Modified date will be missing for files not found in the datastore scan.

Microsoft Hyper-V Server Summary

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points.

This report lists Hyper-V virtual machines - one line for each physical server with usage and status. Drill down to view details including file systems, mount points and LUNs.

Table 20-12 Microsoft Hyper-V Server Summary

Column name	Description
Host Name	Hostname of the Hyper-V server. Click to drill down to the Host Overview.
Name	Display name of the Hyper-V server.
IP Address	IP address of the Hyper-V Server.
Status	Indicator that reflects the condition of the server: White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
System Vendor	Host vendor.
System Model	Model of the host system.
Total Memory	Total memory capacity of the Hyper-V Server.
#CPUs	Total number of processors in the virtual machine.
CPU Vendor	Vendor for the CPU.
CPU Description	Type of processor.
CPU Speed mhz	Processor speed in mhz.
#HBAs	Number of host bus adapters on the Hyper-V Server
#NICs	Number of physical network interface cards.
NIC's Speed (Mbps)	Speed of the interface cards in Mbps.
Cluster Name	Name of the cluster with which the Hyper-V server is associated.
#VMs	Number of VMs configured on this host and in the inventory.
# Active VMs	A VM is active when it is powered on. An inactive VM is in the inventory, but is powered off. This column shows the number of active VMs.
#File System	Total number of Filesystems associated with the Server.
OS Version	The host's specific operating system version, as collected from this vendor product; for example, 6.01.

Table 20-12 Microsoft Hyper-V Server Summary *(continued)*

Column name	Description
Avg CPU Usage	The host's average CPU usage for the past 24 hours for the number of samples taken within that time frame.
Total CPU Mhz	Aggregated CPU resources of all servers, in MHz.
Avg Memory Usage %	The Hyper-Vserver average memory usage for the past 24 hours for the number of samples taken within that time frame.
Total Memory Size	Total memory capacity of the Hyper-V Server.
Storage File System	Used/Total capacity for associated filesystem(s).
Last Updated	The date and time the data was updated in the Portal database.

VM Snapshot Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Snapshot Summary

This report shows all active snapshots by VM, datastore and ESX server. Identifies capacity associated with each snapshot, when the snapshot was created, by guest, ESX server and datastore. Use this report to identify out-of-date snapshots.

Table 20-13 VM Snapshot Summary

Column name	Description
File Name	Name of the snapshot file.
Snapshot Name	Name assigned to the snapshot.
Size	Size of the snapshot.
Last Modified	Timestamp of the last write access. For details about data collection options that impact the Last Modified date, See "Data Collection Options and VM Snapshot Summary" on page 649.

Table 20-13 VM Snapshot Summary *(continued)*

Column name	Description
Datastore Name	The datastore related to the snapshot.
VM Host Name	The name of the server hosting the VMs. Click the name to display the VM Server Detail report.
VM Name	Name of the virtual machine. Click the name to display the VM Detail report.
Create Date	Date snapshot was created.

Data Collection Options and VM Snapshot Summary

If this report is accessed via a drilldown from another report, the values shown represent the type of data collection that was configured to collect this data.

A Virtualization Manager data collection policy can be configured to gather data as follows:

- Inventory (default option): Collect only the data associated with the VMs in the inventory.
- Collect all files in datastore: Perform a full scan of the datastore to discover files that are resident on the datastore, but not associated with a VM in the inventory. This option finds files that may be invisible to VMware, but that are consuming space on the datastore.

When the Inventory option is configured in the data collector policy, it has the following impact on fields shown in this report:

- The Last Modified date will be null.

VM Storage Mapping Summary

This report provides the mapping from the VMs to the physical disks. Use this report to see the relationships among the ESX servers, datastores, and physical storage.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Storage Mapping Summary

Table 20-14 VM Storage Mapping Summary

Column name	Description
VM Name	Click the VM Name link to view VM Details.
VM Server	This is the server on which the VM resides. Click the link to view the VM Server Details.
Datastores	Indicates which datastore contains the VM. Click this link to view the Datastore Utilization Summary.
# Disks	Click this link to view the Physical Disk Utilization details.
# Arrays	This indicates the number of arrays on which this VM has data. Click this numeric link to view the array specifics.

VM Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > VM Summary

Use this report to find the largest Virtual Machines. This report lists virtual guests/hosts with usage and status information.

For Virtualization Manager data collection, VMware Tools (VM Tools) must be installed to enable collection of key properties of a VM Guest, such as the IP address, host name, mount points, disk path, available space on VM guest volumes, and guest operating system of the VM. Whenever data collection cannot retrieve a host name, a VM Guest will not be treated as a host in the Inventory and Virtualization Manager reports will not be populated with host details. For example, a host name may not be available in the following situations: the VM may be down, VM Tools may not be installed on the VM Guest, or a VM template may have been collected.

Note: When using advanced filtering for tabular reports, to filter rows by the Indicator Lights (for example, Connection, Power, or Status) use the color: Red, Yellow, Green, or White.

Table 20-15 VM Summary

Column name	Description
VM Name	Name of the Virtual Machine. See " VM Detail " on page 652.
VM Tags	Displays VMWare tags and custom attributes associated with ESX Hosts and VMs. Tags refer to metadata, such as location or OS, that are assigned to objects to categorize them.
VM Server	The server that is hosting the VM.
Data Center	The data center in which the server resides.
Host Name	Hostname of the guest; value is displayed only if VMware Tools is installed.
IP Address	IP address of the Virtual Machine; value is displayed only if VMware Tools is installed.
State: Connection	None, Connected, or Suspended.
State: Power	Powered Off, Powered On.
Status	This is essentially an alarm indicator that reflects the condition of the resources: White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
Avg CPU Usage Last 24 Hours	The host's average CPU usage for the past 24 hours for the number of samples taken within that time frame.
Avg Memory Usage Last 24 Hrs	The host's average Memory usage for the past 24 hours for the number of samples taken within that time frame.
# of CPU	Total number of processors in the virtual machine
Total Memory	Total memory capacity of the VM.
Size: VM	Size of the virtual machine, which includes VMDK files, log files, and snapshots; the sum of all the files taking up storage by this VM.
Size: VM Disk	Amount of storage that was provisioned to the Guest OS when the virtual disk was originally created for the VM. For thin storage, this can be larger than the Datastore capacity associated with the VM.

Table 20-15 VM Summary *(continued)*

Column name	Description
VM Size (KB)	The sparkline chart provides a visual depiction of the history for the total size of all the VM's datastore files (in KB) to illustrate how much actual space on disk is occupied by the VM. The data points represent the maximum size for a week for the last 60 days, grouped by week. Note, the last data point on the sparkline always reflects the most recent week - even if it is a partial week.
Storage: Volume Usage	This is the disk usage inside of the VM (for example, C:\, D:\)A Volume maps to logical disks, such as C:\ and D:\, as seen by the guest OS; usage for the past 24 hours for the number of samples taken within that time frame. For inactive VMs, the usage is unknown.
Storage: Datastores	Datastores used by this VM.
Storage: Raw Device Map	Indicates if the VM is using Raw Device Mapping.

VM Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Virtualization Manager > Administration Reports > VM Summary

Click the **VM Name** link.

For Virtualization Manager data collection, VMware Tools (VM Tools) must be installed to enable collection of key properties of a VM Guest, such as the IP address, host name, mount points, disk path, available space on VM guest volumes, and guest operating system of the VM. Whenever data collection cannot retrieve a host name, a VM Guest will not be treated as a host in the Inventory and Virtualization Manager reports will not be populated with host details. For example,

a host name may not be available in the following situations: the VM may be down, VM Tools may not be installed on the VM Guest, or a VM template may have been collected.

VM Detail

Table 20-16 VM Detail

Column name	Description
Name	Name of the virtual machine.
Host Name	Hostname of the guest; value is displayed only if VMware Tools is installed.
IP Address	IP address of the Virtual Machine; value is displayed only if VMware Tools is installed.
VM Server	Links to the VM Server Detail report.
Cluster Name	The cluster name associated with the VM Server.
VM Size	Amount of storage the VM takes up.
VMware Tools	Indicates if the VMware tools have been installed for the guest operating system.
Tools Version	Version of the VMware tools that are running on the guest operating system.
Connection State	None, Powered Off, Powered On, or Suspended.
Power State	None, Powered Off, Powered On, or Suspended.
Status	White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
Boot Time	The VM's power-on time.
Suspend Time	The time when the VM was put in a suspended state.
Datastore	The datastore to which the VM belongs. See " Datastore Detail " on page 669.
Resource Pool	The resource pool (compartment) to which the VM belongs.
Last Update Date	The last time the VM configuration was updated.
Current Snapshot	Path and name of the current snapshot.

Table 20-16 VM Detail *(continued)*

Column name	Description
VM Path Name	The file path in datastore format, for example: `[storage] win2k/testing/bue11dtesting.vmx`
System OS	Guest operating system running on the VM.
Notes	The annotation field of the VM.

VM Tags and Custom Attribute Details

This section displays the VM Tags/Custom Attributes and their values.

Table 20-17 VM Detail

Column name	Description
VM Name	Display name of the virtual machine.
Tag Name	Displays VMWare tags and custom attributes associated with ESX Hosts and VMs. Tags refer to metadata, such as location or OS, that are assigned to objects to categorize them.
Tag Type	Identifies the type associated with the tag. Tag is displayed if it is a VM Tag, or Custom is displayed to identify a custom attribute.
Tag Value	Displays VMWare tags and custom attribute values associated with ESX Hosts and VMs. Tag values may appear as blank if no value is available.

Performance Over Time

Note: The default NetBackup IT Analytics polling interval is set to every 5 hours, however, this can be customized.

Table 20-18 VM Detail

Column name	Description
CPU	The CPU average and maximum usage to date.
Memory	The Memory average and maximum usage to date.

VM Size Over Time

Displays the average and maximum VM storage usage to date.

Volumes

This table lists the local disks seen by the virtual machine, such as C:\, D:\.

For a description of the fields in this table, see the following.

See “ Logical Disk Utilization ” on page 661.

Disks

This table provides the details for the physical disks in the VM.

VM Files

This table provides the details for the files in the VM.

See “ VM Files Summary ” on page 634.

Table 20-19 VM Detail

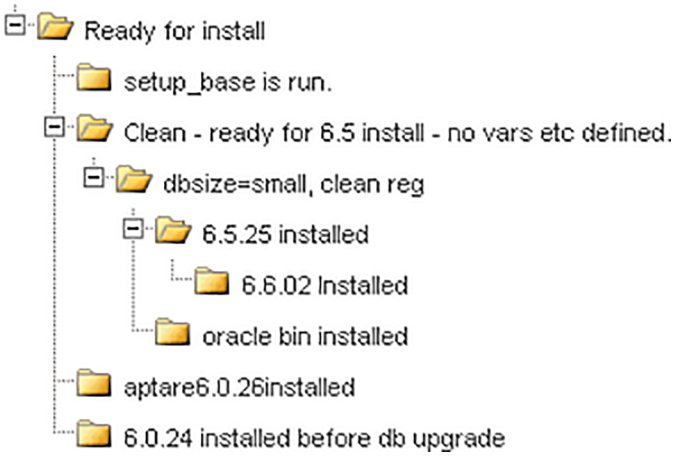
Column name	Description
Disk Label	The label of the physical disk.
Disk Type	RDISK = raw, physical disk; Capacity will be N/A for RDISKsVDISK = virtual disk
Capacity	The storage potential of the virtual disk.
Datastore	The datastore to which the virtual disk belongs.
LUN	The LUN on which the virtual disk’s storage is located.
Array	Array in which the disk resides.
Mode	1 = Persistent - immediate, permanent changes 2 = Non-persistent - changes do not persist once VM is powered off 3 = Undoable - user determines how and when changes are saved 4 = Append - changes appended to a log when VM is powered off

Snapshots

A tree structure represents the Snapshots for this VM.

Example:

Snapshots:



Host Details by Source

Hosts can be collected by any number of sources such as vendor products. This table lists the details that have been collected by source so that you can validate the collected details.

Table 20-20 VM Detail

Column name	Description
Display Name	Name that is displayed in reports.
Host Name	Name of the host as it was collected by the vendor product.
IP Address	IP address of the host.
Product Group	The product group, such as Capacity, for which the host data was collected.
Source	A specific source such as vendor product from which host data was collected; for example, VMware.
Date Created	The date and time the host was created in the Portal database.
Last Updated	The date and time the host data was updated in the Portal database.
OS Platform	Operating system of the host, as collected from this vendor product; for example, Windows Server 2008 R2.

Table 20-20 VM Detail *(continued)*

Column name	Description
OS Version	The host's specific operating system version, as collected from this vendor product; for example, 6.01.
Make	The host's make, as collected from this vendor product; for example,Windows-x64.
Model	The host's model, as collected from this vendor product; for example, PowerEdge 2950.
Backup Server Type	The host's role in a backup relationship, client or server.

Microsoft Hyper-V VM Summary

Use this report to find the largest Virtual Machines. This report lists virtual guests/hosts with usage and status information.

Table 20-21 Microsoft Hyper-V VM Summary

Column name	Description
VM Name	Name of the Virtual Machine.
VM Server	The server that is hosting the VM.
Cluster Name	The cluster name associated with the VM Server.
Host Name	Hostname of the guest; value is displayed only if Host Resource collection has been run.
IP Address	IP address of the Virtual Machine; value is displayed only if Host Resource collection has been run.
Connection	None, Connected, or Suspended.
Health	This is essentially an alarm indicator that reflects the condition of the resources: White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
Status	This is essentially an alarm indicator that reflects the condition of the resources: White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
Disk Provisioned	Physical disk space being used by the VM.

Table 20-21 Microsoft Hyper-V VM Summary *(continued)*

Column name	Description
Allocated	Allocated virtual disk space that the VM recognizes.
Memory Usage	Virtual memory allocated to the VM.
Storage File System	Used/Total capacity for associated filesystem(s).
Mount Point	Mount point on the VM Server.
LUN	If the disk space has been allocated from an array, this is the LUN name that space has been allocated from. Value is displayed only if Host Resource and Array collection has been run.
Last Updated	Last collection time.

VM Detail (Hyper-V)

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal. For Virtualization Manager data collection, VMware Tools (VM Tools) must be installed to enable collection of key properties of a VM Guest, such as the IP address, host name, mount points, disk path, available space on VM guest volumes, and guest operating system of the VM. Whenever data collection cannot retrieve a host name, a VM Guest will not be treated as a host in the Inventory and Virtualization Manager reports will not be populated with host details. For example, a host name may not be available in the following situations: the VM may be down, VM Tools may not be installed on the VM Guest, or a VM template may have been collected.

Table 20-22 VM Detail (Hyper-V)

Column name	Description
Name	Display name of the virtual machine.
VM Name	Name of the virtual machine.

Table 20-22 VM Detail (Hyper-V) *(continued)*

Column name	Description
Host Name	Hostname of the guest; value is displayed only if VMware Tools is installed.
IP Address	IP address of the Virtual Machine; value is displayed only if VMware Tools is installed.
VM Server	Links to the VM Server Detail report.
Cluster Name	The cluster name associated with the VM Server.
Enabled	Powered Off, Powered On.
Health	This is essentially an alarm indicator that reflects the condition of the resources: White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
Status	White circle=Status Unknown, Green check mark=Normal, Yellow triangle=Warning, Red X=Alert, definite problem.
Disk Provisioned	Physical disk space being used by the VM.
Allocated Size	Allocated size to Virtual Hard Disk in Kilobytes.
Memory Usage	Amount of storage the VM takes up.
File System Mount Point	The host's file system mount point for this LUN. It's possible to have zero mount points and file systems, in the scenario where a LUN has been allocated to the host, but the host has not yet made use of it by creating a file system and mounting it.
LUN Name	If the disk space has been allocated from an array, this is the LUN name that space has been allocated from. Value is displayed only if Host Resource and Array collection has been run.
Last Updated	The last time the VM configuration was updated.
OS Version	Guest operating system running on the VM.

Table 20-23 VM Detail (Hyper-V)

Column name	Description
Path Name	Defines the path for the virtual hard disk image files.
FS Type	The type of the file system.
File Size	Actual size on the virtual disk.

Table 20-23 VM Detail (Hyper-V) *(continued)*

Column name	Description
Type	Identifies the mechanism used when provisioning the storage for a virtual hard disk: Fixed is all up front (such as thick provisioning), Dynamic is as needed, (such as thin provisioning), and Differencing is used to show only the changes since a snapshot.
Format	Format for the virtual hard disk image files. Possible values: VHD, VHDX, VHDSet.

VMs Over Time

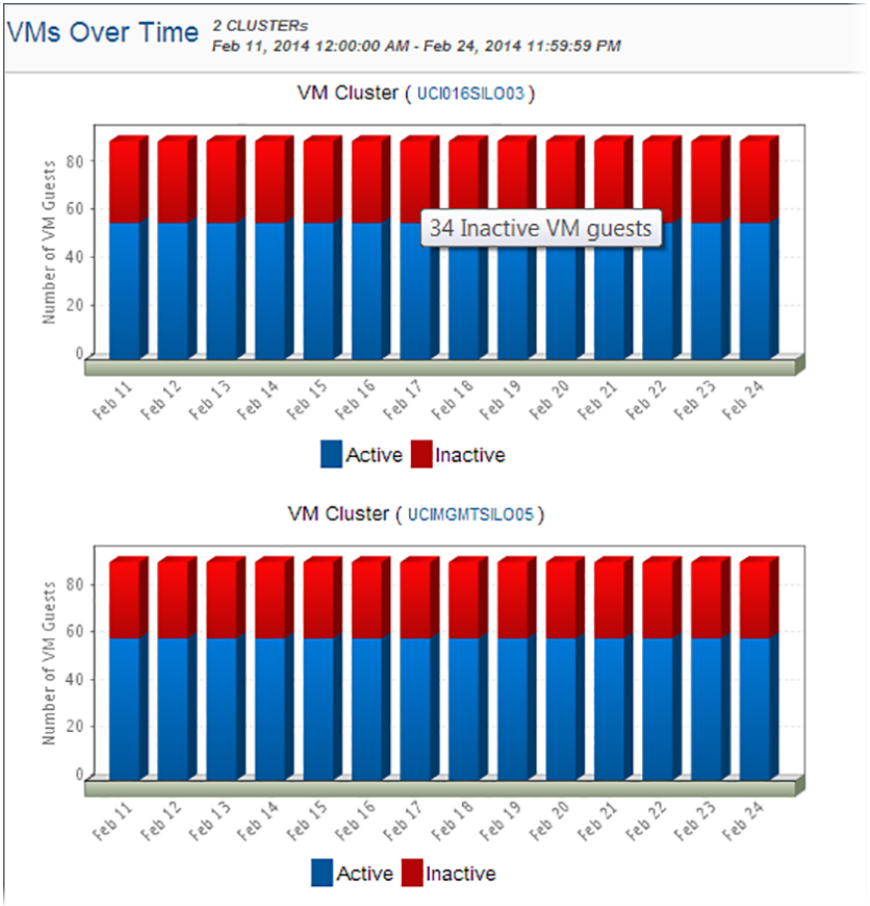
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Forecast & Planning > VMs Over Time

Graphical presentation of VM guests active and inactive over time, grouped by Cluster or Server Group. Offers a historical perspective on the number of VM Guests is displayed by the ESX Cluster or Host Group. You can drilldown to the VM Server Detail for servers and VM Cluster Summary report for clusters. Inactive is defined as a VM Guest/Host that is not powered on.

The screen shot displays a report group by clusters in the scope selector.



Logical Disk Utilization

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Virtualization Manager > Administration Reports > VM Summary

Click the **Volume Usage** thermometer.

Table 20-24 Logical Disk Utilization

Column name	Description
Disk Path	The path to the logical disk.
VM Name	Virtual machine name.
Total Capacity	Total amount of storage allocated to the VM.
Used	Amount of the VM's allocated storage that is in use.
Free	Amount of the VM's allocated storage that is available for use.
Usage	A thermometer represents the percentage of disk usage.

Datastore Usage Breakdown

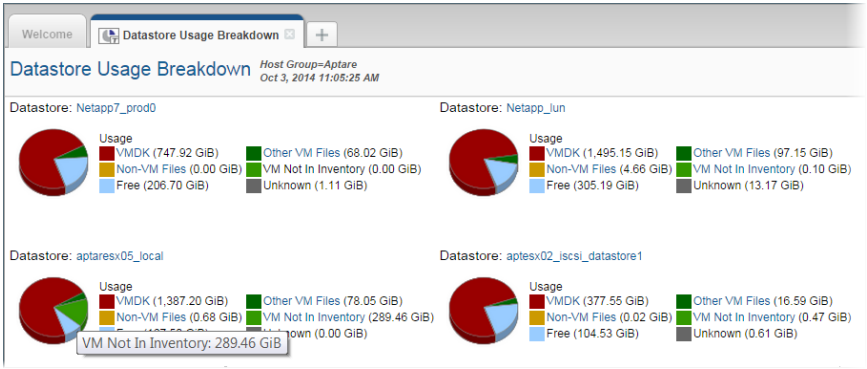
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager> Storage Capacity & Utilization > Datastore Usage Breakdown

Use this report to view a graphical representation of the file types comprising each datastore. Refer to the following for a diagram and description of the types of files that are maintained in a Datastore.

See "Virtualization Manager Overview" on page 629.



Note: The usage value for Non-VM Files includes metadata--system files that comprise system overhead.

A Virtualization Manager data collection policy can be configured to gather data as follows:

- Inventory (default option): Collect only the data associated with the VMs in the inventory.
- Collect all files in datastore: Perform a full scan of the datastore to discover files that are resident on the datastore, but not associated with a VM in the inventory. This option finds files that may be invisible to VMware, but that are consuming space on the datastore.

When the Inventory option is configured in the data collector policy, it has the following impact on fields shown in this report:

- Non-VM Files and VM Not In Inventory will be zero.
- The Unknown category in the pie chart will often be significantly higher than if an all files collection had been done.
- Total VM Used will be reduced by the amount that would have been collected as VM Not In Inventory if an all files collection had been done.

When the Collect all files in datastore option is configured in the data collector policy, it has the following impact on fields shown in this report:

- The Unknown category in the pie chart will usually be non-zero and will often be significantly higher than the default inventory collection.

VM Datastore Overview

For a view of relevant VMware datastore information, go to the Inventory and select **Datastores** or individual datastores in the default hierarchy view (object type and subsystem). Once in the Inventory, the Overview tab presents several reports in dashboard format.

Inventory > Default Hierarchy View > Datastores > Overview Tab

The following reports are listed, prefaced by a VM Datastore Overview for the time period.

- See “ Datastore Usage Breakdown ” on page 662.
- See “Datastore Usage History” on page 665.
- See “Datastore Disk Read/Write Performance” on page 665.
- See “Datastore Disk Latency Performance” on page 666.

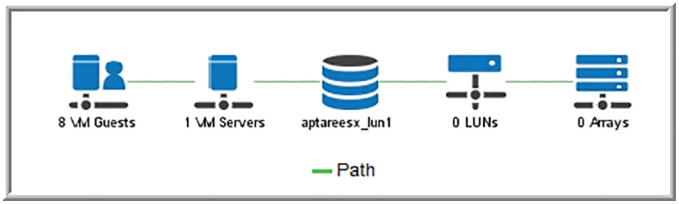


Table 20-25 VM Datastore Overview

Column name	Description
Datastore	The number of datastores in the selected scope. If a single datastore is selected, the name of the datastore is displayed.
Capacity	The total capacity of the datastores, in the selected scope and time period.
# VM Servers	The number of VM servers, in the selected scope and time period. To access additional details, click the link to view the VM Server Summary.
Used	The amount of datastore storage that is in use.
Free	The amount of datastore storage that is available.
Capacity Usage	Capacity Usage = Used/Capacity.
# VM Guests	The number of VM guests associated with the datastores, in the selected scope and time period. To access additional details, click the link to view the VM Summary.

Table 20-25 VM Datastore Overview *(continued)*

Column name	Description
Powered On	The number of VMs that are powered on, in the selected scope and time period. To access additional details, click the link to view the VM Summary.

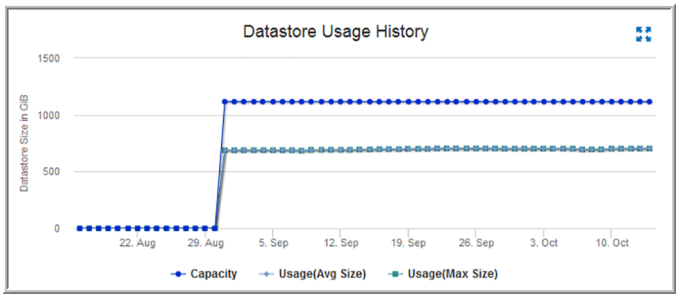
Datastore Usage History and Performance

Several reports the Inventory view provide an at-a-glance overview of VMware datastore capacity and usage over time, as well as disk performance over time.

- See “VM Datastore Overview” on page 664.
- See “ Datastore Usage Breakdown ” on page 662.

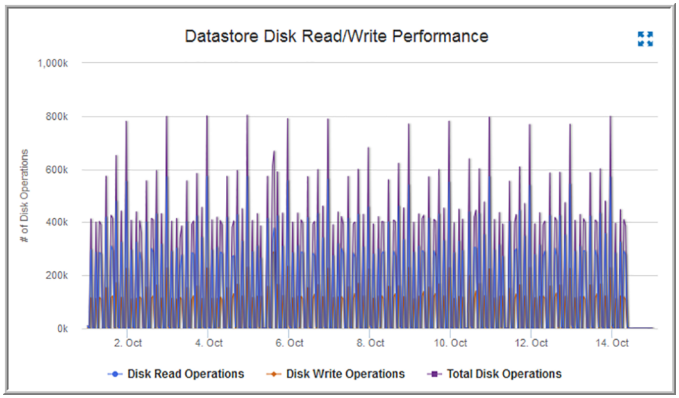
Datastore Usage History

View VMware datastore usage over time, for the selected scope and time period.



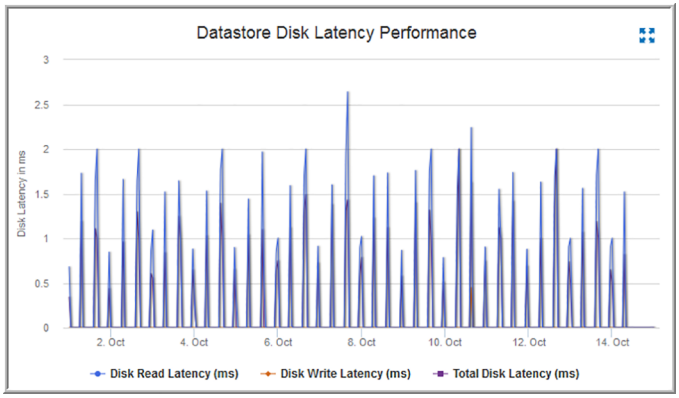
Datastore Disk Read/Write Performance

View VMware datastore disk read/write performance over time, for the selected scope and time period.



Datastore Disk Latency Performance

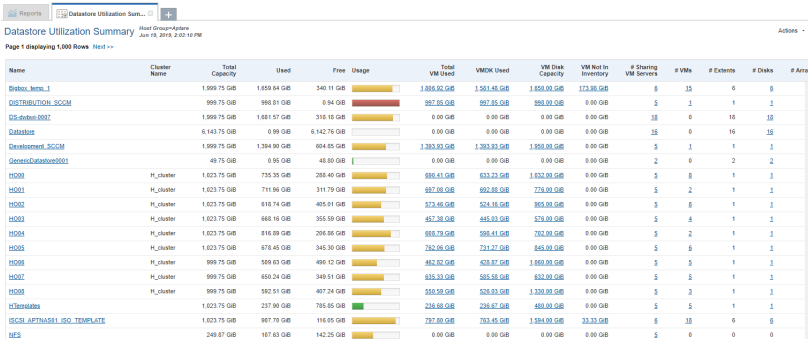
View VMware datastore disk latency over time, for the selected scope and time period.



Datastore Utilization Summary

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Storage Capacity & Utilization > Datastore Utilization Summary



The Datastore Utilization Summary gives you the ability to identify what storage is in use and where it is located. Click on the Datastore link to view the details where you can view the Extents and then link to the corresponding physical details.

See " Datastore Detail " on page 669.

See " Physical Disk Detail " on page 675.

See " Datastore Detail " on page 669.

Table 20-26 Datastore Utilization Summary

Column name	Description
Name	The name of the Datastore. See " Datastore Detail " on page 669.
Cluster Name	The name of the cluster to which the datastore belongs.
Total Capacity	Total capacity of the datastore.
Used	Total amount of the datastore that is in use, both VM usage and other files.
Free	Amount of the datastore that currently is unused and potentially available.
Usage	Mouse over the thermometer to view the % of the total capacity used.
Total VM Used	Size of the virtual machine, which includes VMDK files, log files, and snapshots; the sum of all the files taking up storage by this VM. This value links to the VM Files Summary report, which lists the usage details. For details about data collection options that impact these values, refer to the following report. See "Data Collection Options and Datastore Utilization" on page 669.

Table 20-26 Datastore Utilization Summary *(continued)*

Column name	Description
VMDK Used	The sum of all virtual disks - VMDK (.vmdk files), that are occupying space. This sum does not include snapshot metadata files, however, other outdated snapshot data may be in vmdk files. For this reason, VMDK Used may exceed VM Disk Capacity, the amount that was configured when the VM was created.
VM Disk Capacity	Amount of storage that was provisioned to the Guest OS when the virtual disk was originally created for the VM. For thin storage, this can be larger than the Datastore capacity associated with the VM.
VM Not In Inventory	Indicates VMs that are not currently in the inventory, but are taking up space in the Datastore. These are VMs that are not visible in VMware vCenter. Click this link to view the VM Files Summary. This value will be zero if data collection is configured to collect data from only datastores associated with VMs in the inventory. For details about data collection options that impact these values, refer to the following report. See "Data Collection Options and Datastore Utilization" on page 669.
# Sharing VM Servers	The number of hosts configured to access this datastore.
# VMs	Number of virtual machines stored on this datastore.
# Extents	The number of extents that were added to expand the datastore (up to 32 physical storage extents).
# Disks	The number of disks used by this datastore (up to 32 physical disks).
# Arrays	The number of arrays from which this datastore gets physical storage. See " Array Capacity and Utilization " on page 164.
Is Thin	Indicates whether or not the datastore supports thin provisioning on a per file basis. When thin provisioning is used, backing storage is lazily allocated. This is supported by VMFS3. VMFS2 always allocates storage eagerly. Thus, this value is false for VMFS2. Most NAS systems always use thin provisioning. They do not support configuring this on a per file basis, so for NAS systems this value is also false.
VMDK Max File Size	Some of the capacity of a virtual disk from the point of view of a virtual machine.

Data Collection Options and Datastore Utilization

A Virtualization Manager data collection policy can be configured to gather data as follows:

- Inventory (default option): Collect only the data associated with the VMs in the inventory.
- Collect all files in datastore: Perform a full scan of the datastore to discover files that are resident on the datastore, but not associated with a VM in the inventory. This option finds files that may be invisible to VMware, but that are consuming space on the datastore.

When the Inventory option is configured in the data collector policy, it has the following impact on fields shown in this report:

- VM Not In Inventory will be zero.
- Total VM Used will be reduced by the amount that would have been collected as VM Not In Inventory if an all files collection had been done.

Datastore Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Virtualization Manager > Storage Capacity & Utilization > Datastore Utilization

Click a Datastore name.

Table 20-27 Datastore Detail

Column name	Description
Name	Name of the datastore.
Cluster Name	The name of the cluster to which the datastore belongs.

Table 20-27 Datastore Detail *(continued)*

Column name	Description
Type	Datastore type: a specific NFS file system or NAS.
Server	This field is related to Datastore type and will be populated only when the type is NAS.
Folder	The VM folder in which this Datastore is located. This field is related to Datastore type and will be populated only when the type is NAS.
Multiple Host Access	Indicates if this datastore can be shared by multiple hosts: Yes or No.
Last Updated	Timestamp of the last write access.
VMDK Max File Size	The capacity of a virtual disk from the point of view of a virtual machine.
# VMs	This number links to the VM Summary report, listing the VMs associated with this Datastore.
# Extents	The number of extents that were added to expand the datastore (up to 32 physical storage extents).
# Disks	The number of physical disks associated with the datastore.
# Arrays	The number of arrays from which this datastore gets physical storage links to Array Capacity & Utilization.
VMDK Used	The sum of all virtual disks - VMDK (.vmdk files). This sum does not include snapshot metadata files, however, other outdated snapshot data may be in vmdk files. For this reason, VMDK Used may exceed VM Disk Capacity, the amount that was configured when the VM was created.
Total VM Used	Size of the virtual machine, which includes VMDK files, log files, and snapshots; the sum of all the files taking up storage by this VM. This value links to the VM Files Summary report, which lists the usage details. For details about data collection options that impact these values refer to the following report. See "Data Collection Options and Datastore Utilization" on page 669.

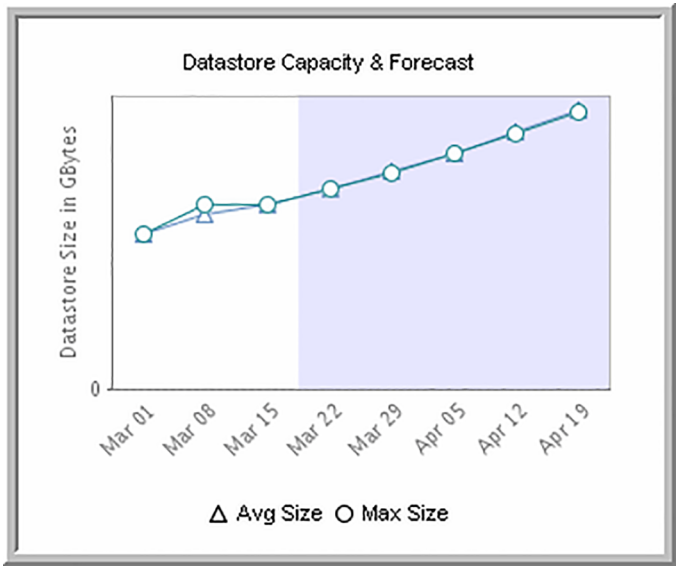
Table 20-27 Datastore Detail *(continued)*

Column name	Description
Is Thin	Indicates whether or not the datastore supports thin provisioning on a per file basis. When thin provisioning is used, backing storage is lazily allocated. This is supported by VMFS3. VMFS2 always allocates storage eagerly. Thus, this value is false for VMFS2. Most NAS systems always use thin provisioning. They do not support configuring this on a per file basis, so for NAS systems this value is also false.
Total Capacity	Total capacity of this Datastore.
Free Capacity	Available capacity in this Datastore
Used Capacity	Amount of this Datastore's capacity already in use
Capacity Usage	Mouse over this thermometer to view the usage percentage
VMDK Allocated	This is the current allocated size of all VMDKs in the datastore.
VM Not In Inventory	Indicates VMs that are not currently in the inventory, but are taking up space in the Datastore. These are VMs that are not visible in VMware vCenter. Click this link to view the VM Files Summary. This value will be zero if data collection is configured to collect data from only datastores associated with VMs in the inventory. For details about data collection options that impact these values, refer to the following report. See "Data Collection Options and Datastore Utilization" on page 669.
Over Provisioned	Indicates what percentage is over-provisioned. This helps you in evaluating possible virtual environment expansions.

Datastore Detail: Datastore Usage Breakdown

See " Physical Disk Detail " on page 675.

Datastore Capacity & Forecast



Extents

Table 20-28 Datastore Detail

Column name	Description
Name	Name of the storage device with the following 4-part name: HBA, SCSI target, SCSI LUN, disk partition EXAMPLE: **vmhba1:1:3:1** If vmhba appears in the name, it indicates a specific physical HBA on the ESX server. If the fourth segment is not included in the name, it indicates that the datastore has consumed the whole disk/LUN.
Type	VMFS or NFS
Mode	physical, logical
Capacity	Capacity of this particular extent
Disk	The disk on which the extent resides

Table 20-28 Datastore Detail *(continued)*

Column name	Description
Array	If there is a LUN mapping, a link to the corresponding array accesses the Array Capacity and Utilization report.

Shared VM Servers

For a description of the fields in this table, refer to the following report.

See " VM Server Summary " on page 639.

VM Utilization

For a description of these fields, refer to the following report.

See " VM Summary " on page 650.

Data Collection Options and Datastore Details

A Virtualization Manager data collection policy can be configured to gather data as follows:

- Inventory (default option): Collect only the data associated with the VMs in the inventory.
- Collect all files in datastore: Perform a full scan of the datastore to discover files that are resident on the datastore, but not associated with a VM in the inventory. This option finds files that may be invisible to VMware, but that are consuming space on the datastore.

When the Inventory option is configured in the data collector policy, it has the following impact on fields shown in this report:

- Non-VM Files and VM Not In Inventory will be zero.
- Total VM Used will be reduced by the amount that would have been collected as VM Not In Inventory if an all files collection had been done.
- VMDK and Other VM Files values typically will not be impacted by the datastore collection type.

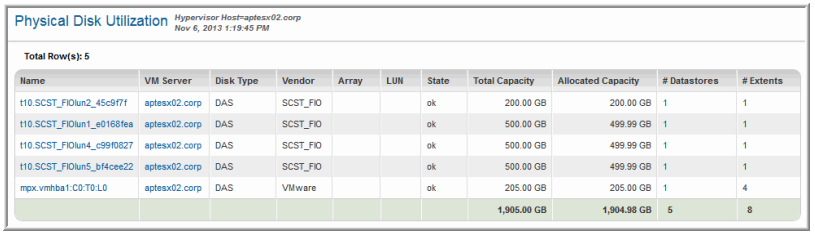
Physical Disk Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Storage Capacity & Utilization > Physical Disk Utilization

Use this report to view what disk space has been allocated to datastores.



Physical Disk Utilization

Hypervisor Host=aptesx02.corp
Nov 6, 2013 1:19:45 PM

Total Row(s): 5

Name	VM Server	Disk Type	Vendor	Array	LUN	State	Total Capacity	Allocated Capacity	# Datastores	# Extents
t10.SCST_FIOlun2_45c9f7f	aptesx02.corp	DAS	SCST_FIO			ok	200.00 GB	200.00 GB	1	1
t10.SCST_FIOlun1_e0168fea	aptesx02.corp	DAS	SCST_FIO			ok	500.00 GB	499.99 GB	1	1
t10.SCST_FIOlun4_c99f0827	aptesx02.corp	DAS	SCST_FIO			ok	500.00 GB	499.99 GB	1	1
t10.SCST_FIOlun5_bf4cee22	aptesx02.corp	DAS	SCST_FIO			ok	500.00 GB	499.99 GB	1	1
mpx.vmhba1:C0:T0:L0	aptesx02.corp	DAS	VMware			ok	205.00 GB	205.00 GB	1	4
							1,905.00 GB	1,904.98 GB	5	8

Table 20-29 Physical Disk Utilization

Column name	Description
Name	The raw device mapping path, with the following components: HBA, SCSI target, SCSI LUN.
VM Server	Name of the VM Server.
Disk Type	SCSI device type, such as disk.
Vendor	Device manufacturer.
Array	The array on which the disk resides; links to the Array Capacity & Utilization report.
LUN	The LUN path, with the following components HBA, SCSI target, SCSI LUN, disk partition. Links to the LUN Utilization Summary.
Total Capacity	Total capacity of the disk.
Allocated Capacity	The amount of the disk that has been allocated.
# Datastores	The number of datastores that are using the disk.
# Extents	The number of extents associated with the disk.

Physical Disk Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Virtualization Manager > Storage Capacity & Utilization > Physical Disk Utilization

Click a **Name** link.

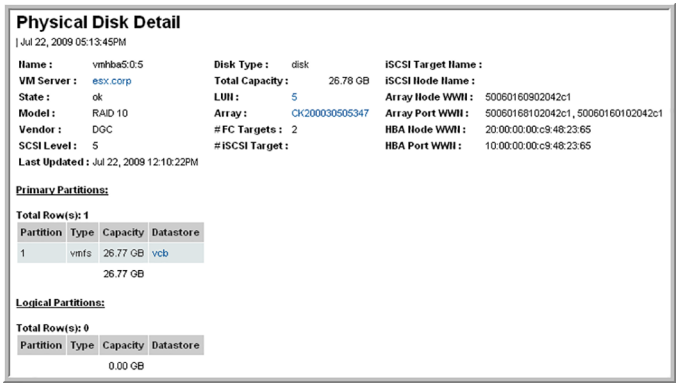


Table 20-30 Physical Disk Detail

Column name	Description
Name	The raw device mapping path, with the following components: HBA, SCSI target, SCSI LUN.
VM Server	Name of the VM Server.
State	Operational state of the disk, such as ok and degraded.
Model	Model of the disk.
Vendor	Device manufacturer.

Table 20-30 Physical Disk Detail *(continued)*

Column name	Description
SCSI Level	The device's SCSI level, an indicator of clock and bus speeds.
Last Updated	Last date and time of data collection.
Disk Type	SCSI device type, such as disk.
Total Capacity	Capacity of this disk.
LUN	The LUN path, with the following components HBA, SCSI target, SCSI LUN, disk partition Links to the LUN Utilization Summary.
Array	Name of the array in which the disk is located; links to the Array Capacity & Utilization report to view the array from with the LUN was provisioned.
# FC Targets	Number of Fibre Channel targets; paths to the array from the host.
# iSCSI Targets	The number of paths to the iSCSI array from the host.
iSCSI Target Name	The actual target name from the array's perspective.
iSCSI Node Name	The actual node name from the array's perspective.
Array Node WWN	Storage array node World Wide Name.
Array Port WWN	Storage array port's World Wide Name.
HBA Node WWN	Host bus adapter node World Wide Name.
HBA Port WWN	Host bus adapter port World Wide Name.

Datastore Performance Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Performance Reports > Datastore Performance Summary

Lists Datastores with disk I/O statistics, including read/write latency. Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.

Note: If some Datastores are missing from this report, it is likely due to virtual systems that have a status of pending deletion and therefore do not have any performance data to display.

Note: Column aggregations are displayed in green on a separate line with the label Aggregation. Rollover the value to display the aggregation type: Avg., Min., Max., or Sum. Column filtering is not supported on reports with aggregated columns.

Table 20-31 Datastore Performance Summary

Column name	Description
Datastore Name	Datastore name links to Datastore Detail.
Disk Name	Disk name links to the Physical Disk Detail.
VM Server	VM Server links to the VM Server Detail.
Disk I/O Data Xfer (KB/sec)	Disk I/O average (read xfer and write xfer) for all disks; the total for this column is the sum of the read and write averages.
Disk I/O Latency (ms)	Disk I/O wait time, which is the average for the report's Group By interval; Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.
Disk Read Data Xfer (KB/sec)	Average disk read transfer rate for the report's Group By interval.
Disk Read # Reads	Sum of the total number of disk reads for the report's Group By interval.
Disk Read Latency (ms)	Disk read wait time, which is the average for the report's Group By interval; Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.
Disk Write Data Xfer (KB/sec)	Average disk write transfer rate for the report's Group By interval.
Disk Write # Writes	Sum of the total number of disk writes for the report's Group By interval.

Table 20-31 Datastore Performance Summary *(continued)*

Column name	Description
Disk Write Latency (ms)	Disk write wait time, which is the average for the report's Group By interval; Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.
Show Chart	Check the boxes for which you want to generate charts. Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host. You can select a Group by option in the scope selector.

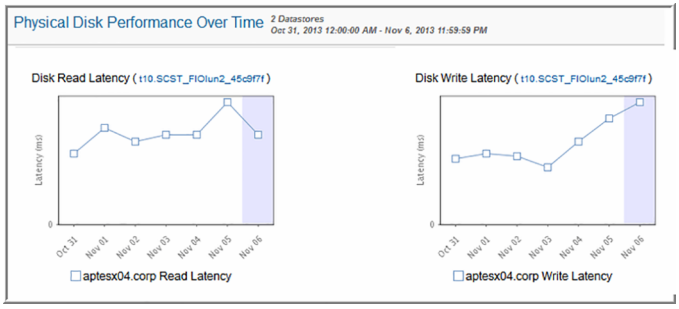
Physical Disk Performance Over Time

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Performance Reports > Physical Disk Performance Over Time

View charts of physical disk performance over time, with drilldowns to VM details. Mouse over points on the graphs to view the number of reads and writes, along with latency values.



Physical Disk Performance Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Performance Reports > Physical Disk Performance Summary

List physical disks with I/O statistics, including read/write latency. Use the Show Chart option to graph the statistics.

Table 20-32 Physical Disk Performance Summary

Column name	Description
Physical Disk Name	Disk name links to the Physical Disk Detail.
VM Server	VM Server links to the VM Server Detail.
# Datastore	Number links to the Datastore Utilization Summary.
Disk I/O Data Xfer (KB/sec)	Disk I/O transfer rate.
Disk I/O Latency (ms)	Disk I/O wait time, which is the average for the report's Group By interval.
Disk I/O Max Queue Depth	Disk I/O maximum queue depth.
Disk Read Data Xfer (KB/sec)	Disk read transfer rate.
Disk Read # Reads	Sum of the number of disk reads for the report's Group By interval.
Disk Read Latency (ms)	Disk read wait time, which is the average for the report's Group By interval.
Disk Read Queue Latency (ms)	Average amount of time taken during the collection interval per SCSI read command in the VMKernel queue.
Disk Write Data Xfer (KB/sec)	Disk Write transfer rate.

Table 20-32 Physical Disk Performance Summary *(continued)*

Column name	Description
Disk Write # Writes	Sum of the number of disk writes for the report's Group By interval
Disk Write Latency (ms)	Disk write wait time, which is the average for the report's Group By interval
Disk Write Queue Latency (ms)	Average amount time taken during the collection interval per SCSI write command in the VMKernel queue.
Show Chart	Check the boxes for which you want to generate charts. Use the Advanced option in the scope selector to set a Group by.

VM Guest to LUN Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Performance Reports > VM Guest to LUN Performance

This report shows connectivity from VM Guests to LUNs with performance metrics and drilldowns to detail reports to facilitate troubleshooting of performance issues within your environment.

Table 20-33 VM Guest to LUN Performance

Column name	Description
VM Guests	
VM Guest	Name of the Virtual Machine (VM). See " VM Detail " on page 652.
Host Name	Name of the host. See " Host Details " on page 838.
Avg Read Data Xfer (KB/sec)	Average disk read transfer rate for the time period selected for the scope of the report.

Table 20-33 VM Guest to LUN Performance *(continued)*

Column name	Description
# Reads	Total number of disk reads for the time period selected for the scope of the report.
Avg Write Data Xfer (KB/sec)	Average disk write transfer rate for the time period selected for the scope of the report.
# Writes	Total number of disk writes the time period selected for the scope of the report.
Datastores	
VM Host Name	Name of the VM Server hosting the virtual server. Click the link to view the details. See " VM Server Detail " on page 642.
Datastore Name	Datastore name. Click the link to view the details. See " Datastore Detail " on page 669.
Avg Latency (ms)	Average disk I/O wait time for the time period selected for the scope of the report.
Max Latency (ms)	Maximum disk I/O wait time for the time period selected for the scope of the report.
Avg Read Data Xfer (KB/sec)	Average disk read transfer rate for the time period selected for the scope of the report.
# Reads	Total number of disk reads for the time period selected for the scope of the report.
Avg Read Latency (ms)	Average disk read wait time for the time period selected for the scope of the report.
Max Read Latency (ms)	Maximum disk read wait time for the time period selected for the scope of the report.
Avg Write Data Xfer (KB/sec)	Average disk write transfer rate for the time period selected for the scope of the report.
# Writes	Total number of disk writes for the time period selected for the scope of the report.
Avg Write Latency (ms)	Average disk write wait time for the time period selected for the scope of the report.

Table 20-33 VM Guest to LUN Performance *(continued)*

Column name	Description
Max Write Latency	Maximum disk write wait time for the time period selected for the scope of the report.
Switches	
# Ports	Number of switch FC ports between VM Host and Storage Array. Click the link to view Switch Port Performance Details.
Arrays	
# Ports	Number of Array ports. Click the link to view Array Port Throughput.
LUNs	
Storage Array	Name of the storage array. Click the link to view Array Capacity & Utilization.
LUN Name	Logical unit name. See "LUN Performance" on page 559.
Avg Read IOPS	Average Read IO throughput per second for the time period selected for the scope of the report.
Avg Write IOPS	Average Write IO throughput per second for the time period selected for the scope of the report.
Avg Total Read IOPS	Average IO throughput per second for the time period selected for the scope of the report.
Max Read Response (ms)	Identifies the LUN with the longest read response time. This is the highest response time for a single LUN. Use this value to help identify the LUN with a performance problem.
Max Write Response (ms)	Identifies the LUN with the longest write response time. This is the highest response time for a single LUN. Use this value to help identify the LUN with a performance problem.
Max Total Response (ms)	Identifies the LUN with the longest total response time. This is the highest response time for a single LUN. Use this value to help identify the LUN with a performance problem.

VM Performance Over Time

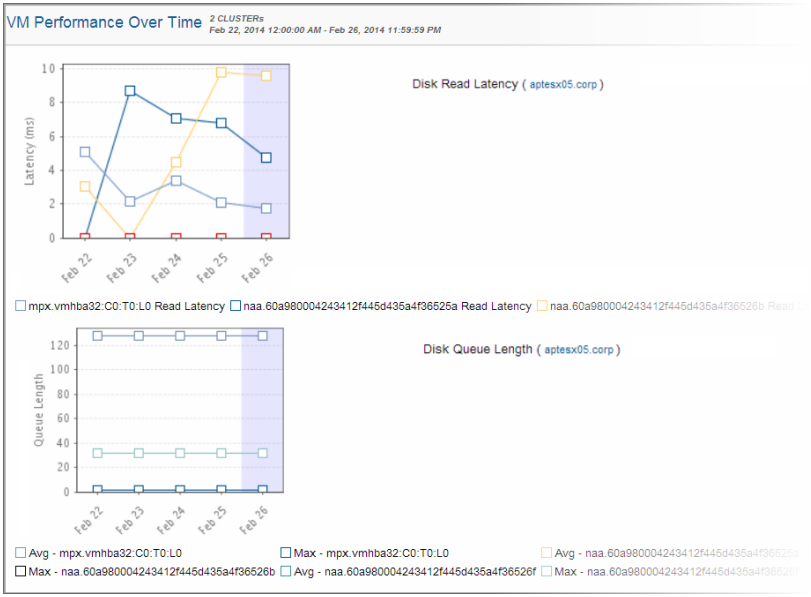
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Performance Reports > VM Performance Over Time

Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.

Use the scope selector to aggregate at the ESX Cluster level or an arbitrary Host Group level. If selected only individual ESX Clusters and Host Groups of ESX servers are reported on. Rendered graphs represent aggregate performance of all individual ESX servers in the Cluster or Host Group. You can drill down to the VM Detail report and VM Server Detail.



VM Performance Summary

To access this report, in the Navigation pane select:

Virtualization Manager > Performance Reports > VM Performance Summary

When a VM exhibits performance issues, it could stem from either the server, the actual storage, or other network factors. Once server bottlenecks are addressed, the next step is to determine if there's an issue with the LUN and disks. While vCenter provides performance statistics, this information is not always readily available to Storage Administrators.

The VM Performance Summary enables a Storage Administrator to have an at-a-glance look at I/O performance statistics: CPU, Memory, and Disk I/O. In addition, Disk I/O is broken out into reads and writes, with latency statistics. In addition, this report provides the necessary CPU Ready States and CPU Wait States to determine if an ESX host can handle another VM Guest.

- Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.
- Note an attempt is made to collect all the statistics listed in the table, however by default most statistics are only available for an hour. If you require a longer schedule, configure the 5-minute summaries to the desired statistics level, otherwise there may be gaps in the data. As a best practice recommendation, enable and configure the 5-minute summaries for at least level 3 statistics. Setting higher levels for intervals may result in increased load and disk usage on the vCenter
- In the vCenter VM ESX Server settings, the statistics level must be set to Level 4 in order for balloon statistics to appear in the VM Performance Summary report. When Level 4 is configured, it will also capture the other statistics, as listed in the Performance Counters and Statistics Level Settings table.
- In the vCenter VM ESX Server settings, the statistics level must be set to the maximum level based on the counter listed in the Performance Counters and Statistics Level Settings table.

  Note that level 1 is always captured, but other levels may need to be set for other performance counters. Configuring a level lower than the required level for a counter will result in a failure to retrieve the relevant data from VSphere and will cause inaccurate values in reports.

  See the section called "vCenter Statistics Configuration" on page 685.

Table 20-34 Performance Counters and Statistics Level Settings

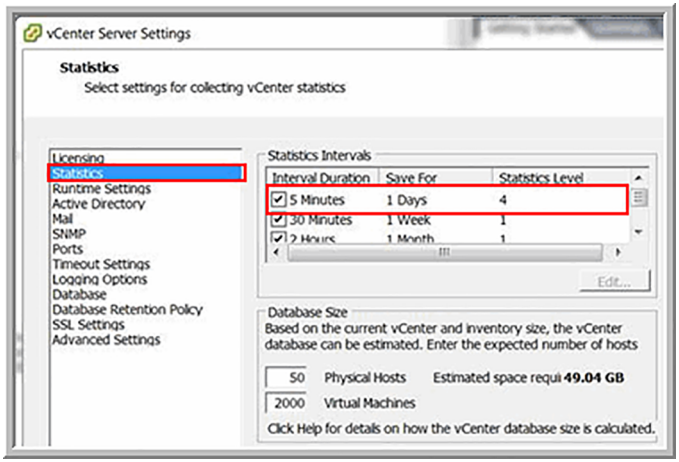
Performance Counters Collected	Statistics Level
Memory (%)	2
Memory Usage (%)	1
Avg CPU Usage (%)	1

Table 20-34 Performance Counters and Statistics Level Settings *(continued)*

Performance Counters Collected	Statistics Level
Avg CPU Usage (KB)	1
Disk Avg Usage (KB/sec)	1
Disk Reads (KB/sec)	2
Disk Writes (KB/sec)	2
Disk # of Reads	3
Disk # of Writes	3
Disk Total Latency (ms)	2
Disk Read Latency (ms)	2
Disk Write Latency (ms)	2
Disk Queue Latency (ms)	2
Disk Queue Read Latency (ms)	2
Disk Queue Write Latency (ms)	2
Maximum Queue Depth	1
Memory Balloon	4

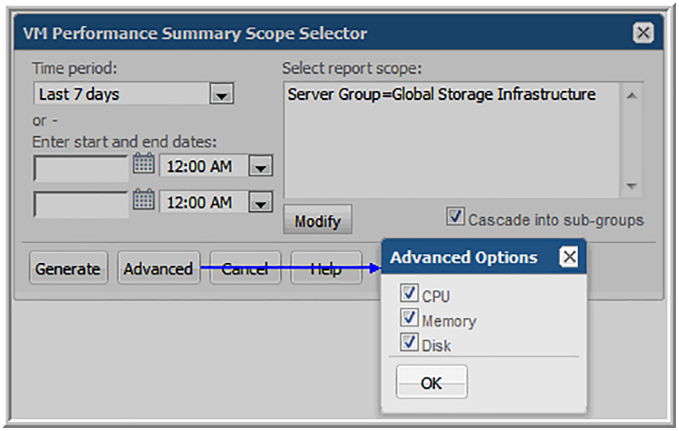
vCenter Statistics Configuration

The following example illustrates how statistics levels can be set in vCenter.



Focusing on Specific I/O Statistics

To isolate a particular set of statistics--for example, Disk I/O--use the Advanced options in the Report Designer and click the relevant checkbox.



All values represent the average for the interval selected in the “Group By” option.

Table 20-35 VM Performance Summary

Column name	Description
Name	Description
Name	The name of the VM. Click on this link to view VM Server Details, if it is a host, or VM Detail, if it is a VM.

Table 20-35 VM Performance Summary *(continued)*

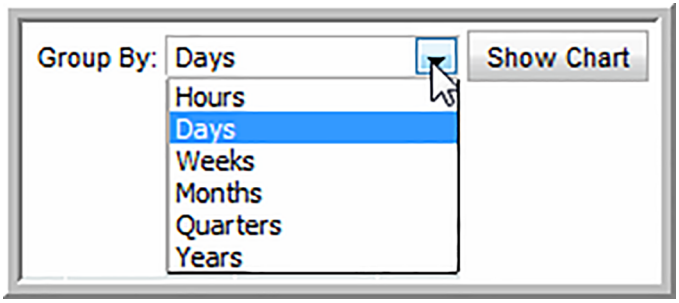
Column name	Description
Type	Type of partition: host server or VM; Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, latency performance charts are available only for partitions of Type=Host.
CPU Usage	
CPU %	The VM's CPU usage percentage.
CPU Max %	Maximum CPU usage percentage for the report's Group By interval.
Avg Wait %	Calculated for VM guests based on the total milliseconds for each data capture interval included in the reporting interval, divided by the number of milliseconds in the reporting interval; used to determine how long VM guests wait to get CPU.
Max Wait %	Calculated for VM guests based on the total milliseconds for each data capture interval included in the reporting interval, divided by the number of milliseconds in the reporting interval; used to determine how long VM guests wait to get CPU.
Avg Ready %	Calculated for VM guests based on the total milliseconds for each data capture interval included in the reporting interval, divided by the number of milliseconds in the reporting interval; used to determine how long VM guests wait to get CPU.
Max Ready %	Calculated for VM guests based on the total milliseconds for each data capture interval included in the reporting interval, divided by the number of milliseconds in the reporting interval; used to determine how long VM guests wait to get CPU.
CPU MHz	Average processor speed for the report's Group By interval.
CPU Max MHz	Maximum processor speed for the report's Group By interval.
Balloon Memory	
Min (KB)	Minimum of all the data collection intervals within the report's time period.
Avg (KB)	Average of all the data collection intervals within the report's time period.
Max (KB)	Maximum of all the data collection intervals within the report's time period.

Table 20-35 VM Performance Summary *(continued)*

Column name	Description
Memory Usage	
Memory %	Average memory usage percentage for the report's Group By interval.
Memory Max %	Maximum memory percentage for the report's Group By interval.
Memory MB	Average memory usage for the report's Group By interval.
Memory Max MB	Maximum memory usage for the report's Group By interval.
Disk I/O	
Disk I/O Data Xfer (KB/sec)	Disk I/O average transfer rate for all disks for the VM.
Disk Read	
Data Xfer (KB/sec)	Disk read transfer rate, average for the report's Group By interval.
# Reads	Sum of the number of disk reads for the report's Group By interval.
Disk Write	
Data Xfer (KB/sec)	Disk write transfer rate for the report's Group By interval.
# Writes	Sum of the number of disk writes for the report's Group By interval.
Show Chart	Check the boxes for which you want to generate charts and click Show Chart to view the VM Performance Over Time graphs. Disk Read and Write Latency is available only for ESX Hosts (not VMs). Therefore, the latency performance charts are available only for partitions of Type=Host. See the section called "Show Chart Parameter" on page 688.

Show Chart Parameter

When selecting the Show Chart option in a VM performance report, a Group By parameter can be applied to the selection.



Note that the default scope selector time period is set to Last 7 Days, so if a Group By greater than Days is selected--that is, Weeks, Months, Quarters, or Years--a chart will not be created.

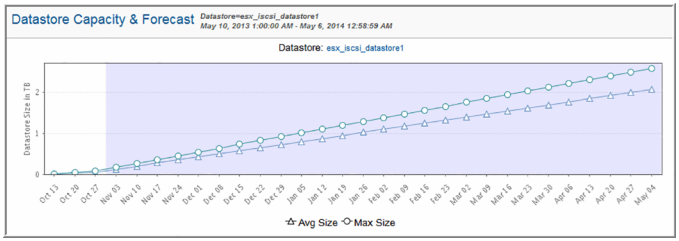
Datastore Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Forecast & Planning > Datastore Capacity Forecast

View average capacity and forecasted usage graphs for each Datastore.



The Datastore name links to the details.

See “ Datastore Detail ” on page 669.

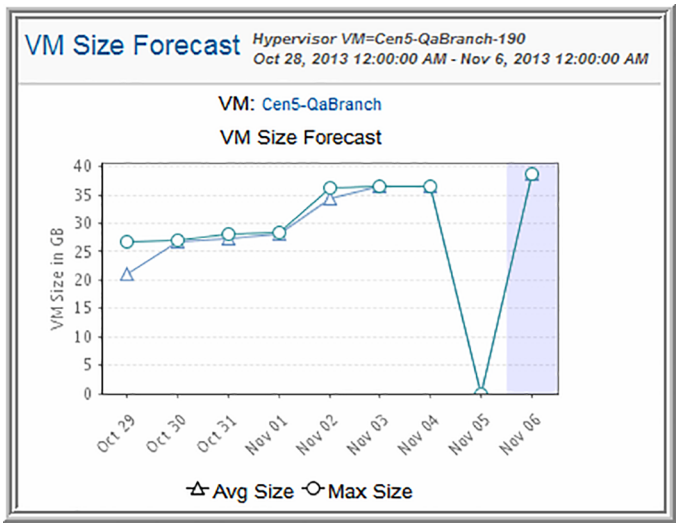
VM Size Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Forecast & Planning > VM Size Forecast

This VM report projects the growth of the VM over time, based on historical data. Mouse over the Avg and Max icons to view the values.



The VM name links to the details.

See “ VM Server Detail ” on page 642.

Chapter 21

Fabric Manager Reports

This chapter includes the following topics:

- Host to Storage Dashboard
- Paths Portlet
- Paths in the Host to Array Connectivity Portlet
- Zones in the Host to Array Connectivity Portlet
- Port Summary
- Port Details
- SAN Fabric Summary
- SAN Fabric Details
- Switch Summary
- Switch Details
- Node Port Summary
- Zone Summary
- Zone Details
- Zone Set Summary
- Zone Members
- Bottom Ranking
- Port Distribution
- SAN Resource Utilization

- Top Ranking
- SAN Fabric Change Status Summary
- SAN Fabric Change Summary
- Port Performance Dashboard
- Ports Portlet
- Port Details Portlet
- Port Performance Summary
- Port Performance Details
- SAN Fabric Usage
- Port Error Rate Forecast
- Port Utilization Forecast

Host to Storage Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Administration Reports > Host to Storage Dashboard

This four-portlet interactive dashboard initially presents a list of host-to-storage paths in the upper left portlet. The other three portlets remain blank until a host is selected. Each time you select a host, the upper right portlet displays a topology map of the host-to-storage configuration.

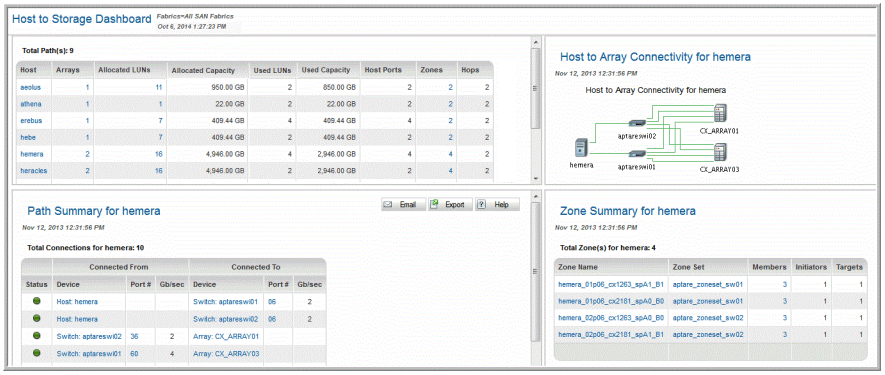
Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Working with Topology

This topology represents the connections from a particular host to arrays providing storage to the host.

Connectivity between objects is shown in a topology diagram. Each object displays the type and the count.

- Mouse-over each device (host, switch, storage array) and paths to view details
- Single click a connection to view the details
- Double click a device to display detailed reports
- Expand the topology window to zoom and navigate through larger diagrams



Paths Portlet

Table 21-1 Paths Portlet

Column name	Description
Host	Click on a host to view the Host-to-Array Connectivity in the portlet to the right of the Paths portlet
Arrays	The number of arrays. Click the link to view the Array Capacity & Utilization.
Allocated LUNs	The number of LUNs allocated to this host. Click to drill down the LUN Utilization Summary.
Allocated Capacity	The total storage capacity allocated.
Used LUNS	The number of LUNs being used by this host.
Used Capacity	The amount of storage capacity used by this host.
Host Ports	The number of HBAs on this host.
Zones	The number of zones in this host.

Table 21-1 Paths Portlet *(continued)*

Column name	Description
Hops	The number of connection made by one or more ISLs between devices.

Paths in the Host to Array Connectivity Portlet

Table 21-2 Paths in the Host to Array Connectivity Portlet

Column name	Description
Status	The status of the connection.
Connected From	
Device	The name of the device from which a connection is made.
Port #	The unique identifier of the Port connected from. Click the link for details. See " Port Details " on page 696.
Gb/Sec	The port speed.
Connected To	
Device	The name of the device to which the connection is made.
Port #	The number of the Port connected to. This is a string, not a numeric field. Click the link for details. See " Port Details " on page 696.
Gb/Sec	The port speed.

Zones in the Host to Array Connectivity Portlet

Table 21-3 Zones in the Host to Array Connectivity Portlet

Column name	Description
Zone Name	The name of the zone. Click the link for details. See " Zone Details " on page 706.
Zone Set	The name of the zone set.
Members	The number of members in each zone.

Table 21-3 Zones in the Host to Array Connectivity Portlet *(continued)*

Column name	Description
Initiators	The number of initiators that send a request to the node.
Targets	The number of receivers of the request.

Port Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Administration Reports > Port Summary

Table 21-4 Port Summary

Column name	Description
Port	
#	The number of the Port in the SAN Fabric. This is a string, not a numeric field. Click the link for details. See " Port Details " on page 696.
Status	The status of the port
Switch Name	The switch associated with the port. Click the link for details. See " Switch Details " on page 701.
Gb/sec	Speed of the port
Type	The type of port
Licensed	Indicates if the port is licensed
WWN	The Worldwide name of the port
Connected to	
Device	The device to which the port is connected. Click the link to view the Array Capacity & Utilization Report.
WWN	The Worldwide name of the port to which it is connected

Table 21-4 Port Summary *(continued)*

Column name	Description
Alias	Corresponds to the WWN of the port.

Port Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Administration Reports > Port Summary

Click a **Port #** link.

Table 21-5 Port Details

Column name	Description
Port #	The number of the port in the SAN Fabric. This is a string, not a numeric field.
Element Name	The name of the port
Status	The status of the port
Type	The type of port
WWN	The Worldwide name of the port
Switch Name	The switch associated with the port. Click the link for details. See " Switch Details " on page 701.
SAN Fabric	The name of the SAN Fabric. Clink the link for details. See " SAN Fabric Details " on page 698.
Connected Device	The device to which the port is connected

Table 21-5 Port Details *(continued)*

Column name	Description
Connected Port WWN	The WWN of the connected port
Caption	The short description of the device
Device Identifier	Unique identifier for the connected device
Alias	Corresponds to the WWN of the port.
Gb/Sec	The speed of the port
Auto Sense	Indicates if the port can automatically determine the speed of attached devices
Licensed	Indicates if there is a license for this port
ISL Mode Enabled	Denotes inter-switch link capabilities
Trunked	The capability for multiple connections defined as a single link, on inter-switch links
Virtual	Indicates if it is a virtual switch (Cisco)
NPIV Enabled	Indicates if a single Fibre Channel port is enabled to appear as multiple, unique ports.
Domain Port #	This port's number within the switch's domain
State	Enabled or Disabled. If the port is not connected, the state is not applicable.
Max Gb/Sec	The maximum speed of the port.
Port Index	The port's position in the slot or switch

SAN Fabric Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Administration Reports > SAN Fabric Summary

A SAN Fabric can range from a basic fibre-channel connection between two devices to a more complex network commonly found in large enterprise environments.

Table 21-6 SAN Fabric Summary

Column name	Description
SAN Fabric	The name of the SAN Fabric. Clink the link for details. See " SAN Fabric Details " on page 698.
State	Denotes the state of the SAN Fabric as either enabled, disabled, or not applicable.
# of Zone Sets	The number of zone sets in the SAN Fabric. Click the link for details. See " Zone Set Summary " on page 708.
# of Zones	The number of zones in the SAN Fabric. Click the link for details. See " Zone Summary " on page 705.
# of Arrays	The number of arrays in the SAN Fabric. Click the link to view the Array Capacity & Utilization.
# of Hosts	The number of hosts in the SAN Fabric. Click the link to view the Host Capacity & Utilization.
# of Switches	The number of switches in the SAN Fabric. Click the link for details. See " Switch Summary " on page 700.
# of Ports	The number of ports in the SAN Fabric. This is a string, not a numeric field. Click the link for details. See " Port Summary " on page 695.
# of Conn Dev	The total number of connected devices. Click the link for details. See " Node Port Summary " on page 704.

SAN Fabric Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Administration Reports > SAN Fabric Summary

Click a **SAN Fabric Name** link.

Table 21-7 SAN Fabric Details

Column name	Description
SAN Fabric	The name of the SAN Fabric
Creating Class Name	The name of the class or subclass for this SAN Fabric
Element Name	The name used to identify the SAN Fabric
Enabled State	Denotes the state of the SAN Fabric as either enabled, disabled, or not applicable
SAN Fabric Identifier	An administrative ID for the SAN Fabric
Interop Mode	Indicates if the switch is operating in interoperability mode.
Is Deleted	Indicates if the switch is offline or that a logical switch was deleted.
Status	Indicates the current status of the SAN Fabric.
Last Updated	The last time the database table was updated with the collected data
Name Format	Identifies how the name of the SAN Fabric is generated.
# of Switches	The number of switches in the SAN Fabric. Click the link details. See " Switch Summary " on page 700.
# of Ports	The number of ports in the SAN Fabric.This is a string, not a numeric field. Click the link for details. See " Port Summary " on page 695.
# of Zone Sets	The number of zone sets in the SAN Fabric. Click the link for details. See " Zone Set Summary " on page 708.

Table 21-7 SAN Fabric Details *(continued)*

Column name	Description
# of Zones	The number of zones in the SAN Fabric. Click the link for details. See " Zone Summary " on page 705.
# of Arrays	The number of arrays in the SAN Fabric. Click the link to view the Array Capacity and Utilization.
# of Hosts	The number of hosts in the SAN Fabric. Click the link to view the Host Capacity and Utilization.

Switch Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Administration Reports > Switch Summary

Table 21-8 Switch Summary

Column name	Description
Switch Name	The name of the switch in the SAN Fabric. Click the link for details. See " Switch Details " on page 701.
Physical Switch Name/Serial No	The name of the physical switch or serial number.
IP Address	The IP address of the switch
Vendor	The switch vendor
Model	The model of the switch
Status	Indicates the status of the switch
WWN	The Worldwide Name of the switch
Total Ports	The total number of ports. Click the link for details. See " Port Summary " on page 695.

Table 21-8 Switch Summary *(continued)*

Column name	Description
Licensed Ports	Click the number of licensed ports on the switch to view the summary. See " Port Summary " on page 695.
Used Ports	The total number of used ports. Click the link to view the summary. See " Port Summary " on page 695.
Unused Ports	The total number of free ports. Click the link to view the summary. See " Port Summary " on page 695.
# of Conn Dev	The total number of connected devices on the switch. Click the link to view the summary. See " Node Port Summary " on page 704.
SAN Fabric	The name of the SAN Fabric
Firmware Version	The version of the software on the switch

Switch Details

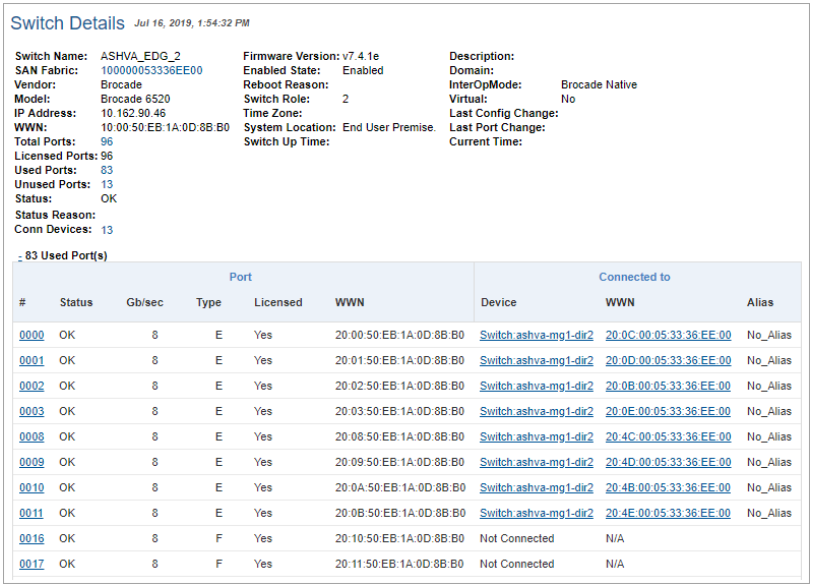
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Administration Reports > Switch Summary

Click on a **Switch Name** link.



Switch Details Jul 16, 2019, 1:54:32 PM

Switch Name:	ASHVA_EDG_2	**Firmware Version:**	v7.4.1e	**Description:**	
SAN Fabric:	100000053336EE00	**Enabled State:**	Enabled	**Domain:**	
Vendor:	Brocade	**Reboot Reason:**		**InterOpMode:**	Brocade Native
Model:	Brocade 6520	**Switch Role:**	2	**Virtual:**	No
IP Address:	10.162.90.46	**Time Zone:**		**Last Config Change:**	
WWN:	10:00:50:EB:1A:0D:8B:B0	**System Location:**	End User Premise.	**Last Port Change:**	
Total Ports:	96	**Switch Up Time:**		**Current Time:**	
Licensed Ports:	96				
Used Ports:	83				
Unused Ports:	13				
Status:	OK				
Status Reason:					
Conn Devices:	13				

83 Used Port(s)

Port						Connected to		
#	Status	Gb/sec	Type	Licensed	WWN	Device	WWN	Alias
0000	OK	8	E	Yes	20:00:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:0C:00:05:33:36:EE:00	No_Alias
0001	OK	8	E	Yes	20:01:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:0D:00:05:33:36:EE:00	No_Alias
0002	OK	8	E	Yes	20:02:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:0B:00:05:33:36:EE:00	No_Alias
0003	OK	8	E	Yes	20:03:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:0E:00:05:33:36:EE:00	No_Alias
0008	OK	8	E	Yes	20:08:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:4C:00:05:33:36:EE:00	No_Alias
0009	OK	8	E	Yes	20:09:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:4D:00:05:33:36:EE:00	No_Alias
0010	OK	8	E	Yes	20:0A:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:4B:00:05:33:36:EE:00	No_Alias
0011	OK	8	E	Yes	20:0B:50:EB:1A:0D:8B:B0	Switch:ashva-mg1-dir2	20:4E:00:05:33:36:EE:00	No_Alias
0016	OK	8	F	Yes	20:10:50:EB:1A:0D:8B:B0	Not Connected	N/A	
0017	OK	8	F	Yes	20:11:50:EB:1A:0D:8B:B0	Not Connected	N/A	

Table 21-9 Switch Details

Column name	Description
Switch Name	The name of the Switch in the SAN Fabric
SAN Fabric	The name of the SAN Fabric
Vendor	The switch's vendor
Model	The model of the switch
IP Address	The IP address of the switch
WWN	The Worldwide name of the switch
Total Ports	The total number of ports. Click the link to view the summary. See " Port Summary " on page 695.
Licensed Ports	The number of licensed ports on the switch.
Used Ports	The total number of used ports. Click the link to view the summary. See " Port Summary " on page 695.
Unused Ports	The total number of free ports. Click the link to view the summary. See " Port Summary " on page 695.

Table 21-9 Switch Details *(continued)*

Column name	Description
Status	Indicates the status of the switch
Status Reason	Additional details related to the status
Conn Devices	The total number of connected devices (HBA Ports, Array Ports, etc) for the switch. Click the link to view the summary. See " Node Port Summary " on page 704.
Firmware Version	The version of the software on the switch
Enabled State	Indicates the state of the switch, either enabled or disabled
Reboot Reason	Indicates why the switch rebooted
Switch Role	Role of the switch
Time Zone	The time zone of where the switch is located
System Location	Where the system is located
Switch Up Time	The time the switch was last turned on
Description	The description of the switch
Domain	The domain name assigned to the switch
InterOpMode	Indicates if the switch is operating in interoperability mode.
Virtual	Indicates if it is a virtual switch (Cisco)
Last Config Change	Date and time of the last change made to the switch
Last Port Change	Date and time of the last change made to any port on the switch
Current Time	The last time the data was collected

Table 21-10 Switch Details

Column name	Description
Ports	
#	The number of the port. Click the link for details. See " Port Details " on page 696.
Status	The status of the port

Table 21-10 Switch Details *(continued)*

Column name	Description
Gb/Sec	The port speed
Type	The type of port
Licensed	The number of licensed ports on the switch.
WWN	The Worldwide Name of the port
Connected To	
Device	The device to which the port is connected. Click the link for details. See " Switch Details " on page 701.
WWN	The Worldwide Name of the port

Node Port Summary

This report enables users to see how each connected device, for example a host or an array, is connected from the SAN Fabric switch. This report is available as a drill down from # of Conn Dev column in the following reports:

- See " SAN Fabric Summary " on page 697.
- See " Switch Summary " on page 700.
- See " Switch Details " on page 701.

Table 21-11 Node Port Summary

Column name	Description
Node Port	
Device	The device to which the port is connected. Click the link to view the details.
WWN	The Worldwide Name of the port.
Alias	Corresponds to the WWN of the port.
Switch Port	
#	The number of the port. Click the link for details. See " Port Details " on page 696.

Table 21-11 Node Port Summary *(continued)*

Column name	Description
Status	The status of the switch port.
Switch Port WWN	The Worldwide Name of the port.
Type	The type of port.
Gb/Sec	The port speed.
Licensed	The number of licensed ports on the switch.
Switch Name	The name of the Switch in the SAN Fabric. Click the link for details. See " Switch Details " on page 701.

Zone Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Administration Reports > Zone Summary

Zones provide the mechanism for partitioning one environment from another in a switched fabric. This segregation supports several scenarios, such as restricting access for security purposes or configuring environments with different intents, such as testing versus production. For example, this segmentation scheme could be deployed to differentiate Linux environments from Windows environments, as storage provisioning differs by operating system; zones offer a way to more easily manage storage resources. Note that a device can be in more than one zone.

To facilitate availability and scalability, zones are deployed in a VSAN (Virtual SAN - partitioned physical storage). A zone is first defined and then activated, with only one zone set active at a time.

Table 21-12 Zone Summary

Column name	Description
Zone Name	The name of the zone in the SAN Fabric. Click the link for details. See " Zone Details " on page 706.

Table 21-12 Zone Summary *(continued)*

Column name	Description
Zone Type	The type of zone
Status	Indicates the status of the zone as either active or inactive
Zone Set	The zone set each zone is in. Click the link to view the summary. See " Zone Set Summary " on page 708.
SAN Fabric	The name of the SAN Fabric
Members	The number of members in each Zone. See " Zone Members " on page 708.
Initiators	The number of initiators that send a request to the node
Targets	The number of receivers of the request

Zone Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Administration Reports > Zone Summary

Click a **Zone Name** link.

Table 21-13 Zone Details

Column name	Description
Zone Name	The name of the zone
Zone Type	The type of zone

Table 21-13 Zone Details *(continued)*

Column name	Description
Status	Indicates the status of the zone as either active or inactive
Caption	Short description of the zone
Description	Description of the zone
Identifier	The identifier contains several components that supply the name, WWN, and active/inactive status.
Other Zone Sub Types	Describes the zone sub-type
Other Zone Types	Describes the zone type
Read Only	Indicates if the zone configuration is changeable
Zone Sub Type	Provides additional details about the zone type
Defined	Indicates if the zone is defined, but not activated
SAN Fabric	The name of the SAN Fabric
Connected Initiators	The number of connected initiators that send a request to the node
Connected Targets	The number of connected receivers of the request

Total Zone Members:

Table 21-14 Zone Details

Column name	Description
Type	The type of the zone member
Name	The name of the zone member
WWN	The member's Worldwide Name
Status	The status of the zone member
Alias Names	A user-assigned alias names
Zone	Zone to which the member belongs
# of Related Zones	Other zones for which this member is also a member; click the numeric link to view additional details

Zone Set Summary

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Administration Reports > Zone Summary

Click a **Zone Set**link.

Table 21-15 Zone Set Summary

Column name	Description
Zone Name	The name of the zone.
Status	The status of the zone set.
Caption	A short description of the zone set.
Description	A description of the zone set.
Identifier	The identifier contains several components that supply the name, WWN, and active/inactive status.
# of Zones	The number of zones in the zone set. Click the link to view the summary. See " Zone Summary " on page 705.

Zone Members

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Administration Reports > Zone Summary

Click a **Members** link.

Table 21-16 Zone Members

Column name	Description
Type	The type of the zone member
Name	The name of the zone member
WWN	The member's Worldwide Name
Status	The status of the zone member
Alias Names	A user-assigned alias names
Zone	Zone to which the member belongs
# of Related Zones	Other zones for which this member is also a member; click the numeric link to view additional details

Bottom Ranking

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Capacity & Utilization > Bottom Ranking

Two versions of this report are available, depending on your selection in the Report Designer:

- Least Used Switches
- Least Used Ports

In addition, you can choose to report on the following metrics:

- Usage
- Errors
- Throughput

Example of a pie chart for the bottom-ranking, least used switches.

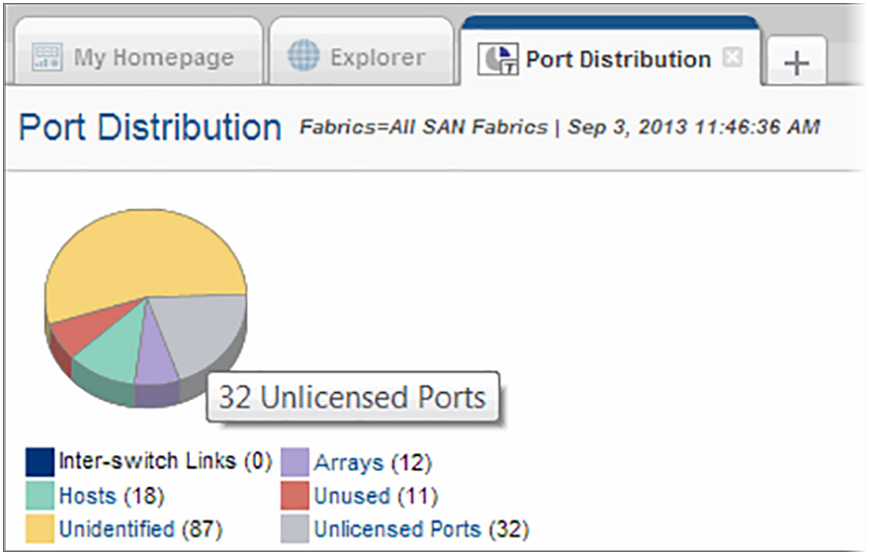


Port Distribution

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Capacity & Utilization > Port Distribution



SAN Resource Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Capacity & Utilization > SAN Resource Utilization

Table 21-17 SAN Resource Utilization

Column name	Description
Switch	
Name	The name of the Switch in the SAN Fabric. Click the link for details. See “ Switch Details ” on page 701.
WWN	Switch WWN
Port	The port on the switch that is connected to the array. This is a string, not a numeric field. Click the link for details. See “ Port Details ” on page 696.

Table 21-17 SAN Resource Utilization *(continued)*

Column name	Description
Array	
Array	The name of the array. Click the link to view the Array Capacity & Utilization Report.
WWN	The Worldwide name of the array. Click the link for details. See " Switch Details " on page 701.
Port	Array port
# of Hosts	The number of hosts. Click the link to view the Host Capacity & Utilization Summary.
Fan out Ratio	Indicates the number of hosts connected to the port. For example, 1-to-2 means 1 port to 2 hosts.

Top Ranking

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Capacity & Utilization > Top Ranking

Two versions of this report are available, depending on your selection in the Report Designer:

- Most Used Switches
- Most Used Ports

In addition, you can choose to report on the following metrics:

- Usage
- Errors
- Throughput

Example of a pie chart for the top-ranking, most used switches.

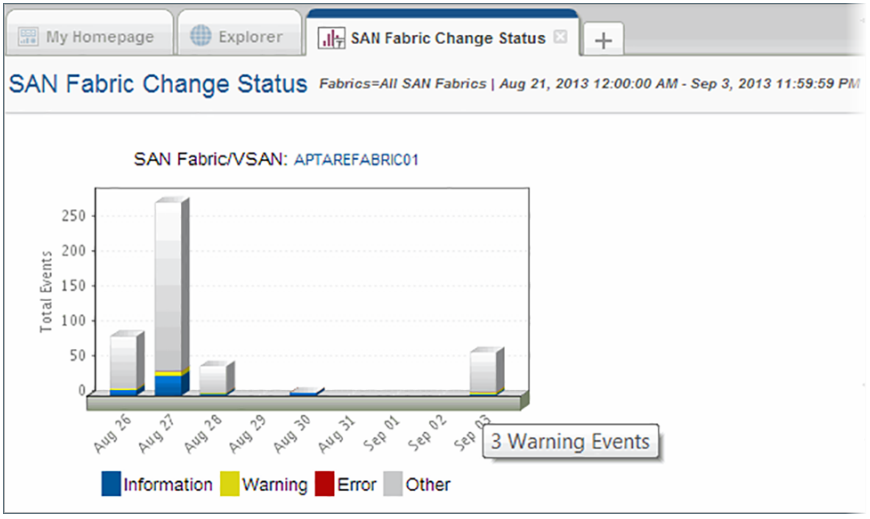


SAN Fabric Change Status Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Change Management Reports > SAN Fabric Change Status Summary



- Mouse over the bars to view specific statistics

- Click a bar segment to view the summary.
  See “ SAN Fabric Change Summary ” on page 714.

SAN Fabric Change Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Change Management Reports > SAN Fabric Change Summary

Table 21-18 SAN Fabric Change Summary

Column name	Description
Status	Indicates the status of the change: information, warning, alert
Severity	Indicates when immediate action is required
SAN Fabric	The name of the SAN Fabric
Date	Date the change occurred
Source	Indicates the status of the change: information, warning, alert
Element Type	Element type
Element	The element that changed
Port #	The number of the port. This is a string, not a numeric field.
Message	Messages associated with the change

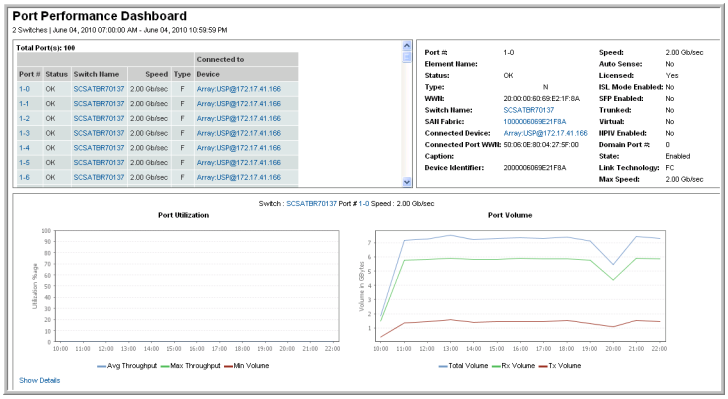
Port Performance Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Performance Reports> Port Performance Dashboard

This four-portlet dashboard initially presents a list of paths in the upper left portlet. The other three portlets remain blank until a path is selected.



Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Ports Portlet

Table 21-19 Ports Portlet

Column name	Description
Port #	The number of the port. This is a string, not a numeric field.
Status	The status of the port.
Switch Name	The name of the Switch. Click the link for details. See " Switch Details " on page 701.
Gb/Sec	The speed of the port.
Type	The type of port.
Licensed	Indicates if the port is licensed
Connected To	

Table 21-19 Ports Portlet *(continued)*

Column name	Description
Device	The device the port is connected to. Click the link to view the Array Capacity & Utilization Report.

Port Details Portlet

Table 21-20 Port Details Portlet

Column name	Description
Port #	The number of the port. This is a string, not a numeric field.
Element Name	Element name
Status	The status of the port; could indicate an error condition
Type	The type of port
WWN	The Worldwide name of the port
Switch Name	The name of the Switch. Click the link for details. See " Switch Details " on page 701.
SAN Fabric	The name of the SAN Fabric. Click the link for details. See " SAN Fabric Details " on page 698.
Connected Device	The device to which the port is connected
Connected Port WWN	The WWN of the port of the connected device
Caption	The short description of the port
Device Identifier	Unique identifier for the connected device
Gb/Sec	The speed of the port
Auto Sense	Indicates if the port can automatically determine the speed of attached devices
Licensed	Indicates if there is a license for this port
ISL Mode Enabled	Denotes inter-switch link capabilities
Trunked	The capability for multiple connections defined as a single link, on inter-switch links

Table 21-20 Port Details Portlet *(continued)*

Column name	Description
Virtual	Indicates if it is a virtual switch (Cisco).
NPIV Enabled	Indicates if a single Fibre Channel port is enabled to appear as multiple, unique ports.
Domain Port #	This port's number within the switch's domain
State	Enabled or Disabled. If the port is not connected, the state is not applicable.
Max Gb/Sec	The maximum speed of the port.
Port Index	The port's position in the slot or switch

Port Performance Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Performance Reports > Port Performance Summary

Table 21-21 Port Performance Summary

Column name	Description
Switch	The name of the Switch in the SAN Fabric. Click the link for details. See " Switch Details " on page 701.
Port #	The number of the port. This is a string, not a numeric field. Click the link for details. See " Port Performance Details " on page 718.
Gb/Sec	The speed of the port.
Received (Rx)	The number of bytes received within the report's interval
Transmitted (Tx)	The number of bytes transmitted within the report's interval

Table 21-21 Port Performance Summary *(continued)*

Column name	Description
Rx+Tx	The total number of bytes transmitted and received within the report's interval
Utilization %	The port's utilization percentage for the reporting interval
Peak Utilization %	The port's peak utilization for the reporting interval
Frames Received	The number of frames received within the report's interval
Frames Transmitted	The number of frames transmitted within the report's interval
CRC Errors	The number of CRC errors for the report's interval
Loss of Sync	The number of loss of sync errors
Loss of Signal	The number of loss of signal errors
Encoding Errors	The number of encoding errors

Port Performance Details

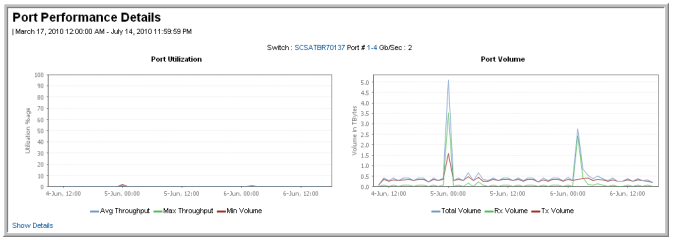
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Fabric Manager > Performance Reports > Port Performance Summary

Click a Port # link.



SAN Fabric Usage

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Billing and Usage> SAN Fabric Usage

SAN Fabric Chargeback policies can be configured to associate a cost with usage. This report shows the calculated costs for the report's interval.

Table 21-22 SAN Fabric Usage

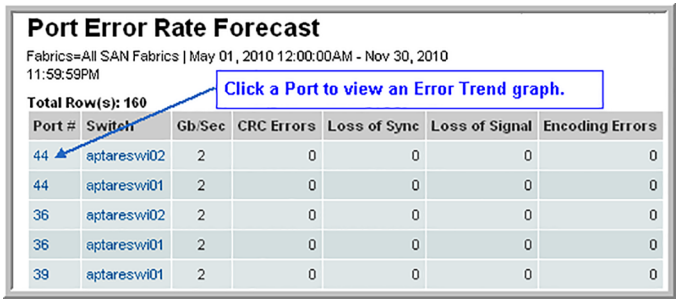
Column name	Description
Host	The host connected to the switch
Policy	The user-configured chargeback policy that defines the cost per TB
SAN Fabric	The name of the SAN Fabric
Switch Name	The name of the switch
Port #	The number of the port. This is a string, not a numeric field.
Gb/Sec	Port speed
Port Cost	Cost associated with the port
Total Usage (TB)	Usage for which charges apply
Usage Cost per TB	The cost defined in the chargeback policy
Total Cost	Total usage times cost per TB

Port Error Rate Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Forecast & Planning > Port Error Rate Forecast

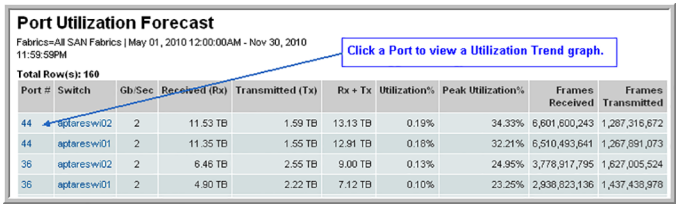


Port Utilization Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Fabric Manager > Forecast & Planning > Port Utilization Forecast



Chapter 22

Backup Manager Management Reports

This chapter includes the following topics:

- Using Backup Manager Reports for Analysis
- Anomalies Summary
- Backup Executive Summary
- Backup Server Performance Dashboard
- Client Summary Report
- Failed Backup Sources
- Consecutive Errors
- Data Domain Backup Summary
- Data Domain File Compression Detail
- Data Domain File Compression Summary
- Deduplication Size Savings Report (Historical)
- Error Log Summary
- HP Data Protector Session Summary
- HP Data Protector Session Details
- Job Duration Report
- Job Duration by Source

- Job Error Code
- Job Status Summary
- Job Status Summary by Client
- Job Summary Report
- Job Summary - Heterogeneous Backup Products
- Job Throughput by Client
- Avamar Job Summary Report (Homogeneous)
- Backup Exec Job Summary Report (Homogeneous)
- Cohesity DataProtect Job Summary Report (Homogeneous)
- Commvault Simpana Job Summary Report (Homogeneous)
- Commvault Simpana ESX Backup Summary
- HP Data Protector Job Summary Report (Homogeneous)
- IBM Spectrum Protect (TSM) Job Summary Report (Homogeneous)
- NetBackup Job Summary Report (Homogeneous)
- NetWorker Job Summary Report (Homogeneous)
- Oracle Recovery Manager (RMAN) Job Summary Report (Homogeneous)
- Veeam Backup & Replication Job Summary Report (Homogeneous)
- Job Details Report
- Avamar Job Details
- Commvault Simpana Job Details
- IBM Spectrum Protect (TSM) Job Details Report
- Veeam and RMAN Job Details Report
- Adding a Note to a Job
- Job Summary by Server
- Job Summary by Source
- Job Type Count
- Job Volume Summary Report

- Largest Backup Volume
- Primary Server Job Throughput
- Monthly Backup Summary
- NetBackup AIR-Replication Import Jobs
- NetBackup deduplication to MSDP savings
- NetBackup Media Server Job Throughput
- NetBackup SLP Status
- NetBackup SLP Details
- NetBackup SLP Status by Destination (AIR Activity)
- NetBackup SLP Status by Client
- NetBackup SLP Status by Image Copy
- NetBackup SLP Status by Image
- NetBackup SLP Status by SLP
- Running and Queued Job Summary
- NetBackup Running and Queued Jobs Summary
- NetWorker Running and Queued Jobs Summary

Using Backup Manager Reports for Analysis

Management Reports offer a variety of ways to view data. Using high-level reports, you can quickly summarize your backup status and then when needed, drill down to the details. The Portal reports will quickly replace your command-line tasks that you use to cobble together event status information.

These management reports can aid in troubleshooting to:

- diagnose malfunctioning backup servers
- determine scheduling conflicts
- analyze throughput issues
- identify storage media degradation

The following table lists several examples of exposure analysis tasks and the reports that support your diagnostic efforts.

Table 22-1 Using Backup Manager Reports for Analysis

What to do:	How to do it:
1. Determine the status of your backup jobs. ■ What failed and how severely? ■ Why did a job take so long? ■ Did another job hold it up?	See " Data Domain Backup Summary " on page 734. Drill down to details.
2. Diagnose backup scheduling issues.	**Job Duration Report** **Job Volume Summary** ■ From these reports, you can click on a bar to drill down to details in the Job Summary report.
3. Diagnose backup server problems.	Error Log Summary by Server
4. Diagnose client problems.	Consecutive Errors By Client ■ Drill down to investigate error conditions. ■ Use the Consecutive Errors by Clientreport in conjunction with the Error Log Summary, which displays specific error code messages.

Anomalies Summary

This report displays a summary anomalies detected for various clients during backup. It appears as a drill down report of the Anomaly Count chart that is displayed on the Backup Primary Server Overview and Ransomware Risk Assessment dashboards. Anomalies are reported provided anomaly detection is configure in NetBackup Web UI (see *Configure anomaly detection settings* section of the *NetBackup Security and Encryption Guide*) and **Security Details** probe is enabled in the NetBackup Policy for the primary server

This report is located here:

Backup Manager > **Management Reports** > **Anomalies Summary**

Table 22-2 Anomalies Summary

Column name	Description
Client	Name of the backup up host. The link drills down to Host Details report. See " Host Details " on page 838.
Server	Name of the Server for the backup job. The link drills down to Host Details report. See " Host Details " on page 838.
Job Id	ID of the backup job. The link drills down to the Job Details report. See " Job Details Report " on page 770.
Anomaly Score	Severity of the anomaly. The score is higher if the severity of the anomaly is more. The link drills down to the Anomaly Details report.
Anomaly Severity	Indicates anomaly severity as High, Medium, or Low, based on the Anomaly Score.
Anomaly Type	Summary of anomalies that are notified for this job.
Anomaly Detection Time	Date and time when the anomaly was first detected.
Job Type	Code assigned to the backup job type.
Policy	Name of the backup policy. Click the link to drill down to Backup Policy Details report. See " Backup Policy Details " on page 968.
Schedule	Schedule name configured for this job. The schedule name uniquely identifies the schedules available in the subsystem.
Job Start Time	Start time of the backup job.
Job Finish Time	End time of the backup job.
False Positive	Indicates whether the false positive is reported for this anomaly or not (Yes/No).

Anomaly Details

This report appears as a drill down report of the Anomalies Summary report. The report displays various details specific to the selected anomaly.

Table 22-3 Anomaly Details

Column name	Description
Client	Name of the backup up host. The link drills down to Host Details report. See " Host Details " on page 838.
Vendor Product	The name of the backup product for which anomaly detection is performed.
Job Id	ID of the backup job.
Server	Name of the Primary Server for the backup job. The link drills down to Host Details report. See " Host Details " on page 838.
Policy	Name of the backup policy. Click the link to drill down to Backup Policy Details report. See " Backup Policy Details " on page 968.
Create Date	Anomaly detection date.
Last Updated	Date and time when the anomaly record was last updated.
Anomaly Summary	
Score	Severity of the anomaly. The score is higher if the severity of the anomaly is more.
Severity	Indicates anomaly severity as High, Medium, or Low, based on the Anomaly Score.
Count	The number of anomalies detected for the job.
False Positive	Indicates whether the false positive is reported for this anomaly or not (Yes/No).
Anomaly Parameters	

Table 22-3 Anomaly Details *(continued)*

Column name	Description
File Count	If this field value is available, then there is anomaly detected on backup taken on number of files.
Image Size	If this field value available, then there is anomaly detected on Image Size.
Data Transferred	If this field value available, then there is anomaly detected on Data Transferred.
Duplication Ratio	If this field value available, then there is anomaly detected on Deduplication Ratio.
Total Time	If this field value available, then there is anomaly detected on Time Taken by Backup Job.

Backup Executive Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Backup Executive Summary

This report provides an Executive Summary overview of the status of backup jobs within defined backup windows. In the Report Designer, a backup window can be selected to narrow the scope of the report.

Table 22-4 Backup Executive Summary

Column name	Description
Backup Window: Start Date	Start date and time of the backup window
Backup Window: Finish Date	End date and time of the backup window

Table 22-4 Backup Executive Summary *(continued)*

Column name	Description
Jobs: Total	Total number of backup jobs that executed within the backup window. Click on the link to view the Job Summary Report.
Jobs: Successful	Total number of successful backup jobs that executed within the backup window. Click on the link to view the Job Summary Report for successful backup jobs for this time period.
Jobs: Failed	Total number of failed backup jobs that executed within the backup window. Click on the link to view the Job Summary Report for failed backup jobs for this window.
Jobs: Success Rate	The success percentage, calculated as Successful jobs divided by Total jobs
Clients: Total	Number of unique clients backed up within the backup window
Clients: Successful	Number of unique clients that were successfully backed up within the backup window.
Clients: Failed	Number of unique clients that failed to be backed up within the backup window. Click on the link to view the Failed Backup Sourcesreport for this time period.
Clients: Success Rate	The success percentage, calculated as Successful clients divided by Total clients
Objects: Total	Number of unique objects backed up within the backup window
Objects: Successful	Number of unique objects that were successfully backed up within the backup window.
Objects: Failed	Number of unique objects that failed to be backed up within the backup window. Click on the link to view the Failed Backup Sourcesreport for this time period.
Objects: Success Rate	The success percentage, calculated as Successful objects divided by Total objects

Backup Server Performance Dashboard

This dashboard is designed to help NetBackup IT Analytics administrators to identify and analyze the appliances or backup servers experiencing slowness and bad performance. It consists of eight portlets that collect data from the backup probes and various host resources. Each portlet is a report in itself and is described below.

This report is located here:**Backup Manager** > **Management Reports** > **Backup Server Performance Dashboard**

Job Throughput by Client

This report shows the average throughput calculated per minute and aggregated for an hour for a given host. The throughput is averaged based on how the jobs are grouped. The report is designed to evaluate a single host; however, if you choose to run it for multiple hosts, it displays a summarized chart of clients from all the hosts. Currently, the report displays the chart of 20 clients causing the highest throughput, so that administrators can focus on specific clients to troubleshoot those responsible for the high throughput.

See “Job Throughput by Client” on page 752.

Job Status Summary

The Job Status Summary report shows a bar chart that aggregates jobs that succeeded, failed, or produced warning messages. This high-level summary provides drill down access to the Job Summary report, where you’ll find specific status messages and scheduling information.

See “ Job Status Summary ” on page 747.

Host Load Average

This report displays the chart based on values derived from hosts running only on UNIX operating systems.

This chart displays historical load averages derived from the uptime command. The uptime command displays the current time, how long the system is running, the number of active logged in users, and the host load averages in the past 1, 5, and 15 minutes.

Host load average is the average of the processes in either pre-run or uninterruptible state. A process in pre-run state either already using the CPU or waiting to use it. However, a process in uninterruptible state is awaiting the use of CPU, which means it is waiting for an I/O access, such as waiting for disk. The averages displayed in the report are taken over the three time intervals mentioned before, but are not normalized for the number of CPUs in a system.

See “Host Load Average” on page 554.

Host CPU Utilization

The report displays a chart of the average CPU usage of the backup host over the time duration specified for the report. In addition, it also plots an additional line indicating the maximum CPU usage for the specified time interval.

Even though the report is designed to evaluate the CPU utilization of a single host, you may view a chart for multiple hosts. In that case, the chart will display values based on the average utilization of all the hosts.

See “Host CPU Utilization” on page 555.

Host Memory Utilization

The report displays a chart of the average memory consumption of the backup host over the time duration specified for the report. In addition, it also plots an additional line indicating the total memory usage for the specified time interval. Although the report is designed to evaluate the memory consumption of a single host, you may view a chart for multiple hosts. In that case, the chart will display values based on the average memory consumption of all the hosts.

See “Host Memory Utilization” on page 556.

Host Memory Swap

The report displays the swap in and swap out of memory per minute for each host within the duration specified while generating the report. This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances.

See “Host Memory Swap” on page 557.

Host Disk reads/writes

This report summarizes the reads and writes of all the disks connected to the host. The values are calculated in Kilobits/second and displayed on the chart for the time duration specified while generating the report. You can evaluate the disk utilization using this report and resolve slowness of applications running on the host.

See “ Host Disk reads/writes” on page 557.

Host Network Packets sent/received

This report displays a chart of summarized values of the packets sent and received over the host network. The chart is a plot of summarized values obtained from all the NICs in the host network. This chart can help you minimize performance issues arising from network congestions.

See “Host Network Packets sent/received” on page 558.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Client Summary Report

This report summarizes several parameters for each client and displays the backup activities. You can use drilldown links to view client detail reports for many of the columns in the table. The drilldown reports provide information about the client and answer questions such as:

- What is the server name, IP address, and operating system type?
- How much volume for a client is backed up?
- How many successful, partially successful, and failed jobs for the client?
- What is the number of restore jobs that run for the client?
- What is the average job throughput?

Note: This report is equivalent to the Client Summary Dashboard of NetBackup OpsCenter reports.

This report is located here: **Backup Manager** > **Management Reports** > **Client Summary Report**.

Table 22-5 Client Summary Report

Column name	Description
Server	Name of the server for the backup job.
Client	Name of the backup host.
Size (GB)	Backup size in GB.
Duration/Sec	Time duration in seconds taken to complete the backup.
KiB/Sec	Backup rate in Kilobytes per second.
Total Jobs	Count of backup jobs. You can drill down from this column to view job details.
Successful Jobs	Number of successful jobs. You can drill down from this column to view job details.
Partial Jobs	Number of partially successful jobs. You can drill down from this column to view job details.
Failed Jobs	Number of failed jobs. You can drill down from this column to view job details.
Success Rate %	Success rate in percentages.

Failed Backup Sources

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > Backup Executive Summary

Click a **Failed Client or Failed Object** link.

This report provides a list of the clients/objects with failed backups for the selected time period, with drill down access to additional details.

Table 22-6 Failed Backup Sources

Column name	Description
Source	Name of the backup client or backup source object. Click the link to view a detail report. If the backup source is a known client, the drill down is to the Host Details report. For other backup sources, the drill down is to the Object Details report.
# of Jobs	The number of failed jobs for the client within the time period. Click the link to view the Job Summary Report.
First Attempted	The first attempted backup that failed during the time period.
Last Attempted	The last attempted backup that failed during the time period.

Consecutive Errors

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports> Consecutive Errors

Use this report to proactively manage clients, servers and objects. Consecutive errors may portend hardware reliability issues. Using this report you can identify not only consecutive errors, but also the number of consecutive days that the errors persisted.

The default order for this listing is descending order, with the source (clients/objects) having the most consecutive errors listed first. Correlate the number of consecutive errors with the number of consecutive days so that you can disregard errors that appear to be intermittent.

The Consecutive Errors report provides the following information:

Table 22-7 Consecutive Errors

Column name	Description
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
IP Address	Source's IP address on the LAN.
Server Name	The server is the system that is managing the storage devices. Click any server name to view the details. See " Host Details " on page 838. Click the Host Group link to determine the host groups to which this server belongs.
# of Consecutive Errors	# of consecutive errors logged within the designated time frame. Click on the number to launch the related Job Summary report where you'll find a detailed listed of the errors.
# of Consecutive Days	# of days in which the client's consecutive errors were logged.
First/Last Error Dates	Date of the first and last errors, within the specified time frame, gleaned from logs. The range between the first and last error dates gives some insight into the severity of the problem. Click on each of the error dates to generate two separate Job Summary reports so that you can compare the snapshots of the start and end of the error events.

Data Domain Backup Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Data Domain Backup Summary

This report lists the NetBackup systems that have client backup data. The Data Domain system acts as an external disk pool. Note that LSU (logical storage unit) data is collected every six hours.

When Data Domain servers are selected for the report scope, the primary server clients are also included in the report.

Table 22-8 Data Domain Backup Summary

Column name	Description
System	Name of the Data Domain appliance.
Client	Name of the client backed up by the primary server.
Server	Name of the primary server
LSU	Logical storage unit.
Total Backups	Total number of jobs. Total number of jobs for the client, job type and logical storage unit.
Size of Backup	Total number of kilobytes. Total number of kilobytes for the client, job_type and logical storage unit.
Last Backup	Finish date of last backup. Finish date of the last successful or partial backup.
Last Full Backup	Finish date of the last full backup. Finish date of the last successful or partial full backup.

Data Domain File Compression Detail

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Data Domain File Compression Detail

View Data Domain file pre- and post-compression values.This information can be instrumental in determining efficient storage strategies and identifying storage that can be reclaimed, thereby reducing archive storage expenses. This data can be used to identify clients with inefficient de-duplication ratios, highlighting where de-deduplication is not an effective approach for certain backed-up files. For example, some hosts may be running database applications that are constantly producing unique bits of data. These hosts can consume much of the expensive Data Domain storage. Data Domain collection now can identify the largest offenders, which can then be moved to less expensive storage to avoid paying premium rates for de-duplication.

See "Data Domain File Compression Summary" on page 737.

Configure EMC Data Domain collection to capture image-level compression rates. Aggregated global and local compression rates for all backup images can be collected for all active Data Domain Server MTrees connected (via DDBOOST) to Primary Servers. These Primary Servers must have an active Data Collector that has successfully completed an initial data collection.

These values are relevant as of the last successful data collection cycle.

Table 22-9 Data Domain File Compression Detail

Column name	Description
System	Name of the Data Domain appliance. Click to view the Data Domain System Details.
Client	Displayed for NetBackup. Backup client name. If backup source is a known host, host name is a link to Host Details. Other backup sources are displayed with link to Object Details.
Source	Backup source name. If backup source is a known host, host name is a link to Host Details. Other backup sources are displayed with link to Object Details.
Source Type	Defines the source type such as resource pool, datastore or a folder.
Server	Name of the host to display in Portal reports. Click the link to display the Host Details or Object Details.
Backup ID	The ID that identifies a specific job.

Table 22-9 Data Domain File Compression Detail *(continued)*

Column name	Description
Date	The date and time that a backup job wrote to the Data Domain appliance.
Storage Unit/Volume	Displayed for NetBackup. Name of the Storage Unit/Volume used in the backup.
Job Finished Date	The date and time this job finished. Veritas NetBackup, Veritas Backup Exec, IBM Spectrum Protect (TSM), Commvault Simpana, HP Data Protector, EMC Avamar, Generic Backup: The date and time this job finished. Veritas NetBackup: This is expressed in the time zone of the primary server. IBM Spectrum Protect (TSM): If the dates are missing the century, 2000 is added to the date.
Policy Name	Name of the policy used in the backup.
Original Size	Amount of logical capacity associated with the backup.
Global Compressed	Deduplication ratio: (raw_bytes / pre_local_compression_size).
Global Deduplication %	Percentage of Global compression (100 - (Global Compressed/Original Size)*100).
Local Compressed	Local compression algorithm: (pre_local_compression_size / post_local_compression_size).
Local Compression %	Percentage of Local compression (100 - (Local Compressed/Original Size)*100)
Metadata	System overhead metadata.
Bytes to Storage Ratio	Bytes to Storage Ratio: bytes/storage_used (local_compression + metadata).
Backup Product	Name of the backup product.
Backup ID	Uniquely identifies the backup job (or similar). For NetBackup, it is the concatenation of the client host name and a long integer representing the starting time of the job, this uniquely identifies the job. For Avamar, it is the client ID (CID) that identifies the client, but not the job.
Path Name	Path name.
Last Collected from Data Domain	Date and time the Data Domain file level compression data was last updated in the Portal database.

Data Domain File Compression Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Data Domain File Compression Summary

View Data Domain file pre- and post-compression values summarized for a client, with aggregated rates. This information can be instrumental in determining efficient storage strategies and identifying storage that can be reclaimed, thereby reducing archive storage expenses. This data can be used to identify clients with inefficient de-duplication ratios, highlighting where de-deduplication is not an effective approach for certain backed-up files. For example, some hosts may be running database applications that are constantly producing unique bits of data. These hosts can consume much of the expensive Data Domain storage. Data Domain collection now can identify the largest offenders, which can then be moved to less expensive storage to avoid paying premium rates for de-duplication.

See "Data Domain File Compression Detail" on page 734.

Configure EMC Data Domain collection to capture image-level compression rates. Aggregated global and local compression rates for all backup images can be collected for all active Data Domain Server MTrees connected (via DDBOOST) to Primary Servers. These Primary Servers must have an active Data Collector that has successfully completed an initial data collection.

These values are aggregated from relevant collections of job details as of the last successful data collection cycle.

Table 22-10 Data Domain File Compression Summary

Column name	Description
System	Name of the Data Domain appliance. Click to view the Data Domain System Details.
Source	Name of the backup source or object backed up by the job during the specified interval. If backup source is a known host, host name is a link to Host Details. Other backup sources are displayed with link to Object Details.

Table 22-10 Data Domain File Compression Summary *(continued)*

Column name	Description
Source Type	Defines the source type such as host, resource pool, datastore or a folder.
Server	The name is derived from the host's DNS name.
# of Images	Click the number of images to view compression details. Click to Data Domain File Compression Detail.
Original Size	Amount of logical capacity associated with backups from this host to the appliance.
Global Compressed	The deduplication ratio: (raw_bytes / pre_compression_size).
Global Deduplication %	Percentage of Global compression (100 - (Global Compressed/Original Size)*100).
Local Compressed	The local compression algorithm: (pre_compression_size / pre_compression_size).
Local Compression %	Percentage of Local compression (100 - (Local Compressed/Original Size)*100)
Metadata	System overhead.
Bytes to Storage Ratio	Bytes to Storage Ratio: bytes/storage_used (local_compression + metadata).

Deduplication Size Savings Report (Historical)

This is a deduplication report created using Media Server Deduplication Pool.

This report can be accessed from Backup Manager > Management Reports > Deduplication Size Savings Report..

This report shows the total backup size saved in GB due to deduplication for one or more NetBackup Primaries. Also, from this report you can drill down and see the Deduplication Jobs details. This report is equivalent to the OpsCenter Deduplication size Savings Report (Historical).

Error Log Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

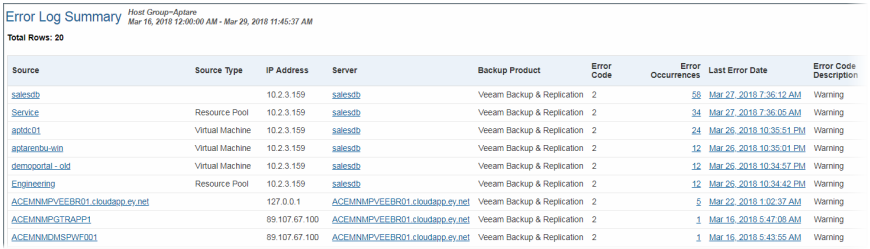
Backup Manager > Management Reports > Error Log Summary

Backup Manager > Management Reports > Error Log Summary by Policy

Backup Manager > Management Reports > Error Log Summary by Server

This report has a different version, for each of the following:

Error Log Summary

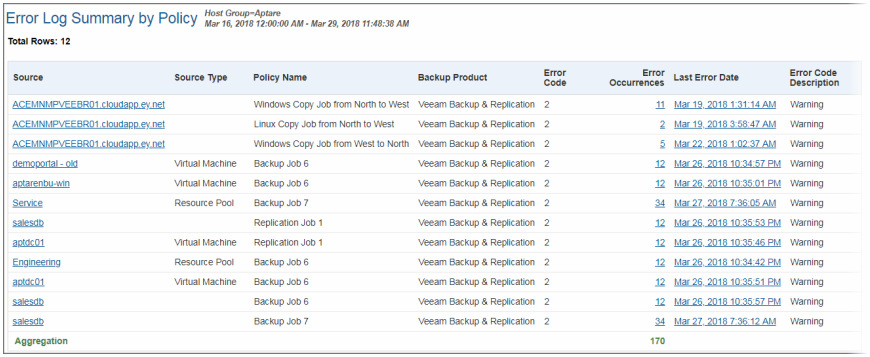


Error Log Summary *Host Group=Aptare Mar 16, 2018 12:00:00 AM - Mar 29, 2018 11:45:37 AM*

Total Rows: 20

Source	Source Type	IP Address	Server	Backup Product	Error Code	Error Occurrences	Last Error Date	Error Code Description
salesdb		10.2.3.159	salesdb	Veeam Backup & Replication	2	58	Mar 27, 2018 7:36:12 AM	Warning
Service	Resource Pool	10.2.3.159	salesdb	Veeam Backup & Replication	2	34	Mar 27, 2018 7:36:05 AM	Warning
aptdc01	Virtual Machine	10.2.3.159	salesdb	Veeam Backup & Replication	2	24	Mar 26, 2018 10:35:51 PM	Warning
aptarenbu-win	Virtual Machine	10.2.3.159	salesdb	Veeam Backup & Replication	2	12	Mar 26, 2018 10:35:01 PM	Warning
demoportal - old	Virtual Machine	10.2.3.159	salesdb	Veeam Backup & Replication	2	12	Mar 26, 2018 10:34:57 PM	Warning
Engineering	Resource Pool	10.2.3.159	salesdb	Veeam Backup & Replication	2	12	Mar 26, 2018 10:34:42 PM	Warning
ACEMNMPVEEBR01.cloudapp.ey.net		127.0.0.1	ACEMNMPVEEBR01.cloudapp.ey.net	Veeam Backup & Replication	2	5	Mar 22, 2018 1:02:37 AM	Warning
ACEMNMPGTRAPP1		89.107.67.100	ACEMNMPVEEBR01.cloudapp.ey.net	Veeam Backup & Replication	2	1	Mar 16, 2018 5:47:08 AM	Warning
ACEMNMDMSPWF001		89.107.67.100	ACEMNMPVEEBR01.cloudapp.ey.net	Veeam Backup & Replication	2	1	Mar 16, 2018 5:43:55 AM	Warning

By Policy

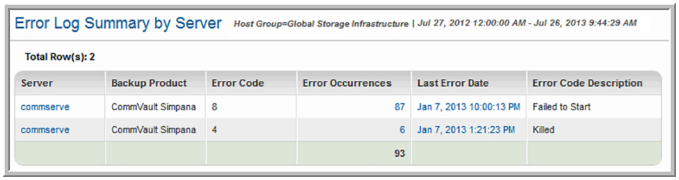


Error Log Summary by Policy *Host Group=Aptare Mar 16, 2018 12:00:00 AM - Mar 29, 2018 11:48:38 AM*

Total Rows: 12

Source	Source Type	Policy Name	Backup Product	Error Code	Error Occurrences	Last Error Date	Error Code Description
ACEMNMPVEEBR01.cloudapp.ey.net		Windows Copy Job from North to West	Veeam Backup & Replication	2	11	Mar 19, 2018 1:31:14 AM	Warning
ACEMNMPVEEBR01.cloudapp.ey.net		Linux Copy Job from North to West	Veeam Backup & Replication	2	2	Mar 19, 2018 3:58:47 AM	Warning
ACEMNMPVEEBR01.cloudapp.ey.net		Windows Copy Job from West to North	Veeam Backup & Replication	2	5	Mar 22, 2018 1:02:37 AM	Warning
demoportal - old	Virtual Machine	Backup Job 6	Veeam Backup & Replication	2	12	Mar 26, 2018 10:34:57 PM	Warning
aptarenbu-win	Virtual Machine	Backup Job 6	Veeam Backup & Replication	2	12	Mar 26, 2018 10:35:01 PM	Warning
Service	Resource Pool	Backup Job 7	Veeam Backup & Replication	2	34	Mar 27, 2018 7:36:05 AM	Warning
salesdb		Replication Job 1	Veeam Backup & Replication	2	12	Mar 26, 2018 10:35:53 PM	Warning
aptdc01	Virtual Machine	Replication Job 1	Veeam Backup & Replication	2	12	Mar 26, 2018 10:35:46 PM	Warning
Engineering	Resource Pool	Backup Job 6	Veeam Backup & Replication	2	12	Mar 26, 2018 10:34:42 PM	Warning
aptdc01	Virtual Machine	Backup Job 6	Veeam Backup & Replication	2	12	Mar 26, 2018 10:35:51 PM	Warning
salesdb		Backup Job 6	Veeam Backup & Replication	2	12	Mar 26, 2018 10:35:57 PM	Warning
salesdb		Backup Job 7	Veeam Backup & Replication	2	34	Mar 27, 2018 7:36:12 AM	Warning
Aggregation					170		

By Server



Error Log Summary by Server *Host Group=Global Storage Infrastructure | Jul 27, 2012 12:00:00 AM - Jul 26, 2013 9:44:29 AM*

Total Row(s): 2

Server	Backup Product	Error Code	Error Occurrences	Last Error Date	Error Code Description
commserve	CommVault Simpana	8	87	Jan 7, 2013 10:00:13 PM	Failed to Start
commserve	CommVault Simpana	4	6	Jan 7, 2013 1:21:23 PM	Killed
			93		

The Error Log Summaryreports backup failure details that have been gleaned from backup error logs. Use this report in conjunction with the Largest Backup Volume report to mitigate future problems by determining which types of errors occurred with the greatest frequency.

See " Largest Backup Volume " on page 793.

The Error Code Description typically provides sufficient details to start your troubleshooting process. For example, a NetBackup error code of 196 would be an indicator that there are not enough resources available and a Virtual Tape Library (VTL) may be needed. The 196 indicates that the client backup cannot begin because drives are not available. Refer to the backup product's reference guides for additional guidance.

This report displays a number of links that enable drill down access to additional details, as described in the following table.

The Error Log Summaryprovides the following information:

Note: Some of the following fields do not appear in every Error Log Summary. Each of the three report types--Client, Server, Policy--contains only the fields that make sense for its group.

Table 22-11 Error Log Summary

Column name	Description
Server	The server is the system that is managing the storage devices.
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.

Table 22-11 Error Log Summary *(continued)*

Column name	Description
IP Address	The IP address of the client, shown only in the Error Log Summary By Client report.
Policy Name	If you selected Group By Policy through the Report Designer, the policy name appears in this column. This column does not appear if you selected Group By Management Server.
Backup Product	Indicates the backup software running on the server where the error occurred. Refer to that product's documentation to correlate the error code to diagnose problems. ■ Commvault Simpana ■ HP Data Protector ■ Veritas Backup Exec ■ IBM Spectrum Protect (TSM) ■ Veritas NetBackup
Error Code	Numeric code reported by the backup product. Refer to the backup product's reference guides for troubleshooting advice.
Error Occurrences	Each row in the table displays the number of occurrences for each error code within the report's selected time frame. Click on any error occurrence to go to the Job Summary report, which lists the events that correspond to that particular error code.
Last Error Date	Click on the Last Error Date to go to the Job Summaryreport, which displays the most recent error events only, including skipped files. If the dates are missing the century, 2000 is added to the date.
Error Code Description	Each row in the table includes a detailed error description that provides you with valuable information. Since the data is derived from backup job error logs, this description is backup-product dependent.

HP Data Protector Session Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > HP DP Session Summary

The Session Summary report provides an overview of the HP Data Protector backup/restore sessions.

Table 22-12 HP Data Protector Session Summary

Column name	Description
Session Name	Backup/restore session name. Click the link for details. See " HP Data Protector Session Details " on page 742.
Event Type	unknown, backup, restore
Backup Type	incremental, differential, transaction, full, or delta
Specification	Name of the HP Data Protector backup specification
Start Time	Backup/restore start time
End Time	Backup/restore end time
Success Percentage	% of backup/restore session that was successful
MB	Amount of data backed up
Status	Session Status: unknown, completed, completed/failures, or failed
Tapes	The tapes used for the backup/restore session

HP Data Protector Session Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > HP DP Session Summary

Click a **Session Name**.

The Session Details report provides details for HP Data Protector backup/restore sessions.

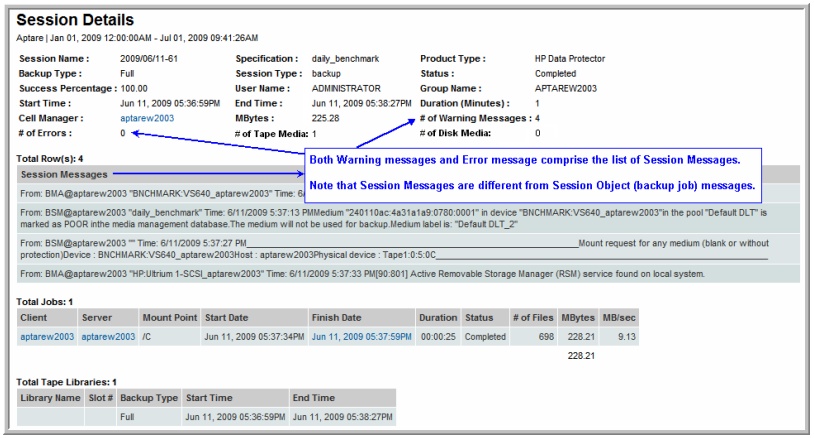


Table 22-13 HP Data Protector Session Details

Column name	Description
Session Name	Backup/restore session name
Backup Type	incremental, differential, transaction, full, or delta
Success Percentage	% of backup/restore session that was successful
Start Time	Backup/restore start time
End Time	Backup/restore end time
Cell Manager	The name of the Cell Manager server. Click the link for details. See “ Host Details ” on page 838.
# of Errors	Indicates the number of session errors. These error messages will be included in the Session Message table.
Specification	The HP Data Protector backup specification links to the HP Data Protector Backup Specification Summary.
User Name	Session user name
MBytes	Amount of data backed up by this session
# of Tape Media	The number of tape media associated with this session

Table 22-13 HP Data Protector Session Details *(continued)*

Column name	Description
Product Type	HP Data Protector
Status	Session Status: unknown, completed, completed/failures, or failed
Group Name	Group Name
Duration	Length of the session
# of Warning Messages	Indicates the number of session warning messages. If there are any warning messages, they will be listed in the table below the session details.
# of Disk Media	Number of disk media associated with this session

Job Duration Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

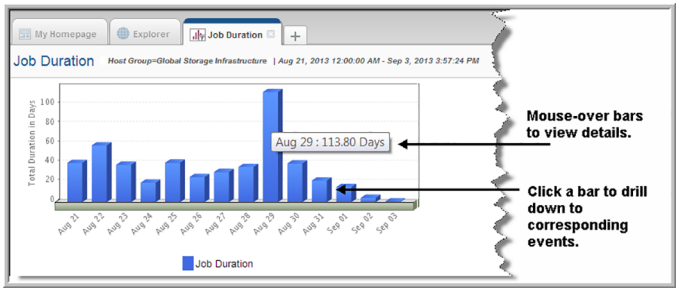
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Duration

The Job Duration report combines into one view, information about both backup and restore duration. If, on a specific day, there were two backup events that completed after 1 hour and 3 hours respectively, and one restore completed after .5 hours, the bar reports a duration of 4.5 hours.

You can also report on a single event type:

- All Backup Events
- Full Backups
- Incremental Backups
- Restores



The Job Durationreport provides the following information:

Table 22-14 Job Duration Report

Column name	Description
Drillable Bars in Chart	Click on any of the bar sections to go to a summary. See “ Job Summary Report ” on page 750.
Pop-up Details	Mouse-over any of the bars or symbols on the chart to display pop-ups that provide details. This pop-up lists the total duration for the selected time period.
Duration	The total amount of time the backup and restore events took to complete. If on a specific day there were two backup events that completed after 1hr and 3hrs respectively, and one restore completed after .5 hours, the bar reports a duration of 4.5 hours.

Use this report as a starting point to determine why certain days show excessive backup durations.

To diagnose excessive backup durations:

1. View the Job Duration bar chart, looking for excessive duration times.
2. See “ Job Summary Report ” on page 750.
3. In the Job Summary Report, click on the Duration column header to sort the data by Duration. Now you’ll see excessive client usage at the top of the list. You might also sort by MBytes/Sec to identify throughput issues.
4. In the Job Summary Report, click on the **Finish Date** of a suspect job to display the details.

   See “ Job Details Report ” on page 770.

By using a combination of reports, you can figure out if some jobs held up others, if there are server problems that require your attention, or if consecutive errors are symptomatic of storage device failure or media degradation.

Job Duration by Source

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Duration by Source

The Job Duration by Source report displays the total duration of backup and restore jobs sorted by Source for a specific period of time. Information can be filtered on vendor job types and backup/restore event types.

Job Error Code

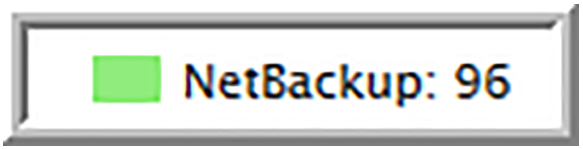
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

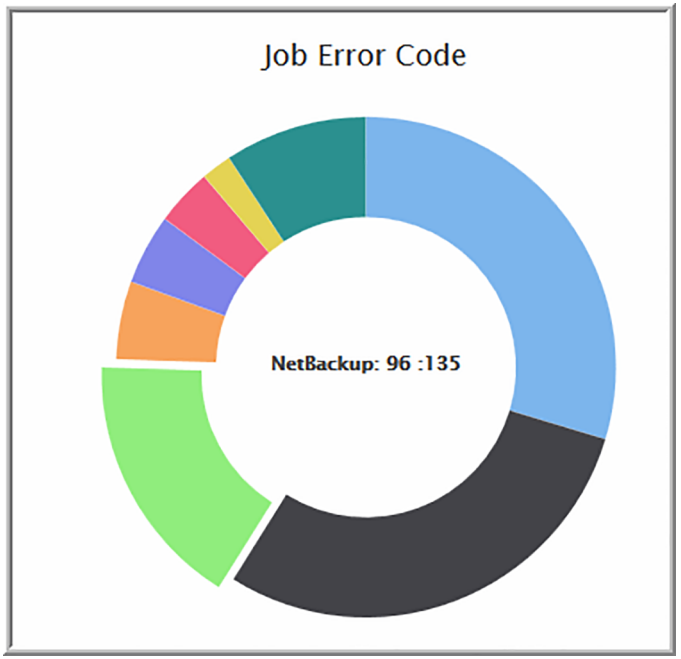
Backup Manager > Management Reports > Job Error Code

The Job Error Code report displays a donut chart that aggregates job error codes. Use this chart for a high-level view of backup and/or restore errors, with access to the relevant details. This high-level summary provides drill down access to the Job Summary report, where you'll find specific status messages and scheduling information.

- Each donut sector represents an error code with the number of error occurrences shown in a tooltip in the following format: <code>:<count>. The total count of all errors is displayed in the center of the donut.
- The Legend for the donut chart has the following format: Color Backup Product: Error Code



- Click a legend item to remove that sector from the rendered chart. The legend items can be toggled to display/remove a sector.
- When you click the donut sector, additional details display the error code and the number of occurrences of that error. The drill down from the sector displays the Job Summary report where you can hover your mouse over the Exit Code to view the reason associated with the code.



Job Status Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Status Summary

The Job Status Summary report shows a bar chart that aggregates jobs that succeeded, failed, or produced warning messages. This high-level summary provides drill down access to the Job Summary report, where you'll find specific status messages and scheduling information. This report is also a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances.

See "Backup Server Performance Dashboard" on page 728.

When you generate this report, several Report Designer options enable you to narrow the scope of the report and also to define how the data will be displayed. Refer to the NetBackup IT Analytics User Guide for Capacity Manager Primary Settings and Backup Manager Primary Settings. In particular, note the **Parent Job Only, Ignore Retries,** and **Backup Window** settings that are available for this report.

Use this procedure to track down backup failure details:

1. View the Job Status Summary report.
2. Click on the top section of a bar in the bar chart to drill down to details.

   The top portion of the bar indicates backup error events, the middle section reports event warnings, and the bottom represents successful backups. Note that often the partial backups represent your biggest exposure, so drilling down on the warning events can help you identify and resolve issues.

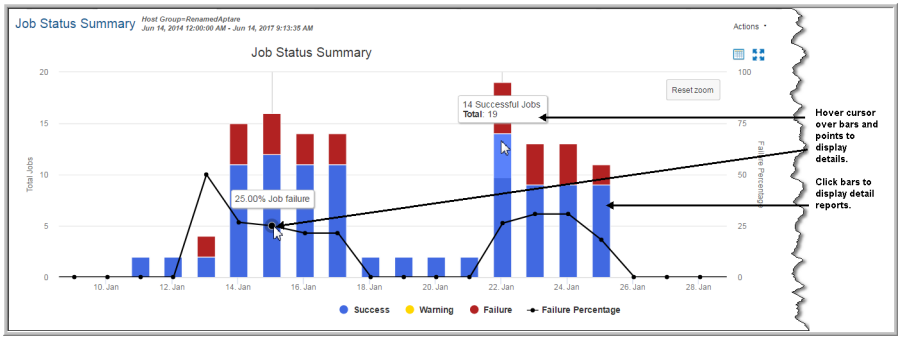
   See " Job Summary Report " on page 750.
3. In the Job Summary report, view the Activity Type column for backup log messages.
4. In the Job Summary report, click on the **Finish Date** to view the details.

   See " Job Details Report " on page 770.

   , This report shows additional message-log entries. Be sure to scroll down to the bottom of the list to find backup attempts and notes related to the job.

Related reports:

- See " Job Summary Report " on page 750.
  This tabular report provides details related to the data represented in the bars shown in the Job Status Summary chart. You can drill down to the job summary details from the Job Status Summary bar chart.
- See "Backup Status SLA" on page 888.
  In addition to job status, this report provides indicators showing a trend line of event successes.



The Job Status Summary provides the following information:

Table 22-15 Job Status Summary

Column name	Description
Drillable Stacked Bar Chart	Click on any of the bar sections to go to a summary. See " Job Summary Report " on page 750. for the backups corresponding to the time period and event status for the bar that you selected. The top portion of each bar indicates backup error events, the middle section reports event warnings, and the bottom section represents successful backups.
Pop-up Details	Mouse-over any of the bars or symbols on the chart to display pop-ups that provide details.
Job Failures	Roll-over the event failures to view as a percentage of the total events for the selected time period.
Job Success	Displays the number of total job successes as well as the percentage.

Job Status Summary by Client

The report displays status summary of backup jobs for the selected clients. The table displays totals for successful, partially successful, and failed jobs for each client that is in scope.

To view the details of a job, click the **Totals** bar displayed in the report. The report displays details of the selected job in a drill down report format.

It is recommended to run this report on specific groups of clients by modifying the scope selector.

This report is equivalent to the NetBackup OpsCenter report(s): Job Success by Client

To view the report, navigate to **Backup Manager** > **Management Reports** > **Job Status Summary**.

Job Summary Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Summary

When you generate this report, several Report Designer options enable you to narrow the scope of the report and also to define how the data will be displayed. Use the Scope Selector to filter the report to just the jobs you want to analyze. In particular, note the **Parent Job Only, Ignore Retries,** and **Backup Window** settings that are available for this report. The **Backup Policy Name Filter** enables you to enter a backup policy name to analyze job status for a specific policy

To view the details available in vendor-specific version of this report, see the following:

- See “Job Summary - Heterogeneous Backup Products” on page 751.
- See “ Avamar Job Summary Report (Homogeneous) ” on page 754.
- See “ Backup Exec Job Summary Report (Homogeneous) ” on page 755.
- See “ Commvault Simpana Job Summary Report (Homogeneous) ” on page 758.
- See “ HP Data Protector Job Summary Report (Homogeneous) ” on page 761.
- See “ NetBackup Job Summary Report (Homogeneous) ” on page 763.
- See “ NetWorker Job Summary Report (Homogeneous) ” on page 766.
- See “ IBM Spectrum Protect (TSM) Job Summary Report (Homogeneous) ” on page 762.

In the scope selector, use the following options to filter the output:

- Event Type: All Backup & Restore Events, All Backup Events, Full Backups, Incremental Backups, Restores.

- Event Status: All Events, Successful Events, Warning Events, Successful or Warning Events, Failed Events
- Not all of the following fields will appear in all Job Summary reports. Many fields are specific to a backup product. For example, NetBackup jobs will display a Policy field that may not be displayed in the Job Summary for other backup products. Also, a job's result code will be labeled according to the backup product's standard. For example, NetBackup uses Exit Code, while Backup Exec uses Result Code.

Job Summary - Heterogeneous Backup Products

The fields listed in this table reflect the columns that are displayed when the report scope applies to multiple backup products.

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary Advanced Options" on page 769.

Table 22-16 Job Summary - Heterogeneous Backup Products

Column name	Description
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the details. See " Host Details " on page 838. For other backup sources, the object name is displayed with a drill down to the details. See " Object Details " on page 860. For Veeam Backup & Replication, if the job is a Parent, a Source is not displayed.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on. For Veeam Backup & Replication, if the job is a Parent, a Source Type is not displayed.
Parent/Child	Displays Parent or Child for those jobs that have a hierarchical relationship. The field is blank if the job is not categorized as a parent or child job.
Server	Name of the host to display in Portal reports.
Product	Vendor-specific backup software.

Table 22-16 Job Summary - Heterogeneous Backup Products *(continued)*

Column name	Description
Type	The type of backup job. For a list of the types that are relevant for each backup product, click the **Advanced** button in the Job Summary Report Designer.
Start Date	Date and time of the beginning of the backup job
Finish Date	Click on this link to launch the Job Details report. For HP Data Protector jobs, multiple items will be reported, representing the HP Data Protector session, which is a collection of clients and their mount points. For these jobs, the finish date is derived from one of the jobs in a session.
Duration	Backup job duration
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
Exit Code	The exit code of the backup job. For failed jobs, place your mouse on top of the red exit code to view the exit code details in a pop-up tool tip.
Status	Indicates the job status: Successful, Partial, or Failed.
# of Files	The number of files written to the backup media.
Tapes	Click on a tape link to launch a Tape Details Report.

Job Throughput by Client

This report shows the average throughput calculated per minute and aggregated for an hour for a given host. The throughput is averaged based on how the jobs are grouped. The report is designed to evaluate a single host; however, if you choose to run it for multiple hosts, it displays a summarized chart of clients from all the hosts. Currently, the report displays the chart of 20 clients causing the highest throughput, so that administrators can focus on specific clients to troubleshoot those responsible for the high throughput.

This report is a component of the Backup Server Performance Dashboard which helps administrators evaluate the reasons for slowness or bad performance in the backup servers or appliances.

To view the report, navigate to *Backup Manager* > *Management Reports* > *Job Throughput by Client*.

This generated report is based on clients of Backup Server(s). Select one or more clients in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 22-17 Report Scope parameters

Parameters	Description	Remark
Host Group	The report can be generated for one or more Host Groups.	For more details on the Report Scope dialog box, see **NetBackup IT Analytics User Guide** > **Generate and maintain reports** > **Report Scope** section.
All Hosts	This report can be generated for all the Backup Hosts or any of its sub-set.	
For Host = All Backup Servers	One or more Backup Servers can be selected from the scope to run the report for all clients corresponding to the Backup Server.	For example: Click **Modify**. Expand **Hosts** >**All Backup Servers** > **Primary Servers**(if any) > **Policies** or **Clients**. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	This report can be filtered on specific Host Attributes associated with the Backup Host(s).	

Note: If you select multiple conditions in **Report Scope** dialog box, NetBackup IT Analytics displays collective data with the selected criteria from the database using AND condition.

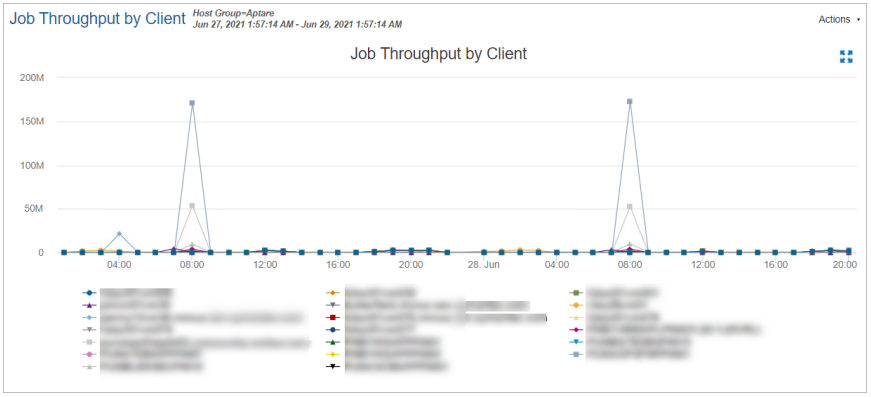
The following are the scenarios in which the report will be displayed based on the criteria selected in **Report Scope** dialog box.

- Scenario 1 - **All Hosts**: If any one Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Client-Backup-Server-Host-Group**: You have a host group named as **All-Client-Backup-Server-Host-Group**which is the parent node. Further, this **All-Client-Backup-Server-Host-Group**node has Two Backup Servers as*Server1*and*Server2*. If you have

selected**All-Client-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be displayed for the*clients of Server1* and *Server2*.

- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

See "Backup Server Performance Dashboard" on page 728.



Avamar Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-18 Avamar Job Summary Report (Homogeneous)

Column name	Description
Client Name	Client that is backed up
Group Name	Links to the Avamar Group Detail
Activity Type	The type of backup, such as On-Demand or Scheduled
Schedule Name	The user assigned name of this job.
Start Date	Actual start date and time of the job
Finish Date	Actual finish date and time of the job
Duration	Amount of time it took for this job to complete
MBytes Sent	Amount of data transferred
MBytes/Sec	Transfer rate
% New Bytes	Percentage of backed up data that was determined to be new
Status Code	Vendor-assigned status code that indicates if the backup was successful. Codes are colored as white, yellow, or red to highlight severity.
Error Code	Vendor-assigned error code for issues that occurred during this backup
Backup Type	The type indicates the recurrence interval, for example, Weekly
Retention	The number of days in the retention period indicates how long the backup will be kept. Backups older than the retention are removed from the system.
# of Files	The number of files written to the backup media
Plugin	OS plug in

Backup Exec Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-19 Backup Exec Job Summary Report (Homogeneous)

Column name	Description
Client	Click on a client name for a job to view the client server details.
Server	The management server name
Job Name	For Backup Exec jobs, this name links to the details. See " Job Details Report " on page 770.
Policy	The name of the backup policy links to the Backup Policy Details.
Schedule Name	Name of this backup schedule.
Type	The type of backup job. For a list of the types that are relevant for each backup product, click the **Advanced** button in the Job Summary Report Designer.
Device Name	For Backup Exec only, the device name.
Resource	For Backup Exec only, the resource name.
Start Date	Date and time of the beginning of the backup job
Finish Date	Click on this link to launch the Job Details report.
Duration	Backup job duration
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
Final Error	The final error code of the Backup Exec job. For failed jobs, place your mouse on top of the red final error code to view the exit code details in a pop-up tool tip.
Status	Indicates the job status: Successful, In Process, Partial, Failed
# of Files	The number of files written to the backup media

Cohesity DataProtect Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See “Job Summary Advanced Options” on page 769.

Table 22-20 Cohesity Data Protect Job Summary Report (Homogeneous)

Column name	Description
Source	Name of the backup source. If the backup source is a known host, host name is displayed with a drill down to details. See “ Host Details ” on page 838. For other backup sources, the object is displayed with a drill down to the details. See “ Object Details ” on page 860. If the job is a **Parent**, a **Source** is not displayed.
Source Type	Type of backup source, such as a datastore, folder, database, virtual machine and so on. If the job is a **Parent**, a **Source Type** is not displayed.
Parent/Child	Displays a Parent or Child for those jobs that have a hierarchical relationship. This field is blank if the job is not categorized as a parent or child job.
Server	Name of the host to display in Portal reports. Click the link to drill down to details. See “ Host Details ” on page 838.
Type	Name of the job type for each backup vendor subsystem.
Job Name	Name of this protection policy.
Session Name	Original identifier of this backup.
Start Date	Date and time this job started.
Finish Date	Date and time this job completed. Click the link to drill down to details. See “ Job Details Report ” on page 770.
Duration	Backup job duration.
Read Size	Indicates the total amount of data read from the source to date.
Transfer Size	Indicates the total amount of data transferred from the source to date.
Data Size	Size of backup in MB.
Backup Size	Indicates the size of the data that was backed up in this session in KiB to date.

Table 22-20 Cohesity Data Protect Job Summary Report (Homogeneous) *(continued)*

Column name	Description
MBytes/Sec	Backup throughput.
Total Source Size	Indicates the total physical size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has a physical disk of 100GB, this field will be equal to 100GB.
Used Source Size	Indicates the physical used size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has 60GB of used physical disk space, this field will be equal to 60GB.
Deduplication	Indicates the deduplication ratio for this session to this point in time.
Compression	Indicates the compression ratio for this session to this point in time.
Status	Indicates job status: Partial, Successful or Failed.

Commvault Simpana Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-21 Commvault Simpana Job Summary Report (Homogeneous)

Column name	Description
Job ID	A backup job's unique internal ID. Click this link to view the Job Details report.
CommServe	CommServe database system host name.
Simpana Job ID	Unique Commvault Simpana ID for the backup job. Click this link to view the Job Details report.
Client	Click a client name for a job to view the details. See " Host Details " on page 838.
iDataAgent	Client software for a specific file system, data base, or hypervisor. Resides on the Client. An individual host can host multiple iDataAgents.

Table 22-21 Commvault Simpana Job Summary Report (Homogeneous) *(continued)*

Column name	Description
Policy	The Commvault Simpana storage policy and links to the details.
Backup Copy	Commvault Simpana storage copy which defines where the backup is actually stored. Links to the copy details which includes policy details.
Type	Backup type that was conducted: Differential, Full, Incremental or Synthetic.
Status	Reflects the status of the job. This summary status is derived from our inspection of the data. The status may be: partial, successful, or failed.
Simpana Status	Current status of the job as reported in the Job Controller of the Commvault Simpana Commcell Console. Job status can be success, killed, failed, or failed to start.
Reason	If applicable, this field indicates a vendor-specific code that you can use to track down troubleshooting tips in the backup vendor's reference guides.
Start Date/Time	Date and time of the beginning of the backup job.
End Date/Time	Date and time of the end of the backup job. Click this link to view details. See "Commvault Simpana Job Details" on page 776.
Duration	Backup job duration.
# of Streams	Number of data channels that connect the client file system or database to the storage media.
Total Job Size (MB)	Total job size in megabytes.
MBytes/Sec	Megabytes per second.
# of Objects	Number of directories determined to be available in the backup source at the start of a backup job, or the number of directories selected to be restored.
# of VMs	The number of VMs used for the backup/restore job. Click this link to view the summary. See "Commvault Simpana ESX Backup Summary" on page 760.

Commvault Simpana ESX Backup Summary

Access this report from the Commvault Simpana Job Summary (Homogeneous) by clicking the # of VMs link in that report.

See " Commvault Simpana Job Summary Report (Homogeneous) " on page 758.

This report enables visibility into the VMs that are being backed up by Commvault Simpana. Use this report for more granular backup auditing, down to the VM Guest level.

Note: Under certain circumstances, some columns will be populated while most other values in the same row are null. This could occur, for example, when a VM job couldn't be started or a Commvault Simpana client had been deleted since the job ran. Even though, in these cases, this report does not show what has been backed up, it does provide insight into what was being backed up or possibly failed, which is useful information for service-level agreements (SLAs).

Table 22-22 Commvault Simpana ESX Backup Summary

Column name	Description
VM Name	The internal name of the virtual machine.
Host	The hostname of the ESX server hosting the virtual machine.
Status	Commvault internal status of the virtual machine backup.
Description	The text reason for why this file failed to be backed up or restored.
VM CBT Status	Changed block tracking status of the virtual machine.
VM Size	The size of the virtual machine. Values are stored as KiB in the database and rendered according to your user profile preferences.
VM Used Space	The used space of the virtual machine. Values are stored as KiB in the database and rendered according to your user profile preferences.
VM Guest Size	The guest size of the virtual machine. Values are stored as KiB in the database and rendered according to your user profile preferences.
VM Guest Valid	Indicates if the reported guest size of the virtual machine is valid.
Backup Start Time	The date and time that the backup of the virtual machine started. May be NULL if there were errors during the backup.
Backup Finish Time	The date and time that the backup of the virtual machine finished. May be NULL if there were errors during the backup.

Table 22-22 Commvault Simpana ESX Backup Summary *(continued)*

Column name	Description
Proxy	The Commvault agent used to backup the virtual machine.
VMware Tools	The VMware tools status for the virtual machine, such as Not Installed, Out of Date or Current.

HP Data Protector Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Note: For HP Data Protector jobs, multiple items will be reported, representing the HP Data Protector session, which is a collection of clients and their mount points. For these jobs, the finish date is derived from one of the jobs in a session.

Table 22-23 HP Data Protector Job Summary Report (Homogeneous)

Column name	Description
Client	Click on a client name for a job to view the client server details.
IP Address	The IP Address of the client
Server	The management server name
Mount Point	For HP Data Protector jobs, the mount point is listed.
Start Date	Date and time of the beginning of the backup job
Finish Date	Click this link to launch the details. See " Job Details Report " on page 770. For HP Data Protector jobs, multiple items will be reported, representing the HP Data Protector session, which is a collection of clients and their mount points. For these jobs, the finish date is derived from one of the jobs in a session.
Duration	Backup job duration

Table 22-23 HP Data Protector Job Summary Report (Homogeneous) *(continued)*

Column name	Description
Status	Indicates the job status: Successful, In Process, Partial, Failed
# of Files	The number of files written to the backup media
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
VM Name	The name of the Virtual Machine.

IBM Spectrum Protect (TSM) Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-24 IBM Spectrum Protect (TSM) Job Summary Report (Homogeneous)

Column name	Description
Client	Click on a client name for a job to view the client server details.
Node Name	For IBM Spectrum Protect (TSM) jobs, the Node Name links to the details. See " Host Details " on page 838.
Server	The management server name
Type	The type of backup job. For a list of the types that are relevant for each backup product, click the **Advanced** button in the scope selector.
Policy	The name of the backup policy links to the Backup Policy Details.
Schedule	If a backup schedule is available, this link will display the Policy Domain Details (for TSM).
Start Date	Date and time of the beginning of the backup job

Table 22-24 IBM Spectrum Protect (TSM) Job Summary Report (Homogeneous) *(continued)*

Column name	Description
Finish Date	Click on this link to launch the details. See "IBM Spectrum Protect (TSM) Job Details Report" on page 779.
Duration	Backup job duration
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
Exit Code	The exit code of the backup job. For failed jobs, place your mouse on top of the red exit code to view the exit code details in a pop-up tool tip.
Result Code	For TSM, the Result Code indicates degrees of success. Refer to the TSM documentation for details.
Status	Indicates the job status: Successful, In Process, Partial, or Failed
# of Files	The number of files written to the backup media

NetBackup Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-25 NetBackup Job Summary Report (Homogeneous)

Column name	Description
Job ID	The ID of the backup job. Click this link to view the details. See " Job Details Report " on page 770.
Client	Click on a client name for a job to view the client server details.
Backup ID	The Backup ID links to the details. See " Job Details Report " on page 770.
Type	The type of backup job. For a list of the types that are relevant for each backup product, click the **Advanced** button in the scope selector.

Table 22-25 NetBackup Job Summary Report (Homogeneous) *(continued)*

Column name	Description
Policy	The name of the backup policy links to the Backup Policy Details.
Policy Keyword	A free-form field that the user can populate with searchable values. This can be used to relate different policies that would otherwise not be linked.
Schedule/Level Type	Specifies if this is a full or incremental backup schedule.
Schedule Name	The user-assigned name of this schedule.
Start Time	Date and time of the beginning of the NetBackup backup job
Finish Time	Click on this link to launch the NetBackup Job Details report.
Duration	Backup job duration
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
Exit Code	The exit code of the NetBackup job. For failed jobs, place your mouse on top of the red exit code to view the exit code details in a pop-up tool tip.
# of Files	The number of files written to the backup media
Tapes	Click on a tape link to launch a Tape Details Report.
Primary Server	Name of the Primary Server for this job.
Media Server	Name of the Media Server to display in Portal reports.
Job Attempt Count	The number of attempts that this job was performed before succeeding.
Storage Unit Type	The type of storage unit as recognized by NetBackup.
Disk Type	The type of disk volume disk: Pure Disk, Advanced Disk.
Total Post Deduplication Size (MB)	The amount of non-redundant data after a backup is complete and that has been written to storage.
Total Deduplication Savings	The amount of data that has been saved due to the whole deduplication process.
Client/Server Deduplication	

Table 22-25 NetBackup Job Summary Report (Homogeneous) *(continued)*

Column name	Description
Dedupe Location	Indicates where the deduplication took place: client-side or server-side.
Bytes Scanned (MB)	Total amount of data scanned, before deduplication
Stream Rate (MBytes/Sec)	Indicates the speed of the scan. KB scanned divided by scan length.
CR Sent (MB)	This is the actual amount of data that is being protected; the amount that has been transmitted to the storage server.
Dedupe Ratio	Indicates the % of data that already has been stored.
Cache Hits	Indicates that a database query was not required to get deduplication data. This is a % of time that the local cache included a record for the segment.
Accelerator Data Sent (KB)	The total number of accelerator kilobytes sent for each NetBackup job. Number of bytes bpdbjobs reports as being sent by Job Accelerator after being converted to kilobytes.
Accelerator Optimization %	The accelerator optimization percentage. Number of bytes sent divided by total number of bytes, expressed as a percentage.
HDR/TIR Deduplication	
Bytes Scanned (MB)	Total amount of data scanned, before deduplication (client-side deduplication only).
Stream Rate (MBytes/Sec)	Indicates the speed of the scan. KB scanned divided by scan length (client-side deduplication only).
CR Sent (MB)	This is the actual amount of data that is being protected; the amount that has been transmitted to the storage server (client-side deduplication only).
Dedup Ratio	Indicates the % of data that already has been stored (client-side deduplication only).
Cache Hits	Indicates that a database query was not required to get deduplication data. This is a % of time that the local cache included a record for the segment.(client-side deduplication only).

NetWorker Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-26 NetWorker Job Summary Report (Homogeneous)

Column name	Description
Client	Click on a client name for a job to view the client server details.
Save Set	This NetWorker-specific field indicates the backup save set associated with the client's backup job.
Server	The management server name
Group	NetWorker-specific group associated with the backup.
Levels	This NetWorker-specific field indicates the backup level/type: Full, Incremental, Migration, Manual, Consolidate, Copy, Diff, 1-9, Restore
Start Date	Date and time of the beginning of the backup job
Finish Date	Click this link to launch the details. See " Job Details Report " on page 770.
Duration	Backup job duration
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
Status	Indicates the job status: Successful, Partial, or Failed
State	Indicates the progress of the backup job, such as Completed or Incomplete.
Tapes	Click on a tape link to launch a Tape Details Report.
Workflow	Name of the workflow for the backup.

Oracle Recovery Manager (RMAN) Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-27 Oracle Recovery Manager (RMAN) Job Summary Report (Homogeneous)

Column name	Description
Source	Name of the backup source. If the backup source is a known host, host name is displayed with a drill down to details. See " Host Details " on page 838. For other backup sources, the object is displayed with a drill down to the details. See " Object Details " on page 860.
Parent/Child	Displays a Parent or Child for those jobs that have a hierarchical relationship. This field is blank if the job is not categorized as a parent or child job.
Server	Name of the host to display in Portal reports. Click the link to display details. See " Host Details " on page 838.
Backup ID	Original identifier of this backup.
Type	Name of the job type for each backup vendor subsystem.
Start Date	The date and time this job started.
Finish Date	The date and time this job completed. Click the link to display details. See " Job Details Report " on page 770.
Duration	Backup job duration.
MBytes	Size of backup in MB.
MBytes/Sec	Backup throughput.
Status	Indicates the job status: Successful, Partial, or Failed.

Table 22-27 Oracle Recovery Manager (RMAN) Job Summary Report (Homogeneous) *(continued)*

Column name	Description
# Backup Items	Number of backup items per job.

Veeam Backup & Replication Job Summary Report (Homogeneous)

For a general description of the Job Summary report refer to the following:

See " Job Summary Report " on page 750.

See "Job Summary - Heterogeneous Backup Products" on page 751.

See "Job Summary Advanced Options" on page 769.

Table 22-28 Veeam Backup & Replication Job Summary Report (Homogeneous)

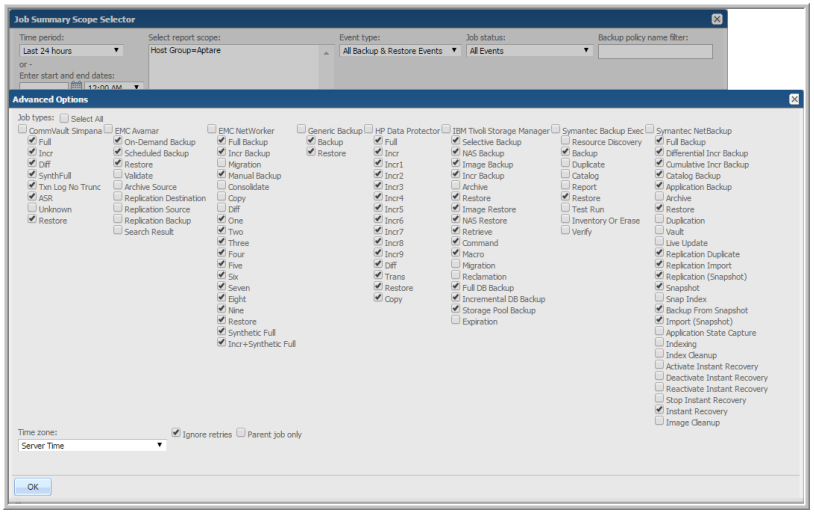
Column name	Description
Source	Name of the backup source. If the backup source is a known host, host name is displayed with a drill down to details. See " Host Details " on page 838. For other backup sources, the object is displayed with a drill down to the details. See " Object Details " on page 860. If the job is a **Parent**, a **Source** is not displayed.
Source Type	Type of backup source, such as a datastore, folder, database, virtual machine and so on. If the job is a **Parent**, a **Source Type** is not displayed.
Parent/Child	Displays a Parent or Child for those jobs that have a hierarchical relationship. This field is blank if the job is not categorized as a parent or child job.
Server	Name of the host to display in Portal reports. Click the link to drill down to details. See " Host Details " on page 838.
Type	The type of backup job. For a list of types that are relevant for each backup product, click **Advanced** in the scope selector.

Table 22-28 Veeam Backup & Replication Job Summary Report (Homogeneous) *(continued)*

Column name	Description
Job Name	Name of the protection policy.
Session Name	Original identifier of this backup.
Start Date	Indicates the date and time this session started.
Finish Date	Indicates the date and time this session completed. Click the link to drill down to details. See " Job Details Report " on page 770.
Duration	Backup job duration.
Read Size	Indicates the total amount of data read from the source to date.
Transfer Size	Indicates the total amount of data transferred from the source to date.
Data Size	Size of the backup in MB.
Backup Size	Indicates the size of the data that was backed up in this session in KiB to date.
MBytes/Sec	Backup throughput.
Total Source Size	Indicates the total physical size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has a physical disk of 100GB, this field will be equal to 100GB.
Used Source Size	Indicates the physical used size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has 60GB of used physical disk space, this field will be equal to 60GB.
Deduplication	Indicates the deduplication ratio for this session to this point in time.
Compression	Indicates the compression ratio for this session to this point in time.
Status	Indicates the job status: Successful, Partial, or Failed.

Job Summary Advanced Options

In the Job Summary scope selector, click **Advanced** to access several advanced options, as shown in the following screen shot.



The report time period is always in the primary server's time. If you select Advanced options and change the time zone, the jobs that are returned are based on the time zone applied to the job start/finish date. The displayed dates are always in the primary server time zone; just the selection of jobs that qualifies in the report time period changes.

Job Details Report

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

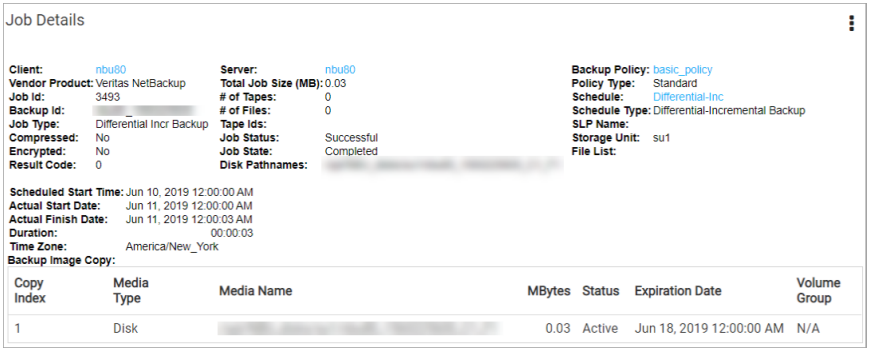
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > Job Summary

Several reports--for example, Job Summary, --contain a **Finish Date** link, which launches a Job Details report.

See " Job Summary Report " on page 750.



- See “Avamar Job Details” on page 773.
- See “Commvault Simpana Job Details” on page 776.
- See “IBM Spectrum Protect (TSM) Job Details Report” on page 779.
- See “ Veeam and RMAN Job Details Report ” on page 781.
- See “Adding a Note to a Job” on page 787.

Table 22-29 Job Details Report

Column name	Description
Client	Click the client name to view the details. See “ Host Details ” on page 838.
Client ID	Displayed only for TSM jobs.
Node Name	Displayed only for TSM jobs.
Vendor Product	Vendor-specific backup software.
Job ID	Backup job unique identifier, For NetBackup jobs, this is the ID assigned by the backup software. For all other backup jobs, this is a unique internal ID.
Backup ID	Unique backup ID.
Job Type	Backup type that was conducted: Differential, Full, Incremental or Synthetic.
Job Status	Reflects the status of the job. This summary status is derived from the data collector’s inspection of the data. The status may be: partial, successful or failed.
Job State	Indicates if the job actually finished running.

Table 22-29 Job Details Report *(continued)*

Column name	Description
Result Code	This vendor-specific code can be used to track down troubleshooting tips in the backup vendor's reference guides.
Scheduled Start Date	The planned date and time of the job's initiation.
Actual Start Date	Date and time of the beginning of the backup job.
Actual Finish Date	Date and time when the backup job ended.
Duration	Backup job duration.
MBytes/Sec	Megabytes per second.
Server Time Zone	The time zone where the job ran.
Server	Click the Server name to launch the details. See " Host Details " on page 838.
Total Job Size (MB)	Job size in MB.
# of Tapes	Number of tapes used by this backup job.
# of Files	Number of files backed up by this job.
Tape IDs	Which tapes were used by the job.
Storage Unit	Physical device name.
File List	Directory path of the file list location.
Disk Pathnames	Disk pathnames
Backup Policy	Click the policy name to launch the Policy Details screen.
Policy Type	Assigned by the System Administrator.
Policy Domain	Displayed only for TSM jobs.
Schedule	Click the schedule name to launch the Policy Details screen.
Schedule Type	The name of the schedule type
SLP Name	Storage Lifecycle Policy, if relevant for NetBackup jobs
Compressed/Encrypted	Indicates how the backup data was treated.

Table 22-29 Job Details Report *(continued)*

Column name	Description
Paths and Related Jobs	The paths being backed up for the job are displayed in a comma-delimited list. When there are multiple related jobs, each of the jobs and associated paths is displayed in the related jobs listing. Clicking on any of the related job IDs will take you to the Job Details report.
Backup Image	Click on the **Media Name** to access additional details in the Tape Details report.
Job Attempts	When there are re-try attempts for any given job, each attempt will be listed in the attempts section at the bottom of the Job Detail. Each attempt row displays details of that attempt, including but not limited to, the backup ID, backup status, started writing date, finished writing date, and transfer rate.
VM Name	For HP Data Protector jobs only, the name of the Virtual Machine.

Avamar Job Details

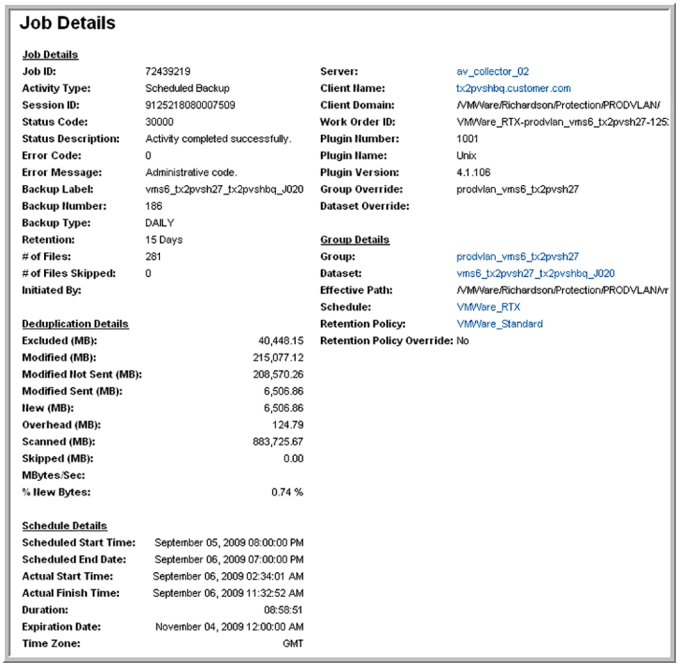
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > Job Summary Report

Click a **Finish Date** link.



Job Details

Job Details

Job ID:	72439219	**Server:**	av_collector_02
Activity Type:	Scheduled Backup	**Client Name:**	tx2pvshbq.customer.com
Session ID:	9125218080007509	**Client Domain:**	/VMWare/Richardson/Protection/PRODVLAN/
Status Code:	30000	**Work Order ID:**	VMWare_RTX-prodvlan_vms6_tx2pvsh27-125
Status Description:	Activity completed successfully.	**Plugin Number:**	1001
Error Code:	0	**Plugin Name:**	Unix
Error Message:	Administrative code.	**Plugin Version:**	4.1.106
Backup Label:	vms6_tx2pvsh27_tx2pvshbq_J020	**Group Override:**	prodvlan_vms6_tx2pvsh27
Backup Number:	186	**Dataset Override:**	
Backup Type:	DAILY		
Retention:	15 Days	**Group Details**	
# of Files:	281	**Group:**	prodvlan_vms6_tx2pvsh27
# of Files Skipped:	0	**Dataset:**	vms6_tx2pvsh27_tx2pvshbq_J020
Initiated By:		**Effective Path:**	/VMWare/Richardson/Protection/PRODVLAN/vr
		Schedule:	VMWare_RTX
Deduplication Details		**Retention Policy:**	VMWare_Standard
Excluded (MB):	40,448.15	**Retention Policy Override:**	No
Modified (MB):	215,077.12		
Modified Not Sent (MB):	208,570.26		
Modified Sent (MB):	6,506.86		
New (MB):	6,506.86		
Overhead (MB):	124.79		
Scanned (MB):	883,725.67		
Skipped (MB):	0.00		
MBytes/Sec:			
% New Bytes:	0.74 %		

Schedule Details

Scheduled Start Time:	September 05, 2009 08:00:00 PM
Scheduled End Date:	September 06, 2009 07:00:00 PM
Actual Start Time:	September 06, 2009 02:34:01 AM
Actual Finish Time:	September 06, 2009 11:32:52 AM
Duration:	08:58:51
Expiration Date:	November 04, 2009 12:00:00 AM
Time Zone:	GMT

Table 22-30 Avamar Job Details

Column name	Description
Job Details	
Job ID	Unique backup job ID
Activity Type	Type of backup activity
Session ID	Backup session ID
Status Code	Indicates if the backup was successful
Status Description	Message associated with the status
Error Code	Vendor-assigned error code for issues that occurred during this backup
Error Message	Message associated with the error code
Backup Label	The backup label associated with this job
Backup Number	Number of the backup
Backup Type	Type of backup job

Table 22-30 Avamar Job Details *(continued)*

Column name	Description
Retention	Indicates how long the backup will be retained
# of Files	Number of files backed up by this job
# of Files Skipped	Number of files skipped by this job
Initiated By	Indicates who initiated this backup job
Server	Avamar server responsible for the backup
Client Name	Client being backed up
Client Domain	Domain of the client being backed up
Work Order ID	Work order ID
Plugin Number	Plug in number
Plugin Name	Descriptive name associated with the plug in number
Plugin Version	Plug in version
Group Override	If set to Yes, Avamar group settings take precedence
Dataset Override	If set to Yes, dataset settings take precedence
Deduplication Details	
Excluded (MB)	Amount of data that was not included to avoid duplication of data
Modified (MB)	Amount of data that was identified as Modified and therefore a candidate for backups
Modified Not Sent (MB)	Amount of the modified data that was determined to be duplicate data and therefore was not sent
Modified Sent (MB)	Amount of modified data that was identified as needing to be backed up
New (MB)	Amount of new data that was never backed up before
Overhead (MB)	Amount of space taken up by overhead data
Scanned (MB)	Amount of data that was scanned to determine if it needed to be backed up
Skipped (MB)	Amount of data that was not backed up because it was determined to be duplicate data

Table 22-30 Avamar Job Details *(continued)*

Column name	Description
MBytes/Sec	Transfer rate
% New Bytes	Percentage of the backed up data that was determined to be new
Schedule Details	
Scheduled Start Time	Date and time that the job was expected to start
Scheduled End Time	Date and time that the job was expected to complete
Actual Start Time	Actual start time of the job
Actual Finish Time	Actual finish time of the job
Duration	Amount of time required for this job to complete
Expiration Date	When is job expires
Time Zone	Time zone where the backup took place
Group Details	
Group	Name of the Avamar group
Dataset	Name of the group dataset
Effective Path	Path to the group dataset
Schedule	Schedule associated with this group
Retention Policy	Retention policy for this group
Retention Policy Override	Indicates if the retention policy can be overridden

Commvault Simpana Job Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > Job Summary Report

On the Job Summary report, click the Job ID or End Date/Time to view the Job Details report.



Table 22-31 Commvault Simpana Job Details

Column name	Description
Client	Click a client name for a job to view the details. See “ Host Details ” on page 838.
Vendor Product	Vendor-specific backup software.
Job ID	A backup job’s unique internal ID.
Job Type	One of the following: Full Backup, Incremental Backup, Restore.
Initiated By	Indicates if the job was initiated by a user through the user interface (GUI) or if it was scheduled.
Is Aged	Indicates if the backup has been aged out.
Simpana Job ID	A backup job’s unique Commvault Simpana internal ID.
Backup Set	Backup set that was protected/recovered during the operation and to which the subclient belongs.
Start Date/Time	Date and time of the beginning of the backup job.

Table 22-31 Commvault Simpana Job Details *(continued)*

Column name	Description
Finish Date/Time	Date and time when the backup job ended.
Duration	Backup job duration.
MBytes/Sec	Megabytes per second.
Time Zone	The time zone where the job ran.
CommServe	Click to launch the Host details.
Total Job Size (MB)	Job size in MB.
Compressed (MB)	Size of the data on the storage media.
Job Status	Reflects the status of the job. This summary status is derived from the data collector's inspection of the data. The status may be: partial, successful, or failed.
Simpana Status	Current status of the job as reported in the Job Controller of the Commvault Simpana Commcell Console. Job status can be success, killed, failed, or failed to start.
iDataAgent	Client software for a specific file system, database, or hyper-visor. Resides on the client. An individual client can host multiple iDataAgents.
Instance	Name of the database instance that will be backed up (for a database iDataAgent).
Policy	Commvault Simpana storage policy.
Backup Copy	Commvault Simpana storage copy which defines where the backup is actually stored. Links to the copy details (which also includes policy details).
# of Streams	Number of data channels that connect the client file system or database to the storage media.
# of Objects	Number of objects involved in the backup or restore. For a file system backup, this is the number of files.
# of Directories	Number of directories determined to be available in the backup source at the start of a backup job, or the number of directories selected to be restored.
# of Skipped Files	Number of files that failed to be backed up.

Table 22-31 Commvault Simpana Job Details *(continued)*

Column name	Description
Subclient	Name of the Commvault Simpana Subclient that defined the backup job.

IBM Spectrum Protect (TSM) Job Details Report

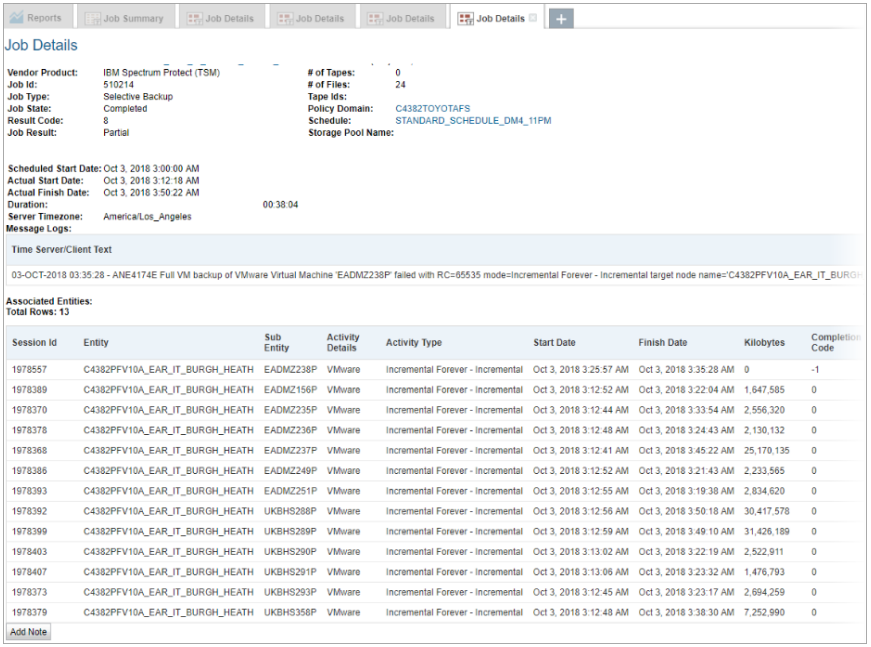
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > Job Summary Report

Click a **Finish Date** link.



Reports | Job Summary | Job Details | Job Details | Job Details | Job Details

Job Details

Vendor Product:	IBM Spectrum Protect (TSM)	**# of Tapes:**	0
Job Id:	510214	**# of Files:**	24
Job Type:	Selective Backup	**Tape Ids:**	
Job State:	Completed	**Policy Domain:**	C4382TOYOTAFS
Result Code:	8	**Schedule:**	STANDARD_SCHEDULE_DM4_11PM
Job Result:	Partial	**Storage Pool Name:**	

Scheduled Start Date: Oct 3, 2018 3:00:00 AM
Actual Start Date: Oct 3, 2018 3:12:18 AM
Actual Finish Date: Oct 3, 2018 3:50:22 AM
Duration: 00:38:04
Server Timezone: America/Los_Angeles
Message Logs:

Time Server/Client Text

03-OCT-2018 03:35:28 - ANE4174E Full VM backup of VMware Virtual Machine 'EADMZ238P' failed with RC=65535 mode=Incremental Forever - Incremental target node name='C4382PFV10A_EAR_IT_BURGH

Associated Entities:
Total Rows: 13

Session Id	Entity	Sub Entity	Activity Details	Activity Type	Start Date	Finish Date	Kilobytes	Completion Code
1978557	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ238P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:25:57 AM	Oct 3, 2018 3:35:28 AM	0	-1
1978389	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ156P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:52 AM	Oct 3, 2018 3:22:04 AM	1,647,585	0
1978370	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ235P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:44 AM	Oct 3, 2018 3:33:54 AM	2,556,320	0
1978378	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ236P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:48 AM	Oct 3, 2018 3:24:43 AM	2,130,132	0
1978368	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ237P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:41 AM	Oct 3, 2018 3:45:22 AM	25,170,135	0
1978386	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ249P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:52 AM	Oct 3, 2018 3:21:43 AM	2,233,565	0
1978393	C4382PFV10A_EAR_IT_BURGH_HEATH	EADMZ251P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:55 AM	Oct 3, 2018 3:19:38 AM	2,834,620	0
1978392	C4382PFV10A_EAR_IT_BURGH_HEATH	UKBHS288P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:56 AM	Oct 3, 2018 3:50:18 AM	30,417,578	0
1978399	C4382PFV10A_EAR_IT_BURGH_HEATH	UKBHS289P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:59 AM	Oct 3, 2018 3:49:10 AM	31,426,189	0
1978403	C4382PFV10A_EAR_IT_BURGH_HEATH	UKBHS290P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:13:02 AM	Oct 3, 2018 3:22:19 AM	2,522,911	0
1978407	C4382PFV10A_EAR_IT_BURGH_HEATH	UKBHS291P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:13:06 AM	Oct 3, 2018 3:23:32 AM	1,476,793	0
1978373	C4382PFV10A_EAR_IT_BURGH_HEATH	UKBHS293P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:45 AM	Oct 3, 2018 3:23:17 AM	2,694,259	0
1978379	C4382PFV10A_EAR_IT_BURGH_HEATH	UKBHS358P	VMware	Incremental Forever - Incremental	Oct 3, 2018 3:12:48 AM	Oct 3, 2018 3:38:30 AM	7,252,990	0

Add Note

Table 22-32 IBM Spectrum Protect (TSM) Job Details Report

Column name	Description
Session ID	Session number; this session number will correspond to the activity log messages generated for this virtual machine.
Entity	The name of the proxy target on whose behalf the activity was performed. Proxy node agent for a BACKUP or RESTORE activity.
Sub Entity	VMware vSphere display name of the virtual machine being backed up or restored.
Activity Details	Context for the type of operation. All Data Protection for VMware records will display a value of VMware.

Table 22-32 IBM Spectrum Protect (TSM) Job Details Report *(continued)*

Column name	Description
Activity Type	Type of activity such as for backup and restore of virtual environments. Valid values include: Incremental Forever - Incremental: The backup was an incremental forever backup, for example, only incremental changes since the last backup of the virtual machine were captured in the backup. Incremental Forever - Full: The backup requested was for a full backup of the virtual machine or the backup requested was an incremental forever backup but the entire (full) virtual machine was captured in the backup. This could be because this was the first backup for the virtual machine, the virtual machine is not capable of tracking incremental changes (for example legacy VMware hardware version), or there was an error encountered in the VMware change block tracking (CBT) mechanism which necessitated a full backup. Periodic Full - Incremental: The backup was a legacy incremental backup, for example, only incremental changes since the last backup of the virtual machine were captured in the backup Periodic Full - Full: The backup was a legacy full backup.
Start Date	Backup start date and time for the virtual machine.
Finish Date	Backup finish date and time for the virtual machine.
Kilobytes	Number of bytes sent transferred to/from the Tivoli Storage Manager Server for a virtual machine backup or restore operation. For backup operations, this number represents the amount of data sent to the Tivoli Storage Manager Server after any data reduction techniques have been applied to the data, e.g., client-side deduplication and/or compression.
Completion Code	This is the return code for the operation as reported in the activity log. For a successful backup or restore, this value will be 0.

Veeam and RMAN Job Details Report

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

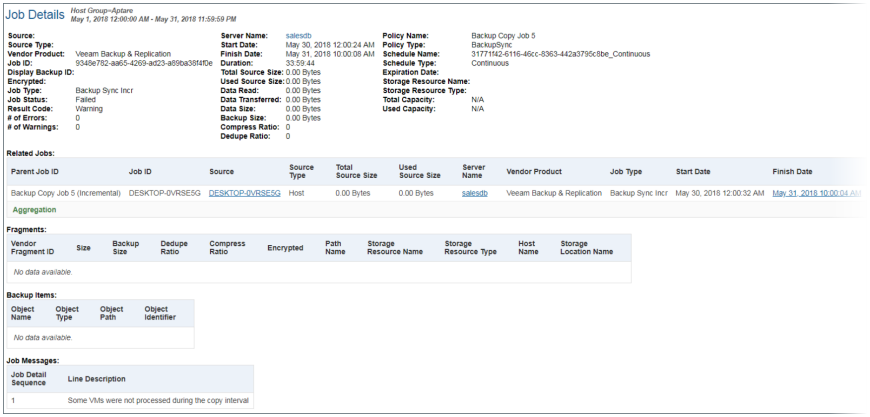
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > Job Summary

Several reports--for example, Job Summary --contain a **Finish Date** link, which launches a Job Details report.

See " Job Summary Report " on page 750.


Job Details Host Group=Aptare
May 1, 2018 12:00:00 AM - May 31, 2018 11:59:59 PM

Source:
Source Type:
Vendor Product: Veeam Backup & Replication
Job ID: 9348e782-aa65-4269-ad23-a69ba38f4f0e
Display Backup ID:
Encrypted:
Job Type: Backup Sync Incr
Job Status: Failed
Result Code: Warning
of Errors: 0
of Warnings: 0

Server Name: salesdb
Start Date: May 30, 2018 12:00:24 AM
Finish Date: May 31, 2018 10:00:08 AM
Duration: 33:59:44
Total Source Size: 0.00 Bytes
Used Source Size: 0.00 Bytes
Data Read: 0.00 Bytes
Data Transferred: 0.00 Bytes
Data Size: 0.00 Bytes
Backup Size: 0.00 Bytes
Compress Ratio: 0
Dedupe Ratio: 0

Policy Name: Backup Copy Job 5
Policy Type: BackupSync
Schedule Name: 31771f42-6116-46cc-8363-442a3795c8be_Continuous
Schedule Type: Continuous
Expiration Date:
Storage Resource Name:
Storage Resource Type:
Total Capacity: N/A
Used Capacity: N/A

Related Jobs:

Parent Job ID	Job ID	Source	Source Type	Total Source Size	Used Source Size	Server Name	Vendor Product	Job Type	Start Date	Finish Date
Backup Copy Job 5 (Incremental)	DESKTOP-0VRSE5G	DESKTOP-0VRSE5G	Host	0.00 Bytes	0.00 Bytes	salesdb	Veeam Backup & Replication	Backup Sync Incr	May 30, 2018 12:00:32 AM	May 31, 2018 10:00:04 AM
Aggregation										

Fragments:

Vendor Fragment ID	Size	Backup Size	Dedupe Ratio	Compress Ratio	Encrypted	Path Name	Storage Resource Name	Storage Resource Type	Host Name	Storage Location Name
No data available.										

Backup Items:

Object Name	Object Type	Object Path	Object Identifier
No data available.			

Job Messages:

Job Detail Sequence	Line Description
1	Some VMs were not processed during the copy interval

Table 22-33 Veeam and RMAN Job Details Report

Column name	Description
Source	Click the Source name to view the details. See " Host Details " on page 838. or See " Object Details " on page 860. reports. If an Object is being backed up, client name will be same as server nam
Source Type	Defines the source type such as resource pool, datastore or a folder.
Vendor Product	Vendor-specific backup software.
Job ID	Backup job unique identifier. This is a unique internal ID. It is a vendor specific job identifier.

Table 22-33 Veeam and RMAN Job Details Report *(continued)*

Column name	Description
Display Backup ID	Original identifier of this backup.
Encrypted	Indicates if this backup job is encrypted (Yes) or (No).
Job Type	Backup type that was conducted: Differential, Full, Incremental or Synthetic.
Job Status	Reflects the status of the job. This summary status is derived from the data collector's inspection of the data. The status may be: partial, successful or failed.
Result Code	This vendor-specific code can be used to track down troubleshooting tips in the backup vendor's reference guides.
# of Errors	The number of errors logged for the backup job within the designated time frame.
# of Warnings	The number of warnings logged for the backup job within the designated time frame.
Server Name	Click the Server name to launch the details. See " Host Details " on page 838.
Start Date	Date and time of the beginning of the backup job.
Finish Date	Date and time when the backup job ended.
Duration	Backup job duration.
Total Source Size	Indicates the total physical size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has a physical disk of 100GB, this field will be equal to 100GB.
Used Source Size	Indicates the physical used size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has 60GB of used physical disk space, this field will be equal to 60GB.
Data Read	Indicates the total amount of data read from the source.
Data Transferred	Indicates the total amount of data transferred from the source.
Data Size	The size of the backup job in kilobytes.
Backup Size	Indicates the size of the data that got backed up in this session in KiB (so far).
Compress Ratio	Indicates the compression ratio for this backup job.

Table 22-33 Veeam and RMAN Job Details Report *(continued)*

Column name	Description
Dedupe Ratio	Indicates the deduplication ratio for this backup job.
Policy Name	The name of protection policy associated to the protected object. A protection policy contains the list of objects to be protected. The protection policy is associated with a policy template that describes the backup retention and archive schedules.
Policy Type	Defines the type of protection policy, for example Backup retention or Archive schedules.
Schedule Name	Schedule name configured for this job. The schedule name uniquely identifies the schedules available in the subsystem.
Schedule Type	Type of this schedule. For example, Daily, Weekly, Monthly, Periodically, Continuous', and ScheduleAfterJob.
Expiration Date	The date when the backup expires.
Storage Resource Name	Name of the storage resource where the backups of this backup job will be saved.
Storage Resource Type	Type of this storage resource. For example, Backup Repository or Scale Out Backup Repository.
Total Capacity	Total capacity of this storage resource in KiB.
Used Capacity	Used capacity of this storage resource in KiB.

Related Jobs

This section of the report displays other jobs within the hierarchy that are related to specified Object Name. These jobs may be parents, children or siblings.

Table 22-34 Veeam and RMAN Job Details Report

Column name	Description
Parent Job ID	Original identifier of this backup job for the Parent.
Job ID	Backup job unique identifier. This is a unique internal ID. It is a vendor specific job identifier.
Source	Name of the backup source object. Click to view the details. See " Object Details " on page 860.

Table 22-34 Veeam and RMAN Job Details Report *(continued)*

Column name	Description
Source Type	Defines the source type such as resource pool, datastore or a folder.
Total Source Size	Indicates the total physical size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has a physical disk of 100GB, this field will be equal to 100GB.
Used Source Size	Indicates the physical used size of the source which is being backed up in this session in KiB. For example, if a VM that is being backed up, has 60GB of used physical disk space, this field will be equal to 60GB.
Server Name	Click the Server name to launch the details. See " Host Details " on page 838.
Vendor Product	Vendor specific backup software.
Job Type	Backup type that was conducted: Differential, Full, Incremental or Synthetic.
Start Date	Date and time of the beginning of the backup job.
Finish Date	Date and time when the backup job ended. Click to view the details. See " Job Details Report " on page 770.
Duration	Backup job duration.
Data Read	Indicates the total amount of data read from the source.
Data Transferred	Indicates the total amount of data transferred from the source.
Data Size	Indicates the size of the data to be backed up in this job in KiB.
Exit Code	The exit code of the backup job. For failed jobs, place your mouse on top of the red exit code to view the exit code details in a pop-up tool tip.

Fragments

This section of the report lists information related to the backup fragment details.

Table 22-35 Veeam and RMAN Job Details Report

Column name	Description
Vendor Fragment ID	Unique identifier of the backup fragment.

Table 22-35 Veeam and RMAN Job Details Report *(continued)*

Column name	Description
Size	Indicates the size of the data to be backed up in this fragment in KiB.
Backup Size	Size of the storage resource where the backups of this fragment is saved.
Dedupe Ratio	Indicates the deduplication ratio for this backup fragment.
Compress Ratio	Indicates the compression ratio for this backup fragment.
Encrypted	Indicates if this backup fragment is encrypted (Yes) or (No).
Path Name	Path of this backup fragment.
Storage Resource Name	Name of the storage resource where the backups of this backup fragment will be saved.
Storage Resource Type	Type of this storage resource. For example, Backup Repository or Scale Out Backup Repository.
Host Name	Host name that the storage location points to. Click to view the details. See “ Host Details ” on page 838.
Storage Location Name	Indicates the storage location, where this backup fragment is stored.

Backup Items

This section of the report lists details for the backup source objects.

Table 22-36 Veeam and RMAN Job Details Report

Column name	Description
Object Name	Name of the backup source object. Click to view the details. See “ Object Details ” on page 860.
Object Type	Defines the object type such as resource pool, datastore or a folder.
Object Path	Provides the hierarchical path of this backup source object.
Object Identifier	Unique identifier of this backup source object.

Job Messages

Table 22-37 Veeam and RMAN Job Details Report

Column name	Description
Job Detail Sequence	Internal sequence which determine the order in which messages are displayed.
Line Description	The output of this session describing the warning, error or success information.

Adding a Note to a Job

As you view a Job Details report, you can add a note that will remain attached to the job. The note automatically supplies your email address and the date the note was added. You supply a message. This feature offers an essential audit trail capability to ensure Sarbanes-Oxley (SOX) compliance.

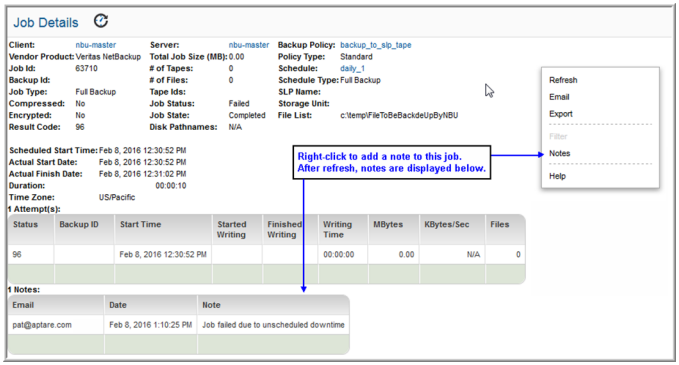
See " Job Details Report " on page 770.

To add a note to a job

1. In the Job Summary report, click a **Finish Date** link to display the details.

   See " Job Details Report " on page 770.

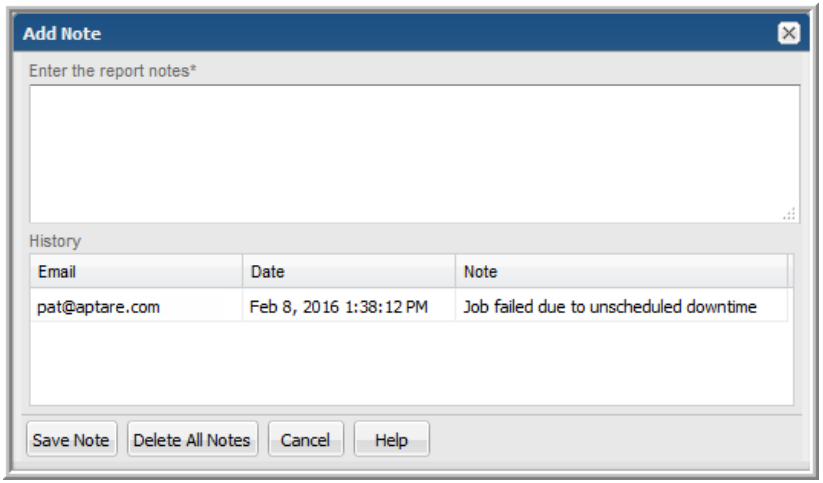
2. In the Job Details report, right-click and then select **Notes**.



1. In the Add Note dialog, enter the note and click **Save Note**.

   In a similar fashion, you can delete all notes.

After a refresh of the Job Details report, the note will appear at the bottom of the report.



Note: Tip: When you add a note to a job or delete all notes, you must refresh the Job Details page so that you can view the results of your action.

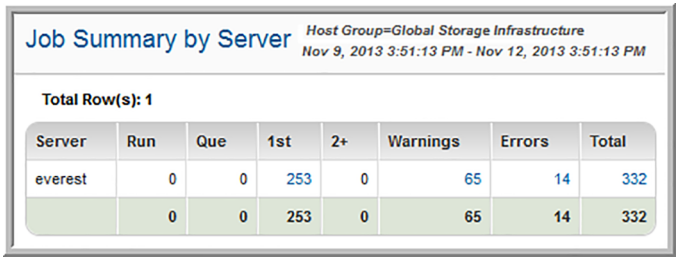
Job Summary by Server

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Summary By Server

This report displays the number of jobs and their status, for jobs started during the time period that you specified for the report.



Job Summary by Server
Host Group=Global Storage Infrastructure
Nov 9, 2013 3:51:13 PM - Nov 12, 2013 3:51:13 PM

Total Row(s): 1

Server	Run	Que	1st	2+	Warnings	Errors	Total
everest	0	0	253	0	65	14	332
	0	0	253	0	65	14	332

Note: Only Cohesity NetBackup, and Commvault Simpana jobs are listed in this report. The report for NetWorker servers is titled, Real Time Save Set Summary.

Table 22-38 Job Summary by Server

Column name	Description
Server	The backup management server.
Run	Jobs that are currently running and have been assigned to a storage unit. this value is the total current running jobs and is not constrained by the selected report time period scope.
Que	Jobs that are queued for an available storage unit. This value is the total currently queued jobs and it is not constrained by the selected report time period scope.
1st	Jobs that completed successfully on their first try.
2+	Jobs that completed successfully after two or more attempts.
Warnings	Jobs that completed, but were partial jobs with one or more warnings.
Errors	Jobs that completed, but a fatal error occurred.
Total	Each numeric column is totaled at the bottom of the table.

Job Summary by Source

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Summary By Source

This report shows the job summary with the rate of success by source. It displays totals for successful, partially successful, and failed jobs for each source. Drilldown on sources, hosts and objects to show details.

Table 22-39 Job Summary by Source

Column name	Description
Source	Backup source name. If backup source is a known host, host name is a link to Host Details. Other backup sources are displayed with link to Object Details.
Source Type	Type of the backup source.
Server	The backup management server. Click drilldown to Host Details.
Product	Subsystem vendor product name.
Total Jobs	Total number of jobs for the backup source in the selected time period. Click to drilldown to the Job Summary report.
Success	Successful jobs for the backup source in the selected time period. Click to drilldown to the Job Summary report.
Failed	Failed jobs for the Source in the selected time period. Click to drilldown to the Job Summary report.
Warning	Jobs that completed, but were partial jobs with one or more warnings in the selected time period. Click to drilldown to the Job Summary report.
Success Rate	Percentage of successful jobs over total jobs.

Job Type Count

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

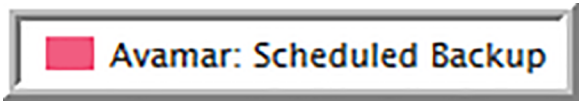
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Type Count

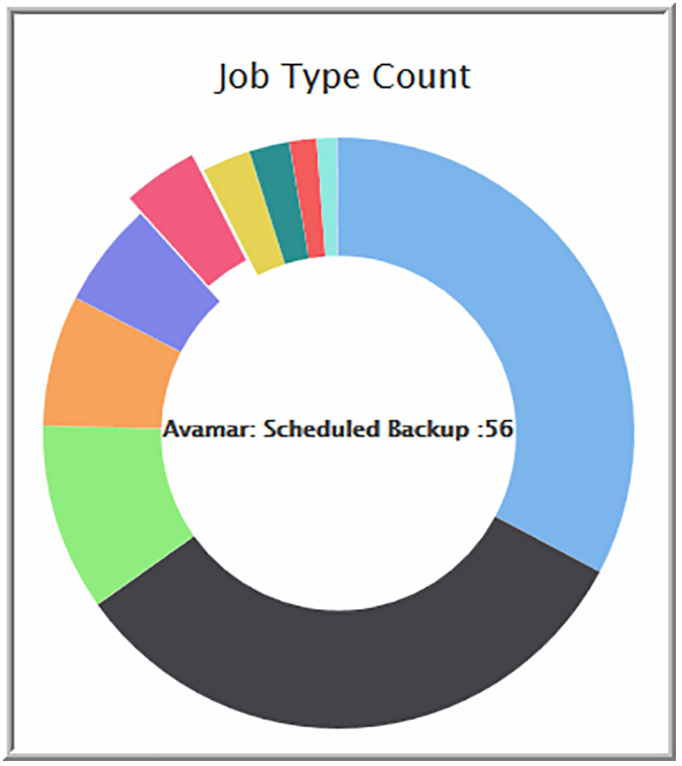
The Job Type Count report displays a donut chart that aggregates job types. Use this chart for a high-level view of backup jobs, with access to the relevant details.

This high-level summary provides drill down access to the Job Summary report, where you'll find specific status messages and scheduling information.

The Legend for the donut chart has the following format: Color Backup Product: Job Type



- Click a legend item to remove that sector from the rendered chart. The legend items can be toggled to display/remove a sector.
- When you click the donut sector, additional details display the job type and the number of occurrences of that type.



Job Volume Summary Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Job Volume Summary

The Job Volume Summary report shows the number of files and amount of data that you backed up for a specified time period. While the Data Domain Backup Summary report approaches the data from an “event status” point of view, the Job Volume Summary report presents data from the perspective of “how much” has been backed up, with a focus on determining if your scheduling is balanced.

See “ Data Domain Backup Summary ” on page 734.

The true value of this high-level report is its ease of use. At a glance--and with a few mouse clicks--you can ascertain problematic time periods, consumption trends, and scheduling peaks.

Use the Scope Selector to filter the report to just the jobs you want to analyze.

Use this procedure to analyze volume usage trends:

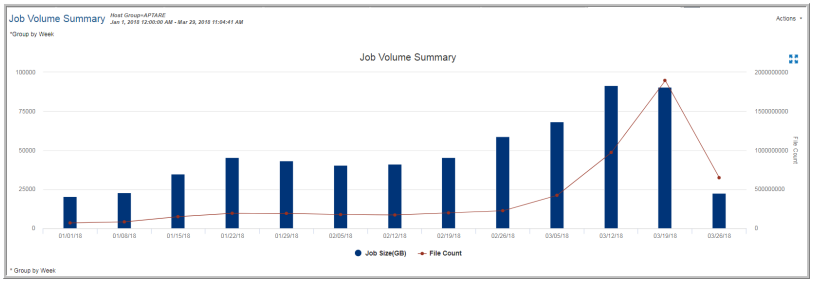
1. View the Job Volume Summary report.
2. Mouse-over a triangle at the top of a bar that indicates a peak usage period.

   A pop-up displays the number of files that were backed up for that period.
3. Click anywhere on the bar to launch the Job Summary.

   In the Job Summary report, pay attention to the data in the following columns: Duration, Written, MBytes/Sec, Exit Code and Job Status. This data gives you an idea of which jobs have over-tasked your storage. Remember that you can sort the report columns to group jobs that took the longest or that failed.

   See “ Job Volume Summary Report ” on page 792.

   See “ Job Summary Report ” on page 750.



The Job Volume Summary report provides the following information:

Table 22-40 Job Volume Summary Report

Column name	Description
Drillable Bar Chart	Click on any of the bars to go to a summary. See " Job Summary Report " on page 750. page for the backups that correspond to the time period of the bar that you selected.
Pop-up Details	Mouse-over any of the bars or symbols on the chart to display pop-ups that provide details. ■ Bar: Displays the volume of data backed up for the associated time period. ■ Triangle symbol: Displays the total number of files backed up.
Backup Volume	Lists the amount of data backed up.
Backup Files	Lists the # of files backed up.

Largest Backup Volume

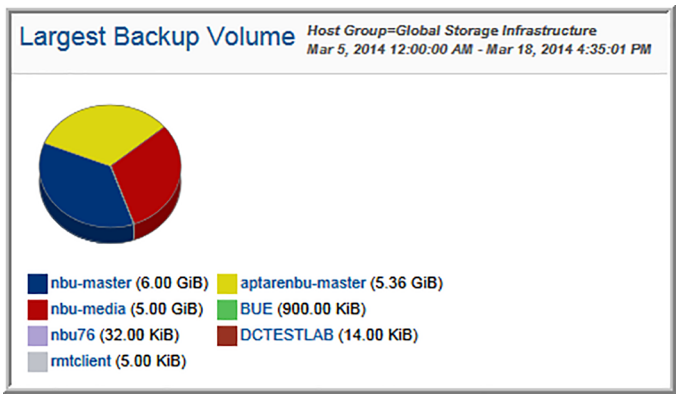
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Largest Backup Volume

Use this report to identify consumers of your most utilized backup volume. A pie chart graphically represents the usage that also is listed in the legend. This report will help you identify greedy servers or applications.

Click the server links in the legend to drill down to details.



The Largest Backup Volume is a real-time snapshot of the largest server backup consumers within the report scope.

- Hold your mouse over any of the pie chart sections to display pertinent server consumption details.
- Click any section of the pie chart to view the Server Consumption Summaryreport for that particular server.

Primary Server Job Throughput

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > **Management Reports** > **Primary Server Job Throughput**

Use this report as a tool to compare performance of various Primary servers. This report displays throughput and job success information for each Primary server. It shows job count and job success rate, and enables you to compare throughputs of different servers to identify any poorly performing servers

Table 22-41 Primary Server Job Throughput

Column name	Description
Primary	host name to display in Portal reports.
Total Job Count	Total number of jobs.
Job Duration	Total duration of the jobs.
Backup Size	The size of the backup.
Throughput	Throughput of the backup server in MBytes per second.
Successful Jobs	Total number of successful jobs.
Failure Jobs	Total number of failed jobs.
Success Rate (%)	Percentage of successful jobs out of total jobs run.

Monthly Backup Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Monthly Backup Summary

Note the following checkbox in the Scope Selector:

- **By Primary Server Only**- Report on Primary Servers rather than clients. For example, if a client is backed up by two Primary Servers, you can avoid having duplicate jobs in the list.

Table 22-42 Monthly Backup Summary

Column name	Description
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.

Table 22-42 Monthly Backup Summary *(continued)*

Column name	Description
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
Total Backups	The total backup jobs that ran for the associated client or server.
MBytes	The amount of data backed up for the client or server.
Last Backup	The Monthly Backup Summary provides you with the date and time of the most recent backup event for each host in their host group. Clicking on the link will take you to the Job Details report where you can view backup details.
Last Full Backup	The Monthly Backup Summary provides you with the date and time of the most recent full backup event for each host in the chosen host group. Clicking on the link will take you to the Job Details report where you can view backup details regarding the full backup, including but not limited to: Job Type, Exit Status, Status Description, # of Tapes, Tape ID #s, and # of Files.

NetBackup AIR-Replication Import Jobs

This report provides information about the NetBackup AIR-Replication Import Jobs. It identifies successful as well as failed AIR-Replication Import Jobs. This report complements the out-of the-box (OOB) AIR-Replication Failed Import Jobs Alert and provides detailed information on which Primary Servers, SLPs, and clients are involved in the failed jobs.

This report is located here:

Backup Manager > **Management Reports** > **NetBackup AIR-Replication Import Jobs**

Table 22-43 NetBackup AIR-Replication Import Jobs

Column name	Description
NBU Job ID	Unique NetBackup ID for the backup job.
Backup Image ID	Unique NetBackup backup image ID.
Source Primary Server	The primary server where the original backup took place.

Table 22-43 NetBackup AIR-Replication Import Jobs *(continued)*

Column name	Description
Destination Primary Server	Primary Server where the NBU Auto Image Replication process imported the Backup that originally took place on the Source Primary Server.
Client	Name of the client backed up by the SLP on the primary server.
Policy Name	Policy associated with the SLP.
Policy Type Name	Indicates the type of data being backed up. For example, **Oracle** indicates data being backed up is from Oracle Database.
Schedule Name	Schedule that is associated with the NetBackup policy.
Lifecycle Policy	Name of the Storage Lifecycle Policy (SLP) configured on the primary server.
Disk Pool Name	Name of the disk pool.
Disk Type	The type of disk volume disk: Pure Disk, Advanced Disk.
Job Type	Backup type that was conducted: Differential, Full, Incremental or Synthetic.
Job Type Name	Indicates the type of Backup. Fo example, **Full Backup** indicates the Job is performing a FULL Backup.
Job Start Date	Actual start date and time of the job.
Job Finish Date	Actual finish date and time of the job.
Job Duration	Amount of time it took for the job to complete.
Job Size	Total job size.
Summary Status	Status of the job - success or failed.
Summary Status Description	Explains the summary status.
Vendor Status	NetBackup Job status code.
Vendor Status Description	NetBackup Job Status code description.

Table 22-43 NetBackup AIR-Replication Import Jobs *(continued)*

Column name	Description
Error Code	Numeric code reported by the backup product. Refer to the backup product's reference guides for troubleshooting advice.
Error Message	Message associated with the error code.

NetBackup deduplication to MSDP savings

These reports provide a view of deduplication trend in your environment and help you to identify the primary servers, clients, and policies across your infrastructure that need improvement in deduplication savings.

You can use the NetBackup deduplication to Media Server Deduplication Pool (MSDP) savings dashboard and reports to evaluate specific benefits described below.

Table 22-44 Report objectives

Report	Objective
NetBackup Deduplication MSDP Savings Dashboard	View deduplication savings by Primary Server, Policy Type, Clients, as well as the deduplication savings trend, all in one place.
NetBackup Deduplication to MSDP Savings - By Primary Servers	Obtain complete visibility into the NetBackup deduplication foot print across your infrastructure
NetBackup Deduplication to MSDP Savings - By Clients	Identify clients with poor deduplication savings.
NetBackup Deduplication to MSDP Savings Trend Over Time	View deduplication performance over time.
NetBackup Deduplication to MSDP Savings - By Policy Type	Evaluate policy types and get information on the types of files being backed up to MSDP.

NetBackup Deduplication to MSDP Savings Dashboard

This dashboard provides a single pane of glass to view deduplication savings by:

- Primary Server
- Policy Type

- Clients

In addition, it also provides the trending data about the deduplication savings, as well as pre- and post-deduplication analytics in the same view.

You can quickly identify areas that need further improvements to obtain better deduplication savings. You can also compare the savings trends before and after making improvements, such as removing frequently changing files, compressed files (that typically account for bad deduplication ratios) from the backups before sending backups to MSDP.

This dashboard comprises of the following reports:

NetBackup Deduplication to MSDP Savings - By Primary Servers

This report provides a view of deduplication savings for each primary server. You can use this report to focus on the primary servers not providing optimal deduplication in your environment.

See "NetBackup Deduplication to MSDP Savings - By Primary Servers" on page 800.

NetBackup Deduplication to MSDP Savings - By Clients

This report provides a view of clients that are receiving poor deduplication savings in your environment. You can sort the **Deduplication Savings %** column in ascending order to filter the clients receiving least deduplication savings at the top. You can investigate such clients to improve their deduplication savings.

See "NetBackup Deduplication to MSDP Savings - By Clients" on page 801.

NetBackup Deduplication to MSDP Savings - By Policy Type

This report provides a view of deduplication savings for each policy type. You can use this report to detect the policy types failing provide optimal deduplication.

See "NetBackup Deduplication to MSDP Savings - By Policy Type" on page 803.

NetBackup Deduplication to MSDP Savings Trend Over Time

You can view the deduplication trend in your environment to evaluate the impact of any improvements you make towards increasing the deduplication ratios.

See "NetBackup Deduplication to MSDP Savings Trend Over Time" on page 803.

NetBackup Pre Vs Post Deduplication Size

You can compare the deduplication sizes of pre- and post backups performed using NetBackup. You can a view the historical trend that allows you to compare the total number of bytes that exist in all the backups performed through NetBackup

You can view the historical trend and compare the total bytes that exist on all backups before and after deduplication for the specified time period.

This report is located here: **Backup Manager** > **Storage Utilization Reports** > **NetBackup Pre Vs Post Deduplication Size**

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

NetBackup Deduplication to MSDP Savings - By Primary Servers

This report provides a view of deduplication savings for each Primary server. You can use this report to focus on the Primary servers not providing optimal deduplication in your environment.

How to use this report

To identify the primary server failing to provide optimal deduplication:

1. Click the **Deduplication Savings %** column header to sort it in ascending order. Along with the least percentage savings, the servers providing least deduplication savings are filtered at the top under **Primary Server** column.
2. Click the respective entries in the **Total Deduplication Jobs** column to drilldown and view the job summary and evaluate the associated policy for poor deduplication savings.

Note: Since compressed files or files with net new changes, such as Oracle REDO log files are prone to poor deduplication savings, you must use a separate policy and assign a non-MSDP pool to save on their backup costs.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > **Management Reports** > **NetBackup Deduplication to MSDP Savings - By Primary Servers**

Table 22-45 NetBackup Deduplication to MSDP Savings - By PrimaryServers

Column name	Description
Primary Server	The NetBackup Primary server that is initiating backups. Click the link to view the host details.
Primary IP	Primary IP address of the source.
Total Deduplication Jobs	Total count of deduplication jobs.
Total Data Scanned Before Dedupe	Amount of data scanned before applying deduplication.
CR Sent (Data Sent to MSDP after Dedupe)	Actual amount of data sent to the MSDP after deduplication.
Deduplication Savings %	The percentage of data saved in percentage due to the deduplication process of the MSDP.

This report is also a component of NetBackup Deduplication MSDP Savings Dashboard that provides a view of the deduplication savings trend and savings by Primary Server, Policy Type, and Clients at the same place.

See "NetBackup Deduplication to MSDP Savings Dashboard" on page 798.

NetBackup Deduplication to MSDP Savings - By Clients

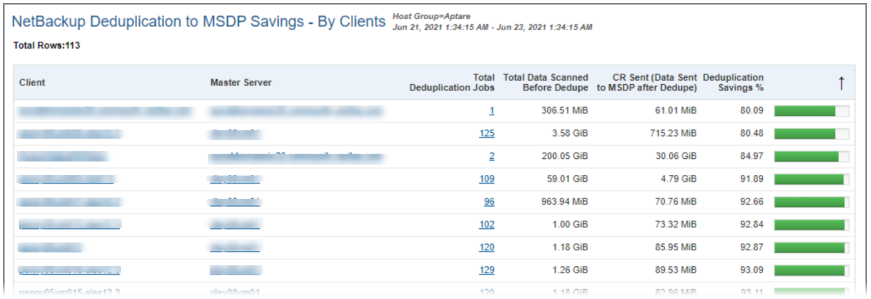
This report provides a view of clients that are receiving poor deduplication savings in your environment. You can sort the **Deduplication Savings %** column in ascending order to filter the clients receiving least deduplication savings at the top. You can investigate such clients to improve their deduplication savings.

How to use the report

To identify the clients receiving poor deduplication:

1. Click the **Deduplication Savings %** column header to sort it in ascending order. Along with the least percentage savings, the clients receiving least deduplication savings are filtered at the in the **Client** column.
2. Click the respective entries in the **Client** and **Primary Server** columns to drilldown and check the host details.
3. Click the respective entries in the **Total Deduplication Jobs** column to drilldown and check the policy name and the types of files being protected in the MSDP.

Note: Since compressed files or files with net new changes, such as Oracle REDO log files are prone to poor deduplication savings, you must use a separate policy and assign a non-MSDP pool to save on their backup costs.


NetBackup Deduplication to MSDP Savings - By Clients

Host Group=Aptare
Jun 21, 2021 1:34:15 AM - Jun 23, 2021 1:34:15 AM

Total Rows:113

Client	Master Server	Total Deduplication Jobs	Total Data Scanned Before Dedupe	CR Sent (Data Sent to MSDP after Dedupe)	Deduplication Savings %
[redacted]	[redacted]	1	306.51 MiB	61.01 MiB	80.09
[redacted]	[redacted]	125	3.58 GiB	715.23 MiB	80.48
[redacted]	[redacted]	2	200.05 GiB	30.06 GiB	84.97
[redacted]	[redacted]	109	59.01 GiB	4.79 GiB	91.89
[redacted]	[redacted]	96	963.94 MiB	70.76 MiB	92.66
[redacted]	[redacted]	102	1.00 GiB	73.32 MiB	92.84
[redacted]	[redacted]	120	1.18 GiB	85.95 MiB	92.87
[redacted]	[redacted]	129	1.26 GiB	89.53 MiB	93.09

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > **Management Reports** > **NetBackup Deduplication to MSDP Savings - By Clients**

Table 22-46 NetBackup Deduplication to MSDP Savings - By Clients

Column name	Description
Client	Names of clients backed up by the NetBackup policy. Click the client name to view its details.
Primary Server	The NetBackup primary server that is initiating backups. Click the link to view the host details.
Total Deduplication Jobs	Total count of deduplication jobs.
Total Data Scanned Before Dedupe	Amount of data scanned before applying deduplication.
CR Sent (Data Sent to MSDP after Dedupe)	Actual amount of data sent to the MSDP after deduplication.
Deduplication Savings %	The percentage of data saved in percentage due to the deduplication process of the MSDP.

This report is also a component of NetBackup Deduplication MSDP Savings Dashboard that provides a view of the deduplication savings trend and savings by Primary Server, Policy Type, and Clients at the same place.

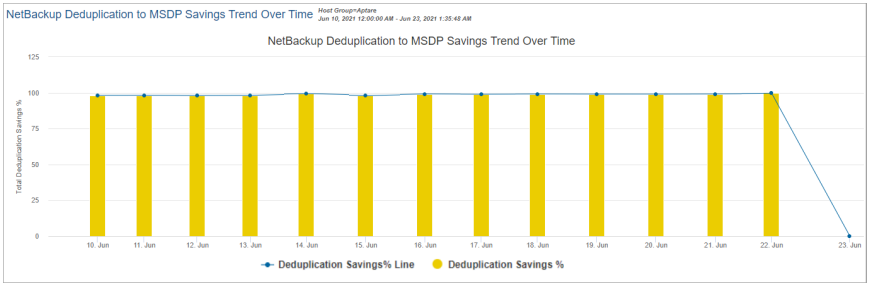
See "NetBackup Deduplication to MSDP Savings Dashboard" on page 798.

NetBackup Deduplication to MSDP Savings Trend Over Time

You can view the deduplication trend in your environment to evaluate the impact of any improvements you make towards increasing the deduplication ratios.

Use the **Reports** tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > **Management Reports** > **NetBackup Deduplication to MSDP Savings - By Clients**



This report is also a component of NetBackup Deduplication MSDP Savings Dashboard that provides a view of the deduplication savings trend and savings by Primary Server, Policy Type, and Clients at the same place.

See "NetBackup Deduplication to MSDP Savings Dashboard" on page 798.

NetBackup Deduplication to MSDP Savings - By Policy Type

This report provides a view of deduplication savings for each policy type. You can use this report to detect the policy types failing provide optimal deduplication.

How to use the report

To identify the policy type failing to provide optimal deduplication:

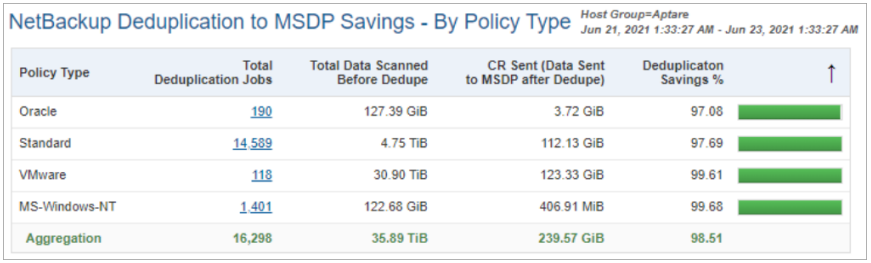
1. Click the **Deduplication Savings %** column header to sort it in ascending order. The policies providing with least deduplication savings are filtered at the top under **Policy Type**.

2. Click the respective entries in the **Total Deduplication Jobs** column to drilldown and check the policy type, policy name, and the types of files being protected in the MSDP.

Note: Since compressed files or files with net new changes, such as Oracle REDO log files are prone to poor deduplication savings, you must use a separate policy and assign a non-MSDP pool to save on their backup costs.

Use the **Reports** tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > **Management Reports** > **NetBackup Deduplication to MSDP Savings - By Policy Type**



NetBackup Deduplication to MSDP Savings - By Policy Type

Host Group=Aptare
Jun 21, 2021 1:33:27 AM - Jun 23, 2021 1:33:27 AM

Policy Type	Total Deduplication Jobs	Total Data Scanned Before Dedupe	CR Sent (Data Sent to MSDP after Dedupe)	Deduplicaton Savings %
Oracle	190	127.39 GiB	3.72 GiB	97.08
Standard	14,589	4.75 TiB	112.13 GiB	97.69
VMware	118	30.90 TiB	123.33 GiB	99.61
MS-Windows-NT	1,401	122.68 GiB	406.91 MiB	99.68
Aggregation	16,298	35.89 TiB	239.57 GiB	98.51

Table 22-47 NetBackup Deduplication to MSDP Savings - By Policy Type

Column name	Description
Policy Type	Defines the type of protection policy assigned by the administrator.
Total Deduplication Jobs	Total count of deduplication jobs.
Total Data Scanned Before Dedupe	Amount of data scanned before applying deduplication.
CR Sent (Data Sent to MSDP after Dedupe)	Actual amount of data sent to the MSDP after deduplication.
Deduplication Savings %	The percentage of data saved in percentage due to the deduplication process of the MSDP.

This report is also a component of NetBackup Deduplication MSDP Savings Dashboard that provides a view of the deduplication savings trend and savings by Primary Server, Policy Type, and Clients at the same place.

See "NetBackup Deduplication to MSDP Savings Dashboard" on page 798.

NetBackup Media Server Job Throughput

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > NetBackup Media Server Job Throughput

This report provides details of data that is backed up by NetBackup media servers and shows server throughput. The report is a useful tool to compare performance of the media servers that are present in your Veritas NetBackup environment.

Table 22-48 NetBackup Media Server Job Throughput

Column name	Description
NetBackup Media Server Job Throughput	Media server name to display in portal reports.
Primary	Primary host name to display in portal reports.
Total Job Count	Total number of jobs.
Job Duration	Total Duration of the job in decimal minutes.
Backup Image Size	The number of bytes occupied by backup image.
Throughput	Throughput of the backup server in MBytes per second.
Successful Jobs	Total number of successful jobs.
Failure Jobs	Total number of failed jobs.
Success Rate (%)	Percentage of successful jobs out of total jobs run.

NetBackup SLP Status

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

To view the report, navigate to *Backup Manager* > *Management Reports* > *NetBackup SLP Status*.

This generated report is based on Backup Server(s). Select one or more Backup Servers in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 22-49 Report Scope parameters

Parameter	Description	Remark
Host Groups	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for one or more the Backup Servers as Hosts from the list.	
For Hosts=All Backup Servers	One or more Backup Servers can be selected from this scope to run the report for all Clients and Polices corresponding to the Backup Server(s).	For example: Click **Modify**. Expand *All Backup Servers* > *Primary Server* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	The report can be filtered on specific Host Attributes associated with backup Host(s).	

Note: If you select multiple conditions in **Report Scope** dialog box, NetBackup IT Analytics displays collective data with the selected criteria from the database using AND condition.

The following are the scenarios in which the report will be displayed based on the criteria selected in **Report Scope** dialog box.

- Scenario 1 - **All Hosts**: If any one or more than one Backup Server as Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Backup-Server-Host-Group**: You have a host group named as **All-Backup-Server-Host-Group** which is the parent node. Further, this **All-Backup-Server-Host-Group** node has Two Backup Servers as*Server1*and*Server2*. If you have selected **All-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be *displayed for Server1* and *Server2*, which are the Backup Servers.
- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

This report lists the Storage Lifecycle Policy (SLP) status for NetBackup jobs, which can be used as an overview of the health of the SLPs for each NetBackup Primary Server.

Using this report's completion statistics, you can determine if the SLP is performing according to schedule, if additional copies have been made, and also if the backlog is increasing. These statistics enable you to identify issues associated with the processing of SLPs.

Table 22-50 NetBackup SLP Status

Column name	Description
Primary Server	This is where the SLP is located. Click the link to drill down to the SLP Status by SLP report. This column lists all the primary servers that have associated SLPs. An SLP may be on the origin primary server or on any of the destination primary servers (in the case of duplication to a remote primary. See " NetBackup SLP Status by SLP " on page 823.
Oldest Unfinished Image	The oldest image that is not SLP-complete for any SLP on the associated primary server. This represents the image end date and time, which can be used to identify possible SLP bottlenecks. The SLP state <> 3 (complete). See " NetBackup SLP Status by Image Copy " on page 817.
# Clients	Total unique clients backed up by the SLP. If a client is backed up more than once by an SLP, the client is only counted once. See " NetBackup SLP Status by Client " on page 813.
# Existing Images	The total number of images for SLPs on the associated primary server. This includes both SLP complete and incomplete images. See " NetBackup SLP Status by Image " on page 820.
Images % SLP Complete	Percentage of images that are SLP complete on the primary server.
Images SLP Complete	Number of images that are SLP complete on the primary server. See " NetBackup SLP Status by Image " on page 820.
Images Not SLP Complete	Number of images that are not SLP complete on the primary server. See " NetBackup SLP Status by Image " on page 820.
# Expected Copies	Total number of image copies for SLPs on the primary includes both complete and incomplete copies. See " NetBackup SLP Status by Image Copy " on page 817.
% Copy Complete	Percentage of completed copies for SLPs on the selected primary server.
Copies Completed	Number of completed copies for SLPs on the selected primary server. See " NetBackup SLP Status by Image Copy " on page 817.

Table 22-50 NetBackup SLP Status *(continued)*

Column name	Description
Copies Not Complete	Number of incomplete copies for SLPs on the selected primary server. See " NetBackup SLP Status by Image Copy " on page 817.
Total Expected Size	This value represents the total for all copies for all SLPs. Total Expected Size for all copies = (Number of copies for SLPs) * (backup copy size for all SLPs) **Note:** The total expected size for all copies is a front-end size and is derived from the number of copies multiplied by the size of the backup copy. The effects of deduplication, compression, or optimized deduplication on the total expected size are not considered in the SLP reports.
Completed Copy Size %	Percentage of completed copy size for SLPs on the selected primary server. The Completed Copy Size is a front-end size. The effects of deduplication, compression, or optimized deduplication on completed copy size are not considered in the SLP reports.
Completed Copy Size	Aggregate size of all the completed copies in the database for the SLP and the specific SLP destination. This value is calculated as: (Completed Copy Count for the specific SLP) x (Size of backup copy for the specific SLP).
Remaining Copy Size	Aggregate size of all remaining copies that are to be created for the specific SLP and image on the selected primary server. This value is calculated as: (Remaining Copy Count) x (Size of Backup Copy1).

NetBackup SLP Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Management Reports > NetBackup SLP Status

Note: In the report scope **All Hosts** and **All Data Domain Servers** are NOT valid inputs for report generation.

This report lists the additional details for a specific NetBackup Storage Lifecycle Policy (SLP). SLPs are retention policies associated with backup policies.

To get to this detailed report, go to any NetBackup SLP report that displays a Version hyperlink.

See " NetBackup SLP Status by Image " on page 820.

See " NetBackup SLP Status by SLP " on page 823.

See " NetBackup SLP Status by Image Copy " on page 817.

Table 22-51 NetBackup SLP Details

Column name	Description
Operation	The type of operation, such as Backup or Duplication. Duplication jobs can be added as hierarchically lower destinations to backups or other hierarchically higher duplication jobs. Operations in this table are indented to represent this hierarchy.
Storage	
Storage Unit	The name of the storage unit where the operation's backup images are stored.
Volume Pool	The volume pool associated with the storage unit.
Media Owner	The owner of the storage media.
Retention	
Type	Retention type, such as Fixed, Staged Capacity Managed, Expire after Duplication, Remote Retention, or Maximum Snapshot Limit.
Period	Period the data will be retained, such as 2 weeks.
Duplication	
Alternate Read Server	The user-defined alternate NetBackup Primary Servers allowed to use the backup image.
Preserve Multiplexing	Indicates if multiplexing is preserved: Yes or No. This is relevant for tape media only.

NetBackup SLP Status by Destination (AIR Activity)

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > NetBackup SLP Status

Several steps are required to view this report:

- In the NetBackup SLP Status report click a Primary Server link to view the report by SLP. In the NetBackup SLP Status by SLP, click the**SLP Name** to view the NetBackup SLP Status by Destination (AIR Activity) report.
  See " NetBackup SLP Status " on page 806.
  See " NetBackup SLP Status by SLP " on page 823.
  See " NetBackup SLP Status by SLP " on page 823.
  See " NetBackup SLP Status by Destination (AIR Activity) " on page 811.

This report lists the details related to Auto Image Replication (AIR) activity. Using the statistics available in this report, you can determine the status of the replication process, such as how much replication to a remote site has been completed.

Table 22-52 NetBackup SLP Status by Destination (AIR Activity)

Column name	Description
Primary Server	This is where the SLP is located. Click the link to drill down to the SLP Status by SLP report. This column lists all the primary servers that have associated SLPs. An SLP may be on the origin primary server or on any of the destination primary servers (in the case of duplication to a remote primary. Click the link to view Host Details.
SLP Name	Name of the SLP that is configured on the primary server
SLP Version	Version that is associated with the specific SLP on the associated primary server. Click the link to view details. See " NetBackup SLP Details " on page 809.
Data Classification	Data classification that is associated with the SLP destination.

Table 22-52 NetBackup SLP Status by Destination (AIR Activity) *(continued)*

Column name	Description
Origin primary Server	Where the primary backup copy resides. This is the server that created the image. The SLP can run on the origin primary server or any of the destination primary servers (in the case of duplication to a remote primary). If duplication is to a remote primary server, the origin primary server and primary server where the SLP is currently running may be different.
Operation	The operation that is performed on an image copy for the specific SLP destination. Examples: Duplication, Backup, or Import.
Retention	Retention period that is associated with the specific SLP destination.
Destination	Destination where an image should reside as per the SLP definition.
Average Lag Time	This value, calculated as the copy time minus the backup time, can be used to determine how much time is required to duplicate the primary copy to a particular destination for a specific SLP. It reflects how fast or slow duplications occur. This value is displayed in seconds.
# Expected Copies	Total number of image copies for the specific SLP on the associated primary server (includes both complete and incomplete copies).
% Copy Complete	Percentage of completed copies for the specific SLP on the selected primary server.
Copies Completed	The number of completed copies for the specific SLP on the selected primary server.
Total Expected Size	The size of all copies for the SLP, calculated as: (Number of copies for the specific SLP) * (backup copy size for the specific SLP). Note: The total expected size for all copies is a front-end size and is derived from the number of copies multiplied by the size of the backup copy. The effects of deduplication, compression, and optimized deduplication on the Total Expected Size are not considered in the SLP reports.
Completed Copy Size %	Percentage of completed copy size for the specific SLP on the selected primary server. This value is calculated as: Completed Copy Size / Total Expected Size - all copies * 100. This is derived from the aggregate size of all the completed copies for the specific SLP on the selected primary server.
Completed Copy Size	Aggregate size of all the completed copies in the database for the SLP and the specific SLP destination. This value is calculated as: (Completed Copy Count for the specific SLP) x (Size of backup copy for the specific SLP).

Table 22-52 NetBackup SLP Status by Destination (AIR Activity) *(continued)*

Column name	Description
Remaining Copy Size	Amount that still needs to be copied.

NetBackup SLP Status by Client

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

To view the report, navigate to *Backup Manager* > *Management Reports* > *NetBackup SLP Status by Client*.

This generated report is based on Backup Server(s). Select one or more Backup Servers in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 22-53 Report Scope parameters

Parameter	Description	Remark
Host Groups	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for one or more the Backup Servers as Hosts from the list.	
For Hosts=All Backup Servers	One or more Backup Servers can be selected from this scope to run the report for all Clients and Polices corresponding to the Backup Server(s).	For example: Click **Modify**. Expand *All Backup Servers* > *Primary Server* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	

Table 22-53 Report Scope parameters *(continued)*

Parameter	Description	Remark
Attributes	The report can be filtered on specific Host Attributes associated with backup Host(s).	

Note: If you select multiple conditions in **Report Scope** dialog box, NetBackup IT Analytics displays collective data with the selected criteria from the database using AND condition.

The following are the scenarios in which the report will be displayed based on the criteria selected in **Report Scope** dialog box.

- Scenario 1 - **All Hosts**: If any one or more than one Backup Server as Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Backup-Server-Host-Group**: You have a host group named as **All-Backup-Server-Host-Group** which is the parent node. Further, this **All-Backup-Server-Host-Group** node has Two Backup Servers as*Server1*and*Server2*. If you have selected **All-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be *displayed for Server1* and *Server2*, which are the Backup Servers.
- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

In the NetBackup SLP Status report, click the **# Clients** link to view the NetBackup SLP Status by Client report.

This report lists the Storage Lifecycle Policy (SLP) status for NetBackup jobs, from the client's perspective.

Table 22-54 NetBackup SLP Status by Client

Column name	Description
Primary Server	This is where the SLP is located. Click the link to drill down to the SLP Status by SLP report. This column lists all the primary servers that have associated SLPs. An SLP may be on the origin primary server or on any of the destination primary servers (in the case of duplication to a remote primary. Click the link to view Host Details.
SLP Name	Name of the SLP that is configured on the primary server
SLP Version	Version that is associated with the specific SLP on the associated primary server. Click the link to view details. See " NetBackup SLP Details " on page 809.
Client Name	Client name that is backed up by the specific SLP on the selected primary server. Click the link to view Host Details.
Origin Primary Server	Where the primary backup copy resides. This is the server that created the image. The SLP can run on the origin primary server or any of the destination primary servers (in the case of duplication to a remote primary). If duplication is to a remote primary server, the origin primary server and primary server where the SLP is currently running may be different.
Oldest Unfinished Image	The oldest image that is not SLP-complete for any SLP on the associated primary server. This represents the image end date and time, which can be used to identify possible SLP bottlenecks. The SLP state <> 3 (complete). See " NetBackup SLP Status by Image Copy " on page 817. report.
# Existing Images	The total number of images that are for SLPs on the selected primary server. This includes both SLP-complete images and SLP-incomplete images. See " NetBackup SLP Status by Image " on page 820.

Table 22-54 NetBackup SLP Status by Client *(continued)*

Column name	Description
Images % SLP Complete	Percentage of images that are SLP-complete for SLPs on the selected primary server, calculated as: Images SLP Complete / Total Existing Images * 100.
Images SLP Complete	Number of images that are SLP-complete for SLPs on the selected primary server, based on: Storage Lifecycle State = 3 (Complete). See " NetBackup SLP Status by Image " on page 820.
Images Not SLP Complete	Number of images that are SLP-incomplete for SLPs on the selected primary server, calculated as: Total Existing Images - Images SLP Complete. See " NetBackup SLP Status by Image " on page 820.
# Expected Copies	Total number of image copies for SLPs on the selected primary server (includes both complete copies and incomplete copies). See " NetBackup SLP Status by Image Copy " on page 817.
% Copy Complete	Percentage of completed copies for SLPs on the selected primary server, calculated as: Copies Completed / Total Copies * 100.
Copies Completed	The number of completed copies for SLPs on the selected primary server, based on Copy State = 3 (Complete). See " NetBackup SLP Status by Image Copy " on page 817.
Copies Not Completed	The number of incomplete copies for SLPs on the selected primary server, calculated as: Copies Not Complete = Total Copies - Copies Completed. See " NetBackup SLP Status by Image Copy " on page 817.
Total Expected Size	This is the expected size of all the copies.

Table 22-54 NetBackup SLP Status by Client *(continued)*

Column name	Description
Completed Copy Size %	Percentage of completed copy size for the specific SLP on the selected primary server. This value is calculated as: Completed Copy Size / Total Expected Size - all copies * 100. This is derived from the aggregate size of all the completed copies for the specific SLP on the selected primary server.
Completed Copy Size (MB)	Aggregate size of all the completed copies in the database for the SLP and the specific SLP destination. This value is calculated as: (Completed Copy Count for the specific SLP) x (Size of backup copy for the specific SLP).
Remaining Copy Size (MB)	Aggregate size of all remaining copies that are to be created for the specific SLP and image on the selected primary server. This value is calculated as: (Remaining Copy Count) x (Size of Backup Copy1).

NetBackup SLP Status by Image Copy

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

To view the report, navigate to *Backup Manager* > *Management Reports* > *NetBackup SLP Status by Image Copy*.

This generated report is based on Backup Server(s). Select one or more Backup Servers in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 22-55 Report Scope parameters

Parameter	Description	Remark
Host Groups	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for one or more the Backup Servers as Hosts from the list.	
For Hosts=All Backup Servers	One or more Backup Servers can be selected from this scope to run the report for all Clients and Polices corresponding to the Backup Server(s).	For example: Click **Modify**. Expand *All Backup Servers* > *Primary Server* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	The report can be filtered on specific Host Attributes associated with backup Host(s).	

Note: If you select multiple conditions in **Report Scope** dialog box, NetBackup IT Analytics displays collective data with the selected criteria from the database using AND condition.

The following are the scenarios in which the report will be displayed based on the criteria selected in **Report Scope** dialog box.

- Scenario 1 - **All Hosts**: If any one or more than one Backup Server as Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Backup-Server-Host-Group**: You have a host group named as **All-Backup-Server-Host-Group** which is the parent node. Further, this **All-Backup-Server-Host-Group** node has Two Backup Servers as*Server1*and*Server2*. If you have selected **All-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be *displayed for Server1* and *Server2*, which are the Backup Servers.

- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

To get to this report, go to any NetBackup SLP report that displays the following hyperlinks: Oldest Unfinished Image, Copies Completed, Copies Not Completed, or Total Copies.

See " NetBackup SLP Status by Client " on page 813.

See " NetBackup SLP Status by Image " on page 820.

This report lists the Storage Lifecycle Policy (SLP) status for NetBackup jobs, from the Image Copy perspective.

Table 22-56 NetBackup SLP Status by Image Copy

Column name	Description
Primary Server	This is where the SLP is located. Click the link to drill down to the SLP Status by SLP report. This column lists all the primary servers that have associated SLPs. An SLP may be on the origin primary server or on any of the destination primary servers (in the case of duplication to a remote primary. Click the link to view Host Details.
SLP Name	Name of the SLP that is configured on the primary server
SLP Version	Version that is associated with the specific SLP on the associated primary server. See " NetBackup SLP Details " on page 809.
Data Classification	This represents the value or importance of the data that has been backed up.
Policy Name	Policy that is associated with the SLP.
Schedule Name	Schedule that is associated with the NetBackup policy.

Table 22-56 NetBackup SLP Status by Image Copy *(continued)*

Column name	Description
Client Name	Client name that is backed up by the SLP for the selected primary server.
Image ID	Image ID associated with the image. The image ID is unique for each image.
Copy Job ID	Job ID associated with the specific operation.
Origin Primary Server	Where the primary backup copy resides. This is the server that created the image. The SLP can run on the origin primary server or any of the destination primary servers (in the case of duplication to a remote primary). If duplication is to a remote primary server, the origin primary server and primary server where the SLP is currently running may be different.
Operation	The operation that is performed on an image copy for the specific SLP destination. Examples: Duplication, Backup, or Import.
State	Current status of the copy, such as Complete, Not Started, In Process, Deleted, or Paused.
Destination	Destination where a copy should reside as per the SLP definition.
Expiration Time	Date and time when the copy expires.
Copy/Import Completion Time	Date and time when the copy is completed.
Lag Time	This value, calculated as the SLP completion time minus the backup time, can be used to determine how much time is required to duplicate the primary copy to a particular destination for a specific SLP. It reflects how fast or slow duplications occur. This value is displayed in seconds.
Completed Copy Size	Aggregate size of all the completed copies in the database for the SLP and the specific SLP destination. This value is calculated as: (Completed Copy Count for the specific SLP) x (Size of backup copy for the specific SLP).

NetBackup SLP Status by Image

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

Ti view the report, navigate to *Backup Manager* > *Management Reports* > *NetBackup SLP Status by Image*.

To get to this report, go to any NetBackup SLP report that displays the following hyperlinks: # Existing Images, Images SLP Complete, and Images SLP Not Complete.

See " NetBackup SLP Status by Client " on page 813.

See " NetBackup SLP Status " on page 806.

This report lists the Storage Lifecycle Policy (SLP) status for NetBackup jobs, from the Image perspective.

Table 22-57 NetBackup SLP Status by Image

Column name	Description
Primary Server	This is where the SLP is located. Click the link to drill down to the SLP Status by SLP report. This column lists all the primary servers that have associated SLPs. An SLP may be on the origin primary server or on any of the destination primary servers (in the case of duplication to a remote primary . Click the link to view Host Details.
SLP Name	Name of the SLP that is configured on the primary server.
SLP Version	Version that is associated with the specific SLP on the associated primary server. Click the link to view details. See " NetBackup SLP Details " on page 809.
Policy Name	Policy that is associated with the SLP.
Schedule Name	Schedule that is associated with the NetBackup policy.
Client Name	Client name that is backed up by the SLP for the selected primary server.

Table 22-57 NetBackup SLP Status by Image *(continued)*

Column name	Description
Origin Primary Server	Where the primary backup copy resides. The SLP may run on the origin primary server or any of the destination primary servers (in the case of duplication to a remote primary). If duplication is to a remote primary server, the origin primary server and primary server where the SLP lives (or is currently running) may be different.
Image ID	Image ID associated with the image. The image ID is unique for each image.
Backup Time	Date and time the primary backup started.
SLP Completion Time	Date and time when the image is SLP-complete.
Lag Time	This value calculated as the SLP completion time minus the backup time, can be used to determine how much time is required to duplicate the primary copy to a particular destination for a specific SLP. It reflects how fast or slow duplications occur. This value is displayed in milliseconds.
Total Copies	Total number of copies for the specific SLP and image on the selected primary server (includes both complete copies and incomplete copies). See " NetBackup SLP Status by Image Copy " on page 817.
Copies Completed	The number of completed copies that are present for the specific SLP and image on the selected primary server. See " NetBackup SLP Status by Image Copy " on page 817.
Copies Not Completed	The number of incomplete copies for the specific SLP and image on the selected primary server, calculated as: Copies not Complete = Total Copies - Copies Completed. See " NetBackup SLP Status by Image Copy " on page 817.

Table 22-57 NetBackup SLP Status by Image *(continued)*

Column name	Description
% Copy Complete	Percentage of completed copies for the specific SLP and image on the selected primary server., calculated as: Copies Completed / Total Copies * 100.
Total Expected Size	The size of all copies for the SLP, calculated as: (Number of copies for the specific SLP) * (backup copy size for the specific SLP). Note: The total expected size for all copies is a front-end size and is derived from the number of copies multiplied by the size of the backup copy. The effects of deduplication, compression, and optimized deduplication on the Total Expected Size are not considered in the SLP reports.
Completed Copy Size	Aggregate size of all the completed copies for the specific SLP and image on the selected primary server.
Remaining Copy Size	Aggregate size of all remaining copies that are to be created for the specific SLP and image on the selected primary server. This value is calculated as: (Remaining Copy Count) x (Size of Backup Copy1).
Completed Copy Size %	Percentage of completed copy size for the specific SLP on the selected primary server. This value is calculated as: Completed Copy Size / Total Expected Size - all copies * 100. This is derived from the aggregate size of all the completed copies for the specific SLP on the selected primary server.

NetBackup SLP Status by SLP

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

To view the report, navigate to *Backup Manager* > *Management Reports* > *NetBackup SLP Status*.

This generated report is based on Backup Server(s). Select one or more Backup Servers in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 22-58 Report Scope parameters

Parameter	Description	Remark
Host Groups	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for one or more the Backup Servers as Hosts from the list.	
For Hosts=All Backup Servers	One or more Backup Servers can be selected from this scope to run the report for all Clients and Polices corresponding to the Backup Server(s).	For example: Click **Modify**. Expand *All Backup Servers* > *Primary Server* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	The report can be filtered on specific Host Attributes associated with backup Host(s).	

Note: If you select multiple conditions in **Report Scope** dialog box, NetBackup IT Analytics displays collective data with the selected criteria from the database using AND condition.

The following are the scenarios in which the report will be displayed based on the criteria selected in **Report Scope** dialog box.

- Scenario 1 - **All Hosts**: If any one or more than one Backup Server as Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Backup-Server-Host-Group**: You have a host group named as **All-Backup-Server-Host-Group** which is the parent node. Further, this **All-Backup-Server-Host-Group** node has Two Backup Servers

as*Server1*and*Server2*. If you have selected **All-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be *displayed for Server1* and *Server2*, which are the Backup Servers.

- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

In the NetBackup SLP Status report, click the Primary Server link to view the NetBackup SLP Status by SLP report.

This report lists the Storage Lifecycle Policy (SLP) status for NetBackup jobs, grouped by SLP.

Table 22-59 NetBackup SLP Status by SLP

Column name	Description
Server	This is where the SLP is located. Click the link to drill down to the SLP Status by SLP report. This column lists all the primary servers that have associated SLPs. An SLP may be on the origin primary server or on any of the destination primary servers (in the case of duplication to a remote primary. Click the link to view Host Details.
SLP Name	Name of the SLP that is configured on the primary server. See " NetBackup SLP Status by Destination (AIR Activity) " on page 811.
SLP Version	Version that is associated with the specific SLP on the selected primary server. See " NetBackup SLP Details " on page 809.
Oldest Unfinished Image	The oldest image that is not SLP-complete for any SLP on the associated primary server. This represents the image end date and time, which can be used to identify possible SLP bottlenecks. See " NetBackup SLP Status by Image Copy " on page 817.

Table 22-59 NetBackup SLP Status by SLP *(continued)*

Column name	Description
# Clients	Number of unique clients (link) that are backed up by the specific SLP. Note that if the same client is backed up more than once by the SLP, the client is counted only once. See " NetBackup SLP Status by Client " on page 813.
# Existing Images	The total number of images that are for the specific SLP on the selected primary server. This includes both SLP-complete images and SLP-incomplete images. See " NetBackup SLP Status by Image " on page 820.
Images % SLP Complete	Percentage of images that are SLP-complete for the specific SLP on the selected primary server, calculated as: Images SLP Complete / Total Existing Images * 100.
Images SLP Complete	Number of images that are SLP-complete for the specific SLP on the selected primary server. See " NetBackup SLP Status by Image " on page 820.
Images Not SLP Complete	Number of images that are SLP-incomplete for the specific SLP on the selected primary server. See " NetBackup SLP Status by Image " on page 820.
# Expected Copies	Total number of image copies for the specific SLP on the selected primary server (includes both complete copies and incomplete copies). See " NetBackup SLP Status by Image Copy " on page 817.
% Copy Complete	Percentage of completed copies for the specific SLP and image on the selected primary server., calculated as: Copies Completed / Total Copies * 100.
Copies Completed	The number of completed copies that are present for the specific SLP and image on the selected primary server. See " NetBackup SLP Status by Image Copy " on page 817.
Copies Not Complete	The number of incomplete copies for the specific SLP and image on the selected primary server, calculated as: Copies not Complete = Total Copies - Copies Completed. See " NetBackup SLP Status by Image Copy " on page 817.

Table 22-59 NetBackup SLP Status by SLP *(continued)*

Column name	Description
Total Expected Size	The size of all copies for the SLP, calculated as: (Number of copies for the specific SLP) * (backup copy size for the specific SLP). Note: The total expected size for all copies is a front-end size and is derived from the number of copies multiplied by the size of the backup copy. The effects of deduplication, compression, and optimized deduplication on the Total Expected Size are not considered in the SLP reports.
Completed Copy Size %	Percentage of completed copy size for the specific SLP on the selected primary server. This value is calculated as: Completed Copy Size / Total Expected Size - all copies * 100. This is derived from the aggregate size of all the completed copies for the specific SLP on the selected primary server.
Completed Copy Size	Aggregate size of all the completed copies in the database for the SLP and the specific SLP destination. This value is calculated as: (Completed Copy Count for the specific SLP) x (Size of backup copy for the specific SLP).
Remaining Copy Size	Aggregate size of all remaining copies that are to be created for the specific SLP and image on the selected primary server. This value is calculated as: (Remaining Copy Count) x (Size of Backup Copy1).

Running and Queued Job Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Management Reports > Running and Queued Job Summary

Note: This report is applicable to only NetBackup and Networker.

The Running and Queued Jobs Summary displays any job that was running or queued at the time the management servers were polled.

Jobs that require attention (long-running, stalled, or slow jobs) will be highlighted in red and moved to the top of the running and queued jobs list so that they are quickly noted by administrators and handled appropriately. Each of these flagged

jobs will be accompanied by one or more icons to indicate the type of potential problem.

NetBackup Running and Queued Jobs Summary

Table 22-60 NetBackup Running and Queued Jobs Summary

Column name	Description
Server	Enables drill down access to the Host Details
Client	Drill down to Client details
Policy Name	The name of the backup policy associated with the job
Job ID	The NetBackup job ID
Try (MB)	Amount of data attempting to be backed up
Try Start Date	Date of the attempted backup
Duration	The length of the backup job
KB/Sec	Transfer rate
Last Written	The date and time of the last write.
Slow Running Long Running Stalled	Indicators identify long-running (>= 12 hours), stalled (no write activity for >= 15 minutes), or slow-running jobs (kilobytes per second throughput <= 200).

NetWorker Running and Queued Jobs Summary

Table 22-61 NetWorker Running and Queued Jobs Summary

Column name	Description
Server	Enables drill down access to the Host Details
Client/Group	Drill down to Client details
Save Set ID	The NetWorker Save Set ID
Duration	The length of the backup job

Table 22-61 NetWorker Running and Queued Jobs Summary *(continued)*

Column name	Description
Slow Running Long Running Stalled	Indicators identify long-running (>= 12 hours), stalled (no write activity for >= 15 minutes), or slow-running jobs (kilobytes per second throughput <= 200).

Chapter 23

Backup Administration Reports

This chapter includes the following topics:

- Backup RTO RPO Compliance
- Avamar Grid Capacity Dashboard
- Command Center Dashboard
- Data Protection Dashboard
- Host Details
- Host Group
- IBM Spectrum Protect (TSM) Storage Pools Dashboard
- Job Histogram
- Jobs Dashboard
- Message of the Day Report
- Mission Control - Backup
- NetBackup Event Notification Summary - CRITICAL
- NetBackup License Report
- NetBackup Host Certificates
- NetBackup Non-Host Certificates
- Object Details

- Operations Dashboard
- Tailoring the Report Output by Backup Product
- TSM Process Summary
- NetBackup Audit Report
- EEBs on Primary server, Media server and Clients reports

Backup RTO RPO Compliance

This report assists in computing RTO (Recovery Time Objectives) and RPO (Recovery Point Objectives) for backups using the time when the last full backup was performed then adding the time it takes to apply any incremental. The time objectives can be set at report runtime and indicators show you if you are meeting your SLAs. This report appears as a drilldown for the Compliance RTO RPO link on the Ransomware Risk Assessment Dashboard.

To generate the report, click *Reports* tab.

Navigate to *Backup Manager > Administration Reports* then click *Backup RTO RPO Compliance* report.

NetBackup IT Analytics displays scope selection dialog box for the report. Select the appropriate values in the dialog box.

This generated report is based on clients of Backup Server(s). Select one or more clients in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 23-1 Report Scope parameters

Parameters	Description	Remark
Host Group	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide > Generate and maintain reports > Report Scope* section.
All Hosts	This report can be generated for all the Backup Hosts or any of its sub-set.	

Table 23-1 Report Scope parameters *(continued)*

Parameters	Description	Remark
For Host = All Backup Servers	One or more Backup Servers can be selected from the scope to run the report for all clients corresponding to the Backup Server.	For example: Click **Modify**. Expand *Hosts* > *All Backup Servers* > *Primary Servers* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	This report can be filtered on specific Host Attributes associated with the Backup Host(s).	

Note: If multiple conditions are selected in the **Report Scope**, NetBackup IT Analytics displays collective data with the selected criteria from the database using the **AND** condition.

Based on the parameters selected in **Report Scope**, the following are the scenarios in which the data in the report will be displayed.

- Scenario 1 - **All Hosts**: If any one Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Client-Backup-Server-Host-Group**: You have a host group named as **All-Client-Backup-Server-Host-Group**which is the parent node. Further, this **All-Client-Backup-Server-Host-Group**node has Two Backup Servers as*Server1*and*Server2*. If you have selected**All-Client-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be displayed for the*clients of Server1* and *Server2*.
- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).

- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

The following table depicts the report columns and the field values displayed in the report.

Table 23-2 Backup RTO RPO Compliance report columns

Column name	Description
Recovery Status	Indicates successful recovery with green and recovery failure in red.
Product	The product for which recovery was performed.
Client	Name of the host/client that has been backed up.
Restore Plan	The configured restore plan. Examples: Full Only, Incremental, Full + Incremental.
Last Good Full	Date and time of the last successful full backup of source data.
Size	The amount of data backed up during the last successful full backup.
Full Duration	The time taken to successfully complete the last full backup of source data.
Last Incremental After Full	Date and time of the last incremental backup after the successful full backup.
# Incr After Full	The number of incremental backups that completed after the successful full backup.
Incr Size After Full	The amount of data backed up during the incremental backup.
Incr Duration After Full	The time taken to successfully complete the last incremental backup after the successful full backup of source data.
Recovery Date	The date and time of the expected completion of full data recovery, based on historical times and sizes.
Recovery Point	The date and time up to which successful recovery was performed.
Lost Time	The time lost in recovering the data.

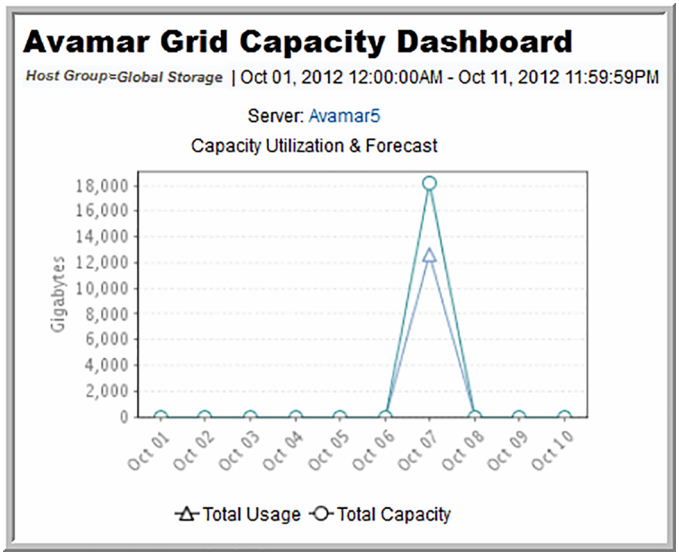
Avamar Grid Capacity Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Administration Reports > Avamar Grid Capacity Dashboard

Use this report to see trends and forecast storage utilization.



The Capacity Utilization & Forecast is derived from the node disk utilization. The graph represents the total usage and capacity for all disks on the server for the point in time.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Command Center Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Administration Reports > Command Center Dashboard

The Command Center Dashboard provides four portlets with a real-time view of your backup environment: Real Time Job Summary, Real Time Library and Drive Status, Real Time Storage Unit Utilization, and Running and Queued Jobs Summary.

Note: Only Veritas NetBackup and EMC NetWorker jobs are represented in this dashboard.

Table 23-3 Command Center Dashboard

Column name	Description
Job Summary by Server	Displays the number of jobs and the job status for the time period that you selected in the report scope: You can drill down on any non-zero value to view the summary for those corresponding jobs. Additionally, you can drill down further to the more comprehensive for those corresponding jobs. Additionally, you can drill down further to the more comprehensive for any individual jobs listed in the summary. See “ Job Summary Report ” on page 750. See “ Job Details Report ” on page 770.
Real Time Library and Drive Status	Summarizes each tape library in your environment. The data will only be as current as the time of the most recent poll of your management server(s) by the Data Collector. Monitor the health of each library by tracking the number of drives up, down, in use and available within each tape library. When the number of down drives in a library is equal to or greater than 25%, the text for the library's row is flagged in red.
	The Real Time Storage Unit Component displays current high-level information for each storage unit in your environment. Each storage unit listed provides you with several columns of pertinent information.

Table 23-3 Command Center Dashboard *(continued)*

Column name	Description
Real Time Storage Unit Utilization	■ Storage Unit: Name of the unit. ■ Media Server: Click this link to display the Server Details report. ■ Type: The type of unit, such disk, disk staging, media manager, or NDMP (Network Data Management Protocol) ■ Total Drives: The configuration of the unit. ■ Jobs Run: Jobs that are currently running and have been assigned to a storage unit. This value is the total current running hobs and is not constrained by the selected report time period scope. ■ Max Drive: Maximum drives per job. ■ Max Mulitplex/Drives: The maximum number of concurrent drives.
	The Drives Up column combined with the Jobs Queued and Jobs Running columns help you determine whether a storage unit is functioning properly. For instance, if there are jobs in the queue for a storage unit, and there are also drives available within a storage unit, there may be something preventing the queued jobs from moving to the running status.
Running and Queued Jobs Summary	Displays any job that was running or queued at the time the Data Collector last polled the management server(s). Click any Job ID # in the list to retrieve comprehensive detail about a specific job. The total number of running jobs displayed on the dashboard is limited. If the number of running and queued jobs exceeds the limit, all the running and queued jobs can be viewed in a separate report by clicking the Total Running and Queued Jobs link. The following jobs require attention, so the report highlights them in red and moves them to the top of the running and queued jobs list. Each of these flagged jobs will be accompanied by one or more icons to indicate the potential problem:

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Data Protection Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

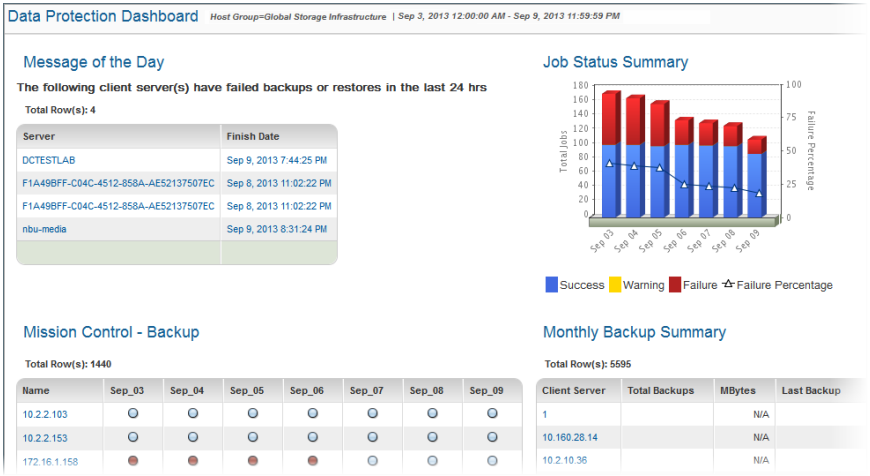
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Administration Reports > Data Protection Dashboard

The Data Protection Dashboard contains the following reports:

- Message of the Day
- Job Status Summary
- Mission Control - Backup
- Monthly Backup Summary

In the Data Protection Dashboard Report Designer, in addition to specifying a time period, you also can choose to omit backup retries.



Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Host Details

Detail reports are related to a specific enterprise object, such as a backup job or SAN fabric. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. Therefore, they will not be available in search results.

Use Quick Search to find the main report or a dashboard by name. Search is case insensitive, supports partial entries, and will display a list of potential matches.

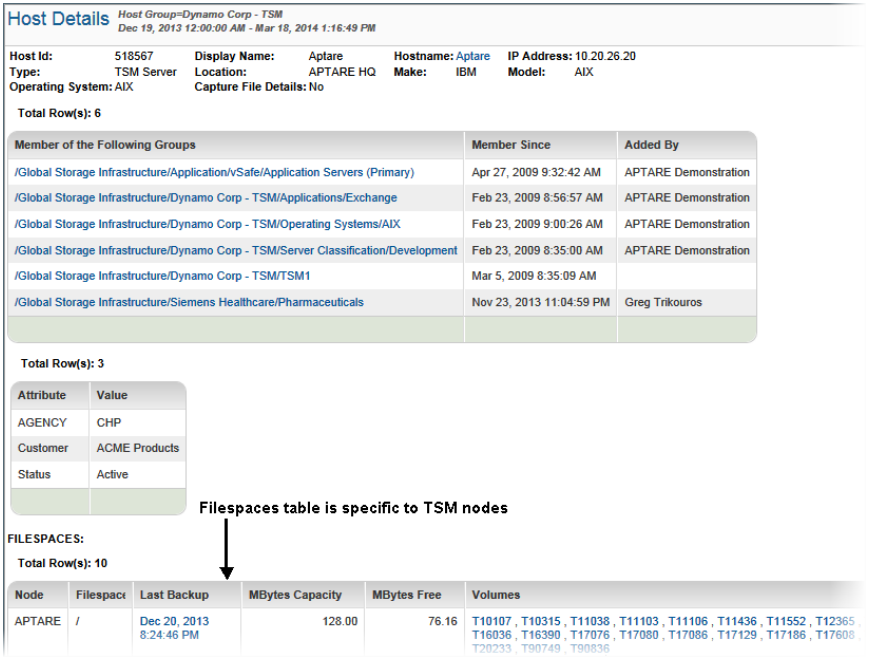
As you enter the report or dashboard name in the Quick Search field, up to 10 potential matches are displayed. If the report or dashboard is shown, you can select and run it directly from the match list. You can also select All Items to display the full search results page and further filter your results.

Access this report from numerous main reports using either Client or Server name links, such as:

- Job Summary Report
- Consecutive Errors by Client
- Disk Usage and Performance
- Billing and Chargeback Summary

Click the server name to go to the Host Details screen for additional information.

Note: The output is specific to the server type, as shown in the following examples. For VMware ESX servers, the host operating system is populated with the VM server's OS and version.



Host Details *Host Group=Dynamo Corp - TSM*
Dec 19, 2013 12:00:00 AM - Mar 18, 2014 1:16:49 PM

Host Id: 518567 **Display Name:** Aptare **Hostname:** Aptare **IP Address:** 10.20.26.20
Type: TSM Server **Location:** APTARE HQ **Make:** IBM **Model:** AIX
Operating System: AIX **Capture File Details:** No

Total Row(s): 6

Member of the Following Groups	Member Since	Added By
/Global Storage Infrastructure/Application/vSafe/Application Servers (Primary)	Apr 27, 2009 9:32:42 AM	APTARE Demonstration
/Global Storage Infrastructure/Dynamo Corp - TSM/Applications/Exchange	Feb 23, 2009 8:56:57 AM	APTARE Demonstration
/Global Storage Infrastructure/Dynamo Corp - TSM/Operating Systems/AIX	Feb 23, 2009 9:00:26 AM	APTARE Demonstration
/Global Storage Infrastructure/Dynamo Corp - TSM/Server Classification/Development	Feb 23, 2009 8:35:00 AM	APTARE Demonstration
/Global Storage Infrastructure/Dynamo Corp - TSM/TSM1	Mar 5, 2009 8:35:09 AM	
/Global Storage Infrastructure/Siemens Healthcare/Pharmaceuticals	Nov 23, 2013 11:04:59 PM	Greg Trikouros

Total Row(s): 3

Attribute	Value
AGENCY	CHP
Customer	ACME Products
Status	Active

Filespaces table is specific to TSM nodes

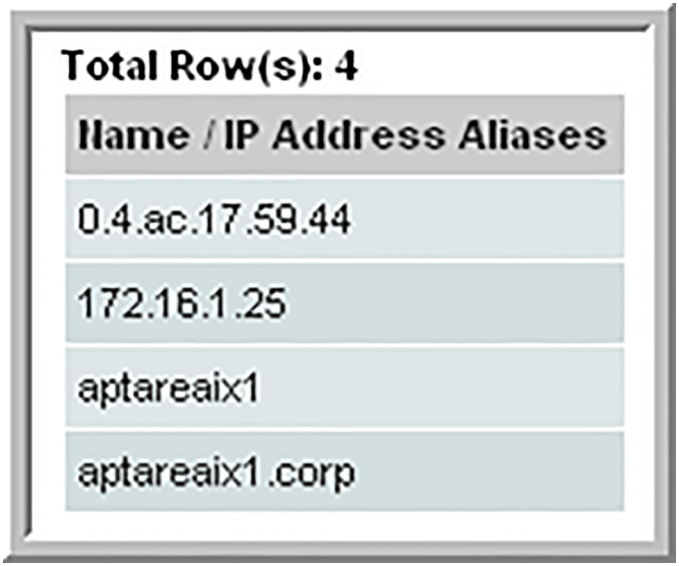
FILESPACES:

Total Row(s): 10

Node	Filespace	Last Backup	MBytes Capacity	MBytes Free	Volumes
APTARE	/	Dec 20, 2013 8:24:46 PM	128.00	76.16	T10107 , T10315 , T11038 , T11103 , T11106 , T11436 , T11552 , T12365 , T16036 , T16390 , T17076 , T17080 , T17086 , T17129 , T17186 , T17608 , T20233 , T90749 , T90836

Host Aliases

Some hosts/servers also may show a list of host aliases in this Host Details report:



Edit Host Details

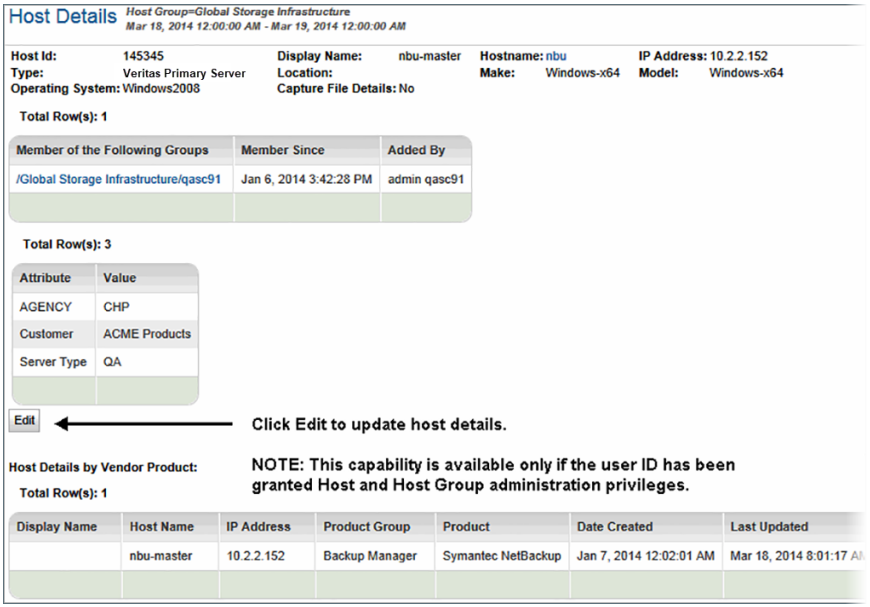


Table 23-4 Edit Host Details

Column name	Description
Host Id	The internal numeric ID of the server
IP Address	The server's IP address
Location	Server location, as supplied by the System Administrator
Make	Vendor name, such as IBM or Sun
Model	The product's alphanumeric identifier
Operating System	The specific version of the operating system. For VMware ESX servers, the host operating system is populated with the VM server's OS and version. This enables aggregation of ESX servers for more relevant host inventory reporting in ESX environments.
Host Group Membership	Click a group link to access the host group listing. See " Host Group " on page 842.
Filespaces	The filespace section at the bottom of this report provides the following data (for TSM only): Node lists the name of the physical server. Filespace lists the directory of the filespace. Last Backup provides a link to the last backup's job details. MBytes Capacity displays the capacity of the filespace. MBytes Free shows the available space in the filespace. Volumes: Click links to access specific volumes to display the Tape Media Detail. Storage Pool lists the Storage Pools to which the Node belongs.

Host Details by Source

Hosts can be collected by any number of sources such as vendor products. This table lists the details that have been collected by a source so that you can validate the collected details. You may want to re-prioritize how hosts are being listed, for example, prioritized by the source that is important for your reporting needs. To re-prioritize hosts, refer to the following.

Table 23-5 Host Details by Source

Column name	Description
Display Name	Name that is displayed in reports.
Host Name	Name of the host as it was collected by the vendor product.
IP Address	IP address of the host.
Product Group	The group, such as Capacity, for which the host data was collected.
Source	A specific source from which host data was collected; for example, VMware.
Date Created	The date and time the host was created in the Portal database.
Last Updated	The date and time the host data was updated in the Portal database.
OS Platform	Operating system of the host, as collected from this vendor product; for example, Windows Server 2008 R2.
OS Version	The host's specific operating system version, as collected from this vendor product; for example, 6.01.
Make	The host's make, as collected from this vendor product; for example, Windows-x64.
Model	The host's model, as collected from this vendor product; for example, PowerEdge 2950.
Backup Server Type	The host's role in a backup relationship, client or server.

Host Group

This report can be accessed from the Host Details.

See " Host Details " on page 838.

In the Host Details Report, click a Group Member link to list the clients/servers belonging to that host group.

The Host Group table provides the following information. Note that the columns that are shown may vary, depending on the type of host group:

Table 23-6 Host Group

Column name	Description
Client	Links to the Host Details Report.
Server	Links to the Host Details Report for the backup server.
Last Backup	Links to the Job Summary Report.
Backups this month	The number of the client's backups for the month.
Total Available (MB)	Amount of available space on this client
Capacity (MB)	The host's capacity, if it can be determined. In the case of backup clients, the capacity may be unknown to the Portal.
Backup Product	The backup software that performed the backups.
Node	For IBM Spectrum Protect (TSM) servers, the node name

IBM Spectrum Protect (TSM) Storage Pools Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Administration Reports > IBM Spectrum Protect (TSM) Storage Pools Dashboard

The IBM Spectrum Protect (TSM) Storage Pools Dashboard provides a real-time view into capacity and process status information for the selected list of Storage Pools. In a single pane, you quickly can analyze capacity and process status information for each of the Storage Pools.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Storage Pool Capacity

The pie chart displays the percentage and GBytes used and available for the Storage Pool. The GBytes Used is calculated by multiplying the percentage used by the estimated capacity of the Storage Pool.

Storage Pool Attributes

Table 23-7 Storage Pool Attributes

Column name	Description
Instance	Backup Instance name
Type	Primary, Copy
Device Type	Sequential, Random
Access Type	Read/Write, Read only
Reclaim Threshold	TSM can reclaim volumes that exceed a reclamation threshold--that is, a percentage of unused space on the volume. This threshold value is set by the administrator, within IBM Spectrum Protect (TSM).
Next Storage Pool	The storage pool that is next in line for backups when the current pool reaches capacity--that is, the destination storage pool for migrated data. IBM Spectrum Protect (TSM)automatically requests a storage pool when needed.
Max Scratch	The maximum amount of scratch space allocated to the storage pool. When the number of scratch volumes that IBM Spectrum Protect (TSM)is using for the storage pool exceeds the maximum number of scratch volumes specified, the storage pool can run out of space.
Capacity (GB)	Storage pool capacity
% Utilized	Percentage used
% Migrated	Estimated percentage of storage that can be migrated
Low Mig %	The value used by IBM Spectrum Protect (TSM)to start/stop migration. This percentage indicates when the server should stop migrating files to the next available storage pool.
High Mig %	The value used by IBM Spectrum Protect (TSM)to start/stop migration. This percentage indicates when the server should start migrating files to the next available storage pool.
Max Size (MB)	The maximum size of the storage pool.

Storage Pool Utilization

This historical tracking bar chart displays the history of used and available capacity for the selected storage pool.

Tape Occupancy Summary

This pie chart displays a summary of the GBytes and number of tapes used by each client/server that occupy this Storage Pool. The chart displays the top eight tape media consumers of this Storage Pool. Click on any of the client/server names to view detailed tape media consumption information

Tape Media Summary

This pie chart shows a summary of all the tapes assigned to this Storage Pool. This pie chart displays both the percentage of tapes grouped by media status, and the number of tapes grouped by media status:

- Filling, Full, Pending. Empty, Unknown, StgPool, Private, Pending, Questionable

Storage Pool Process History

This historical tracking bar chart displays the summary status of any TSM process that used this Storage Pool. Examples of TSM processes that are plotted on this bar chart are: Storage Pool Backups, Migrations, Expirations. Click on a bar in this graph to view the summary.

See " TSM Process Summary " on page 865.

Job Histogram

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use **Search** to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

To view the report, navigate to *Backup Manager* > *Administration Reports* > *Job Histogram*.

From the Job Histogram Report Designer, you can select one of the following options:

- Job Throughput Histogram
- Job Activity Histogram

These reports provide hour-by-hour data for the previous 24-hour period, by default.

Using these reports, you can identify bottlenecks and adjust your environment accordingly.

This generated report is based on clients of Backup Server(s). Select one or more clients in the **Report Scope** to filter the report data.

Currently the following parameters are available in the **Report Scope** dialog box.

Table 23-8 Report Scope parameters

Parameter	Description	Remark
Host Group	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for all the Backup Hosts or any of its sub-set.	
For Host = All Backup Servers	One or more Backup Servers can be selected from the scope to run the report for all clients corresponding to the Backup Server.	For example: Click **Modify**. Expand *Hosts* > *All Backup Servers* > *Primary Servers* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers.	Not valid for this report.	
Attributes	This report can be filtered on specific Host Attributes associated with the Backup Host(s).	

Note: If multiple conditions are selected in the **Report Scope**, NetBackup IT Analytics displays collective data with the selected criteria from the database using the **AND** condition.

Based on the parameters selected in the **Report Scope**, the following are the scenarios in which the data for the report will be displayed .

- Scenario 1 - **All Hosts**: If any one Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Client-Backup-Server-Host-Group**: You have a host group named as **All-Client-Backup-Server-Host-Group**which is the parent node.

Further, this **All-Client-Backup-Server-Host-Group**node has Two Backup Servers as*Server1*and*Server2*. If you have selected**All-Client-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be displayed for the*clients of Server1* and *Server2*.

- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

Job Throughput Histogram

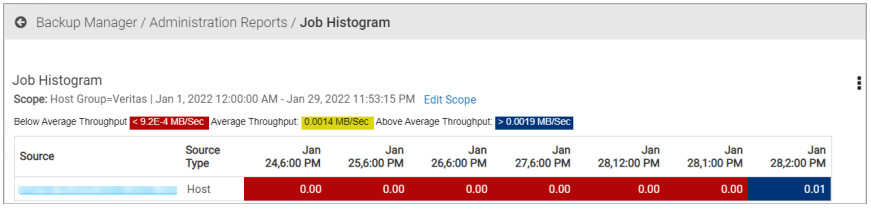
This report provides colored visual cues for quickly determining the backup throughput trends for clients that had backup jobs. This report shows completed jobs that started within the selected time period. These jobs could have a completion status of either successful or warning.

For example:

- **Red** - Indicates below average throughput--less than 5 MB/second
- **Yellow** - Indicates average throughput--greater than 5 MB/second and less than 9 MB/second
- **Blue** - Indicates above average throughput--greater than 9 MB/second

---

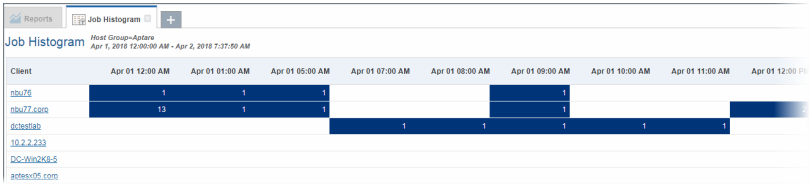
Note: The breakdown is based on the average, not the mean throughput.



Job Activity Histogram

The Job Activity Histogram lists clients that had backup jobs. The colored cells enable you to quickly determine the activity per client.

The report shows the number of completed jobs that started within the selected time period. These jobs could have a completion status of: successful, failed, or warning.



Jobs Dashboard

The Jobs Dashboard provides a view of your backup environment including job status, backup volume, error codes, job types, job volume and job duration.

You can access this dashboard from **Backup Manager** > **Administration Reports** > **Jobs Dashboard**.

Use this dashboard to view your entire backup environment or a subset based on your scope all in a single window. The charts in the dashboard provide drill down reports which give you more detailed information about each data point. The dashboard scope is set at last 72 hours by default, but you can change the scope as needed. You can also zoom-in and adjust other areas of the dashboard and individual reports.

Message of the Day Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

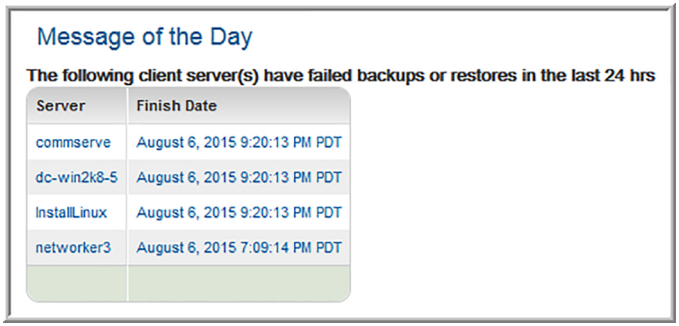
Backup Manager > Administration Reports > Data Protection Dashboard

This report provides a starting point for handling issues that require immediate attention. Specifically, it lists the failed backups and restores per server. From these lists, you can drill down to the details.

See " Host Details " on page 838.

See " Job Summary Report " on page 750.

See " Data Domain Replication History " on page 250.



1. Click a link under the Server column to access the details.

   See " Host Details " on page 838.

2. Click a **Finish Date** link to access Job Summary reports.

Mission Control - Backup

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

To view the report, navigate to *Backup Manager* > *Administration Reports* > *Mission Control - Backup*.

The Mission Control report provides you with a comprehensive snapshot of your backup and restore event status. Using this report, you can easily identify the sources (objects or clients) that have had no backups or failed backups. And, you can drill down to ascertain the root cause of unsuccessful backups. Organizations use this report to ensure Sarbanes-Oxley compliance, as it shows ultimate backup success over time.

When you generate this report, the Report Designer prompts you to select scope parameters, including the options to:

- Ignore retries
- Exclude clients with successful backups

- Exclude clients with no backups -Only the sources with at least one “No Jobs” or “Failed” indicator will be shown. Note that if the last column corresponds to the current day and the indicator is “No Jobs” and all other indicators for the row are successful, the row will be considered a “success” and therefore will be excluded from the report.
- By Primary Server only- Report on Primary Servers only. For example, if a source is backed up by two Primary Servers, you can avoid having duplicate jobs in the list.

This generated report is based on clients of Backup Server(s). Select one or more clients in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

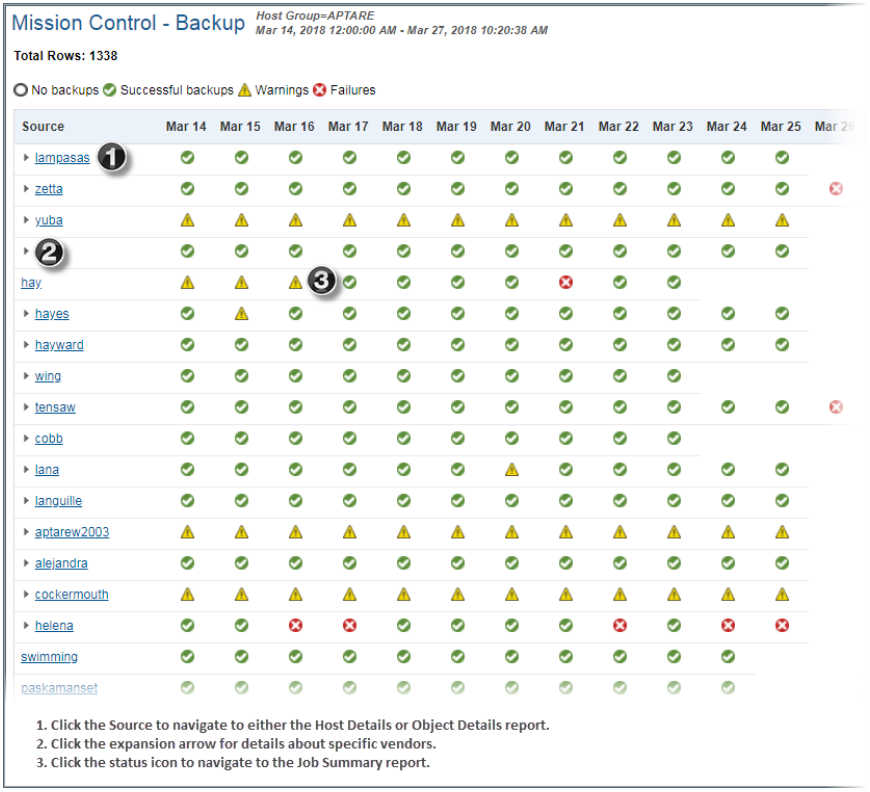
Table 23-9 Report Scope parameters

Parameters	Description	Remark
Host Group	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for all the Backup Hosts or any of its sub-set.	
For Host = All Backup Servers	One or more Backup Servers can be selected from the scope to run the report for all clients corresponding to the Backup Server.	For example: Click **Modify**. Expand *Hosts* > *All Backup Servers* > *Primary Servers* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	This report can be filtered on specific Host Attributes associated with the Backup Host(s).	

Note: If multiple conditions are selected in the **Report Scope**, NetBackup IT Analytics displays collective data with the selected criteria from the database using the **AND** condition.

Based on the parameters selected in the **Report Scope**, the following are the scenarios in which the data for the report will be displayed,.

- Scenario 1 - **All Hosts**: If any one Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Client-Backup-Server-Host-Group**: You have a host group named as **All-Client-Backup-Server-Host-Group**which is the parent node. Further, this **All-Client-Backup-Server-Host-Group**node has Two Backup Servers as*Server1*and*Server2*. If you have selected**All-Client-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be displayed for the*clients of Server1* and *Server2*.
- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).



Note: A row can be expanded to list underlying details for NetBackup, and HP Data Protector servers. For IBM Spectrum Protect (TSM) and Backup Exec, an expansion arrow will not be available. Also, the Flashback policy type (used for VM backups) is not supported; however, an expansion arrow may be displayed, with no underlying details available.

Table 23-10 Mission Control - Backup

Column name	Description
Source	Displays a row for each client or object within the report scope. Click the source to drill down to either the Host Details or Object Details reports. Some Sources will have an expansion arrow next to them, to show the underlying file systems.

Table 23-10 Mission Control - Backup *(continued)*

Column name	Description
Current Date Range	The date range that is currently displayed in the report is always listed at the top of the report. If the date range is beyond the current day, or before your backup environment data started being pushed to the Reporting Database, the event indicator will be “No Jobs.”
Status Indicators	Each column for every source contains an event status indicator. The indicator always reflects the most serious status if there is more than one event for that source on that day. For example, if there are three successful events and one failure, the indicator will be a failure. Click on any of the indicators to go to the summary. See “ Job Summary Report ” on page 750. for all events associated with that server on that day. Once in the Job Summary Report, you can drill down to Job Details and even add a note to indicate the strategy for resolving an issue--thus, providing an audit trail.

NetBackup Event Notification Summary - CRITICAL

Use this report to view a summary of critical event notifications specific to each server, along with their priority, severity, and count. The report also provides a drill-down report of each server to perform an in-depth review of each server and its related notifications.

This report is located here: **Backup Manager** > **Administration Reports** > **NetBackup Event Notification Summary - Critical**.

Table 23-11 NetBackup Event Notification Summary - CRITICAL

Column name	Description
Server Name	Name of the NetBackup primary server. The server name is clickable and opens a drill-down report.
Notification Priority	Priority of the notifications based on the server.
Notification Severity	Severity of the notifications based on the server.

Table 23-11 NetBackup Event Notification Summary - CRITICAL *(continued)*

Column name	Description
No. of Notifications	Count of notifications raised for each server. It is a clickable link and opens a drill-down NetBackup Event Notification Details – CRITICAL report.

NetBackup Event Notification Details - Critical

This is a drill-down report that contains additional details about the server-specific notifications. This report opens when you click a server name in the parent report.

Table 23-12 NetBackup Event Notification Summary - CRITICAL

Column name	Description
NetBackup Primary Server	Name of the NetBackup primary server.
Generated on (UTC)	ISO 8601 formatted time. The timestamp of when the notification was created.
Event Details	Event details for the notification. This string can contain some or all of the extended attributes of the notification.
Notification Id	ID assigned to the notification.
Notification Priority	Priority level associated with the critical notification.
Notification Type	Type of notification.
Producer Name	Host name of the notification producer.
Producer Id	Host ID (UUID) of the notification producer.
Producer Type	Host type of the notification producer.
Producer SubType	Sub-type of the notification producer.

NetBackup License Report

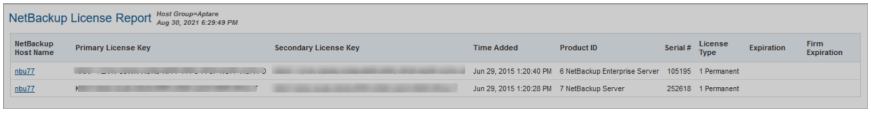
This report provides an overview of NetBackup license status in across the hosts configured in your environment. It helps you to avoid downtime due to expired licenses on your NetBackup Primary Servers.

This report is located at:

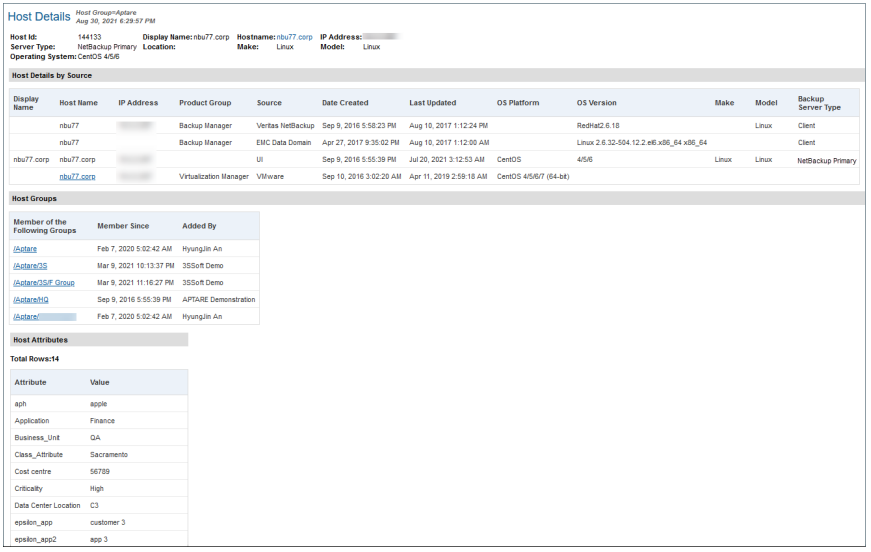
Backup Manager > **Administration Reports**

This report displays the key license information for each primary server selected, which includes:

- Name of the NetBackup host
- Primary and secondary license keys
- Date and time when the host was added
- Product ID
- License type
- License expiration date



You can click the host name to get a drill down view of host details by source, groups, and attributes.



NetBackup Host Certificates

There may be a number of services certificates used in your backup environment. This report shows the status of all the host certificates in your environment to help you evaluate the risk of their expiration.

To view the report, navigate to *Backup Manager* > *Administration Reports* > *NetBackup Host Certificates*.

This generated report is based on clients of Backup Server(s). Select one or more clients in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 23-13 Report Scope parameters

Parameters	Description	Remark
Host Group	The report can be generated for one or more Host Groups.	For more details on the Report Scope dialog box, see **NetBackup IT Analytics User Guide** > **Generate and maintain reports** > **Report Scope** section.
All Hosts	This report can be generated for all the Backup Hosts or any of its sub-set.	
For Host = All Backup Servers	One or more Backup Servers can be selected from the scope to run the report for all clients corresponding to the Backup Server.	For example: Click **Modify**. Expand **Hosts** >**All Backup Servers** > **Primary Servers**(if any) > **Policies** or **Clients**. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	This report can be filtered on specific Host Attributes associated with the Backup Host(s).	

Note: If multiple conditions are selected in the **Report Scope**, NetBackup IT Analytics displays collective data with the selected criteria from the database using the **AND** condition.

Based on the parameters selected in the **Report Scope**, the following are the scenarios in which the data for the report will be displayed .

- Scenario 1 - **All Hosts**: If any one Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Client-Backup-Server-Host-Group**: You have a host group named as **All-Client-Backup-Server-Host-Group**which is the parent node. Further, this **All-Client-Backup-Server-Host-Group**node has Two Backup Servers as*Server1*and*Server2*. If you have selected**All-Client-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be displayed for the*clients of Server1* and *Server2*.
- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

Table 23-14 NetBackup Host Certificates report columns

Column name	Description
Server	Host or server on which the certificate is installed.
Issued To	Alias name to which the certificate is issued.
Serial	Serial number of the certificate installed.
Key Size	Certificate encryption key size.
Valid From	The date from which the certificate was deemed valid.
Valid To	The date of certificate expiry.
Days Left	Number of days remaining for certificate expiry.
Last Used	Last date when the certificate was used.
Active	Indicates whether the certificate is still active.

NetBackup Non-Host Certificates

There are a number of services certificates that may expire causing risk to the backup environment. This report shows all non-host certificates at risk and includes server name, type certificates, size, issuer, owner, and date of expiration.

To view the report, navigate to *Backup Manager* > *Administration Reports* > *NetBackup Non-host Certificates*.

This generated report is based on clients of Backup Server(s). Select one or more clients in the **Report Scope** to filter the report data.

Currently, the following parameters are available in the **Report Scope** dialog box.

Table 23-15 Report Scope parameters

Parameters	Description	Remark
Host Group	The report can be generated for one or more Host Groups.	For more details on the **Report Scope** dialog box, see *NetBackup IT Analytics User Guide* > *Generate and maintain reports* > *Report Scope* section.
All Hosts	This report can be generated for all the Backup Hosts or any of its sub-set.	
For Host = All Backup Servers	One or more Backup Servers can be selected from the scope to run the report for all clients corresponding to the Backup Server.	For example: Click **Modify**. Expand *Hosts* > *All Backup Servers* > *Primary Servers* (if any) > *Policies* or *Clients*. Select any node.
Servers = All Data Domain Servers	Not valid for this report.	
Attributes	This report can be filtered on specific Host Attributes associated with the Backup Host(s).	

Note: If multiple conditions are selected in the **Report Scope**, NetBackup IT Analytics displays collective data with the selected criteria from the database using the **AND** condition.

Based on the parameters selected in the **Report Scope**, the following are the scenarios in which the data for the report will be displayed .

- Scenario 1 - **All Hosts**: If any one Host is selected from parent All Hosts, the data is displayed for the selected Host in the node.
- Scenario 2 - **All-Client-Backup-Server-Host-Group**: You have a host group named as **All-Client-Backup-Server-Host-Group**which is the parent node. Further, this **All-Client-Backup-Server-Host-Group**node has Two Backup Servers as*Server1*and*Server2*. If you have selected**All-Client-Backup-Server-Host-Group**, which is the parent node, then the data in the report will be displayed for the*clients of Server1* and *Server2*.
- Scenario 3 - **All Backup Servers**:
  - If **All Backup Server** node is selected, the data is displayed of the clients.
  - If any one or more than one **Clients** from All Backup Server node is selected, the data is displayed of the selected Client(s).
  - If any one or more than one **Policies** from All Backup Server node is selected, the data is displayed of the selected Policy(s).
- Scenario 4 - **Attributes**: If any one or more than one Attributes is selected from the Report Scope dialog box, the data is displayed of all Backup Hosts for the selected Attribute.
- Scenario 5 **All Data Domain**: This is invalid selection as this report is generated only based on Backup Host(s).

Table 23-16 NetBackup Non-Host Certificates report columns

Column name	Description
Primary Server	Name of the primary server hosting the certificates.
Alias Name	The user-assigned alias name.
Creation Date	Certificate creation date.
Type	Certificate type.
Chain Length	Number of certificates in the non-host certificate chain.
Owner	Certificate owner.
Issuer	Name of the entity that has issued the certificate.
Serial number	Serial number of the certificate.
Valid From	The date from which the certificate was deemed valid.
Valid Until	The date after which the certificate will expire.
Version	Certificate version.

Table 23-16 NetBackup Non-Host Certificates report columns *(continued)*

Column name	Description
Days left	Number of days remaining for certificate expiry.
Keystore	The name of the keystore where the certificates are stored.

Object Details

You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. Therefore, they will not be available in search results.

Use Search to find the main report or a dashboard by name. Search is case insensitive, supports partial entries, and will display a list of potential matches.

The Object Details report provides information about backup sources other than known hosts, such as datastores, folders and so on. Access this report from numerous main reports using the Source link, such as:

- Job Summary Report
- Consecutive Errors by Client
- Data Protection Dashboard
- Failed Backup Sources
- Billing and Chargeback Summary

Table 23-17 Object Details

Column name	Description
Object Name	Name of the backup source object.
Object Identifier	Unique identifier for this protection source.
Root Object Name	If a hierarchy structure exists, displays the name of the top most parent of the object. If a hierarchy does not exist, for example in a file or a directory, this displays the object name.
Creation Date	Object creation date.
Last Collected Date	The date the last collection occurred.
Object Type	Defines the object type such as resource pool, datastore or a folder.

Table 23-17 Object Details *(continued)*

Column name	Description
VM Host Name	Host name of the virtual machine which is VMWare host name. This is dispay only when VMWare is installed.
Last Backup Date	Displays the date for the last backup of the object.
Last Successful Backup	Displays the date for the last successful backup of the object.
Protected By	
Policy Name	The name of protection policy associated to the protected object. A protection policy contains the list of objects to be protected. The protection policy is associated with a policy template that describes the backup retention and archive schedules.
Protected by Object Name	Name of the object protected by the policy. In case of the hierarchy, this protected object may result in protecting other objects.
Protect by Object Type	Product specific type of this protection source. For example, virtual machine, resource pool, and so on.
Policy Last Backup Date	Last backup date for the policy.
Is Enabled	Indicates if this policy is enabled.
Enabled Date	Date when this protection policy was enabled.
Retention Days	Specifies the number of days a backup (based on this policy template) needs to be retained.
Schedule Name	Schedule name configured for this policy template. The schedule name uniquely identifies the schedules available in the subsystem.
Schedule Type	Type of this schedule. For example, Daily, Weekly, Monthly, Periodically, Continuous', and ScheduleAfterJob.
Storage Resource Name	Name of the storage resource where the backups of this protection policy will be saved.
Storage Resource Type	Type of this storage resource. For example, Backup Repository or Scale Out Backup Repository.
Storage Resource Capacity	Total capacity of this storage resource in KiB.

Table 23-17 Object Details *(continued)*

Column name	Description
Storage Resource Used Capacity	Used capacity of this storage resource in KiB.

Operations Dashboard

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Administration Reports > Operations Dashboard

The Operations Dashboard provides you with multiple views of your backup environment on a single page.

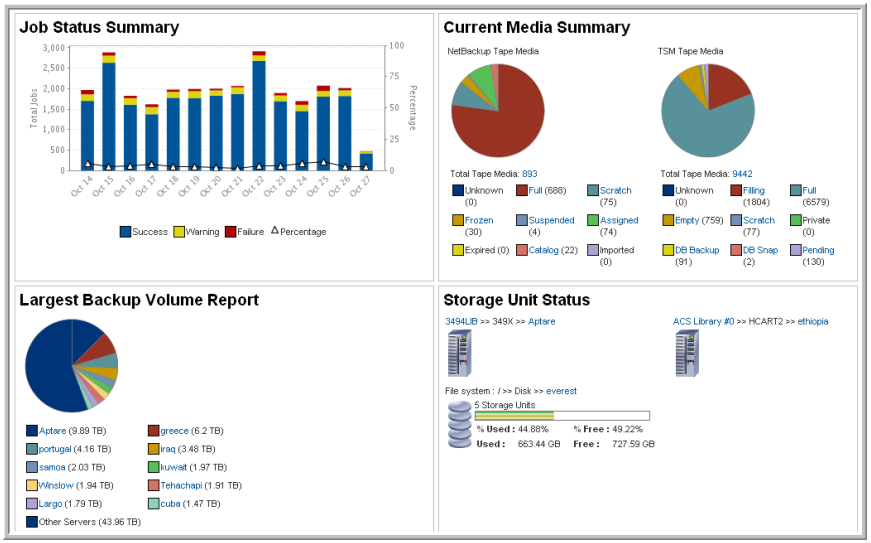


Table 23-18 Operations Dashboard

Report	Description
1. Job Status Summary	Displays the number of successful, warning and error events for each time period. The dashed yellow line displays the SLA success goal while the triangle symbols track actual achievement for each day. Clicking on any of the bars will take you to the summary for the corresponding event status and time frame. See " Job Summary Report " on page 750.
2. Current Media Summary	Provides a real-time snapshot of all the media utilized by the servers within the scope. Mouse-over any of the pie-chart sections to display pop-ups that provide pertinent details about the media of the corresponding status. Click on any section of the pie-chart or the to go to the summary for all media with the corresponding status. See " Tape Media Summary Report " on page 875. Click the Total Tape Media link to go to the Tape Media Summary report for all the servers within the report scope.
3. Largest Backup Volume Report	Provides a real-time snapshot of the largest server backup consumers within the report scope. Mouse-over any of the bars or symbols on the chart to display pop-ups that provide details. Click on any section of the pie-chart to go to the summary for a particular server. See " Server Consumption Summary " on page 938.

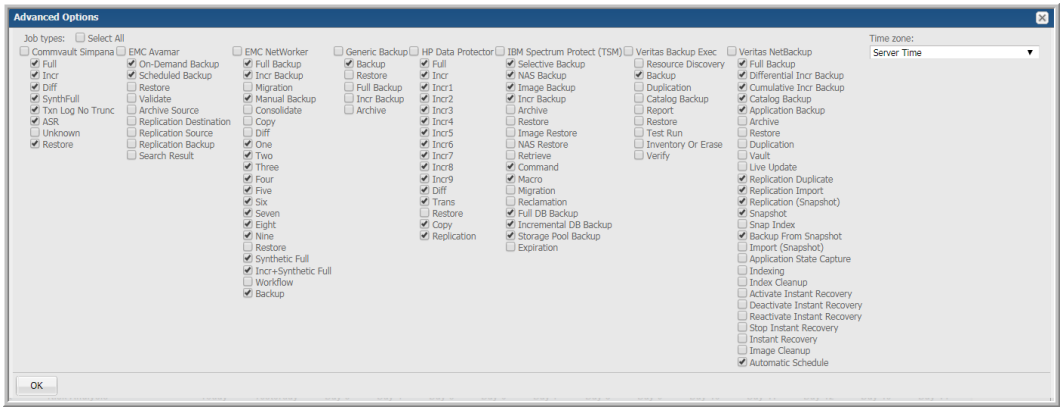
Table 23-18 Operations Dashboard *(continued)*

Report	Description
4. Storage Unit Status	Click on the Total Servers link to go to the Server Consumption Summary report for all the servers within the report scope. Each file system is identified by the file system path, with its management server displayed above the file system graphic. The total number of Storage Units sharing the same file system is also shown. The display of this report is dependent on the backup products in your environment and the report scope. Each file system is accompanied by a thermometer, which indicates the percentage of space occupied and available on the file system. The Used value is as reported by the operating system, and includes space occupied by all backup images as well as any other files that may reside on that file system. The Free value also is reported by the operating system. Used and free values of **Unknown** indicate the Discovery module has not been configured to probe the media servers' file systems for their physical characteristics.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Tailoring the Report Output by Backup Product

In the scope selector, click **Advanced**.



Note that the options available in this window depend on the backup product data collection options that have been enabled for your Portal.

TSM Process Summary

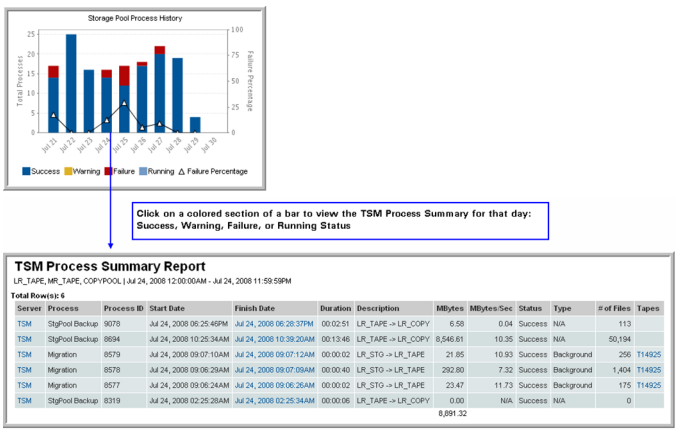
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Administration Reports > TSM Storage Pools Dashboard

Click a bar in a Storage Pool Process History graph.



This report provides indicators of the type of job, as well as, the number of files and the relevant tape IDs.

NetBackup Audit Report

In NetBackup, the audit report provides extensive audit details of Policies, Jobs, storage unit changes and so on.

The reports module in NetBackup IT Analytics is enhanced with a new report **NetBackup Audit report** to fetch these audit details.

This report is located here: **Reports** > **Backup Manager** > **Administration Reports** > **NetBackup Audit Report**

Note: This report is available only to System Administrator users.

NetBackup compatibility and pre-requisites

- NetBackup IT Analytics supports NetBackup version 8.1.2 and above.
- Audit feature should be enabled on NetBackup.
- NetBackup IT Analytics uses SSH/WMI Based collection method for the report.
- Configure **Audit Events** under **Veritas NetBackup Data Collector policy** screen.

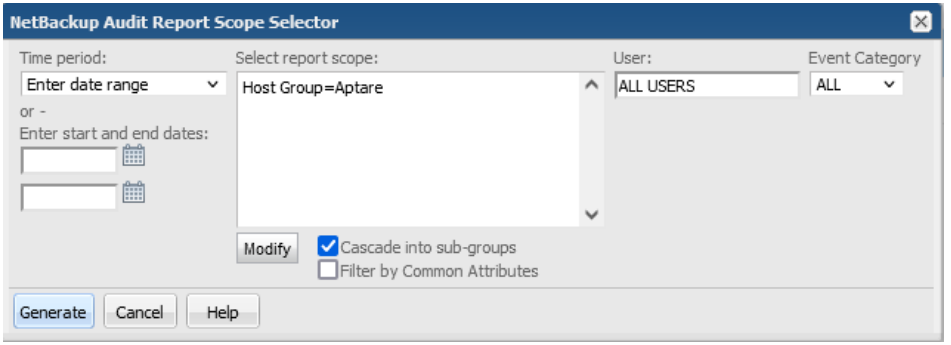
  Note: For more information, see *NetBackup IT Analytics Data Collector Installation Guide for Backup Manager >> Pre-Installation setup for Veritas NetBackup >> Adding a Veritas NetBackup Data Collector policy* section.

NetBackup Audit Report scope selector

Using the scope selector, a system administrator user can define the scope of the report.

The following are the parameters available for the reports.

- Time Period
- Select report scope
- User
- Event category
- Cascade into sub-groups
- Filter by common attributes



NetBackup Audit Report

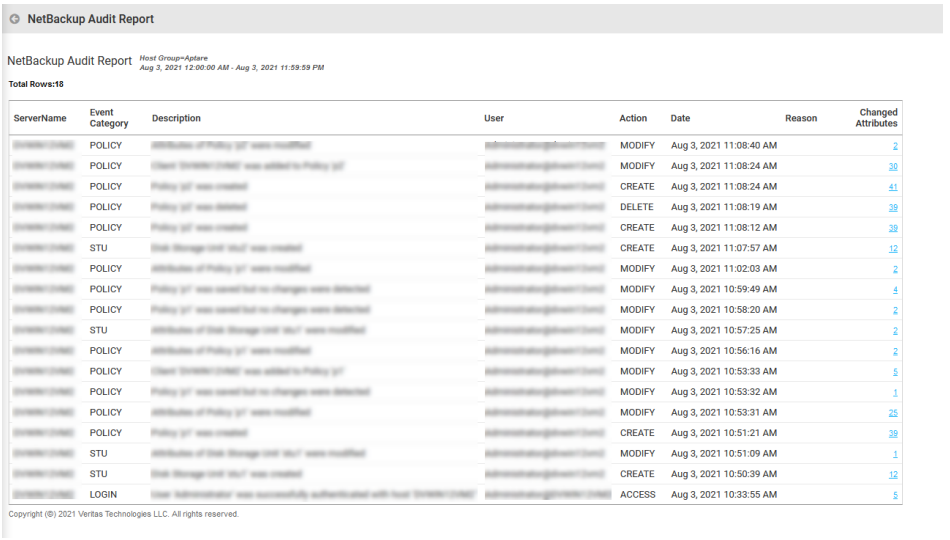


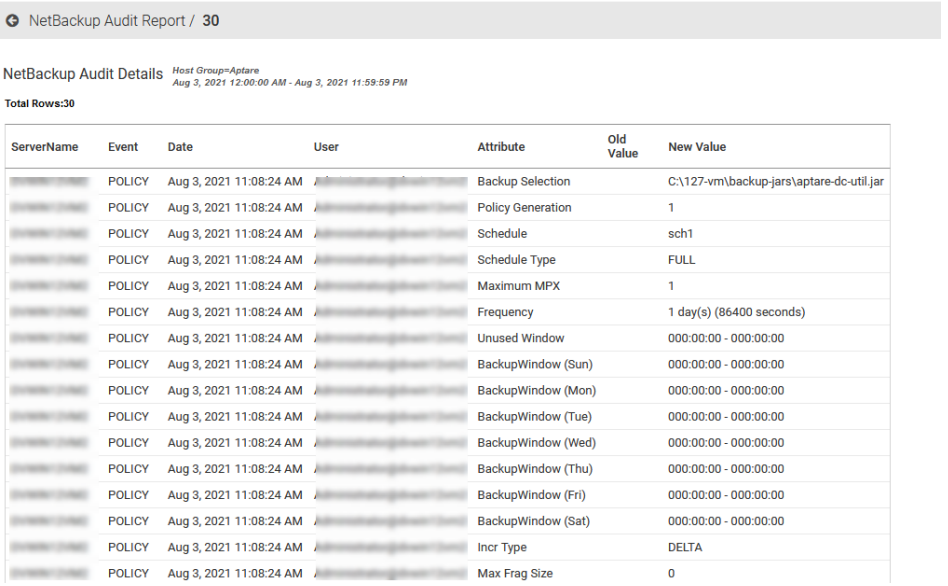
Table 23-19 NetBackup Audit report

Column name	Description
Server Name	The NetBackup Server name that has been updated or modified.
Event Category	Event category of the logged event. Some of the examples are – LOGIN, LICENSING, and POLICY.
Description	A summary of the type of event that occurred.
User	Username of the user who performed the audited action.

Table 23-19 NetBackup Audit report *(continued)*

Column name	Description
Action	The type of action taken by the user that generated the event. Some of the actions are – ACCESS, MODIFY, CREATE, and DELETE.
Date	Timestamp of event logged.
Reason	The reason for the generated event.
Changed Attributes	The number of changed attributes for the generated event. It redirects to “NetBackup Audit Details” report for providing more details about Changed Attributes.

NetBackup Audit drill down details report


NetBackup Audit Report / 30

NetBackup Audit Details *Host Group=Aptare*
Aug 3, 2021 12:00:00 AM - Aug 3, 2021 11:59:59 PM

Total Rows:30

ServerName	Event	Date	User	Attribute	Old Value	New Value
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Backup Selection		C:\127-vm\backup-jars\aptare-dc-util.jar
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Policy Generation		1
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Schedule		sch1
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Schedule Type		FULL
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Maximum MPX		1
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Frequency		1 day(s) (86400 seconds)
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Unused Window		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Sun)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Mon)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Tue)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Wed)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Thu)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Fri)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	BackupWindow (Sat)		000:00:00 - 000:00:00
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Incr Type		DELTA
[redacted]	POLICY	Aug 3, 2021 11:08:24 AM	[redacted]	Max Frag Size		0

EEBs on Primary server, Media server and Clients reports

The new REST API, in NetBackup v10.2 release, displays the list of the installed EEBs in the NetBackup domain, including primary server, media server, and clients.

The following are the two APIs:

- /config/all-hosts - Lists the installed EEBs per host
- /deployment/eebs - Lists all the EEBs installed in the NetBackup Domain as a single list.

NetBackup IT Analytics persists the data of the installed EEBs for all the hosts in the domain. The following reports are incorporated to display the details.

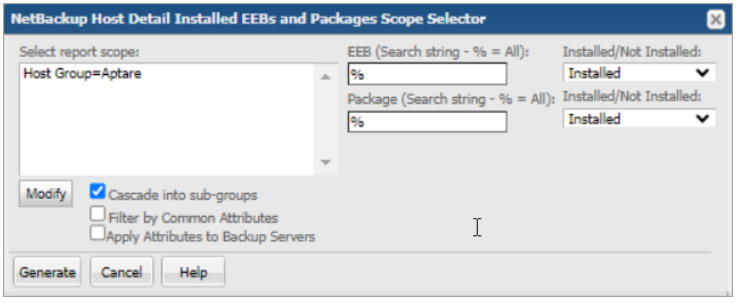
Navigate to *Reports* >> *Backup Manager* >> *Administration Reports* >> **NetBackup Host Detail Installed EEBs and Packages** report

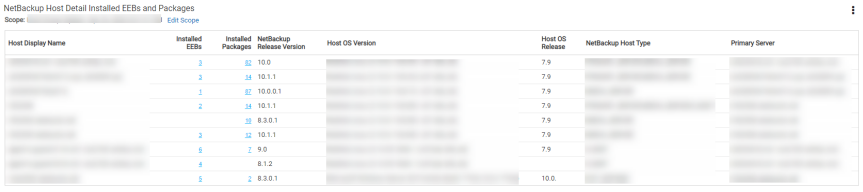
NetBackup Host Detail Installed EEBs and Packages

The report displays Installed EEBs and Packages for the selected hosts, providing a count of EEBs and Packages, and other Host version information.

NetBackup Host Detail Installed EEBs and Packages scope selector

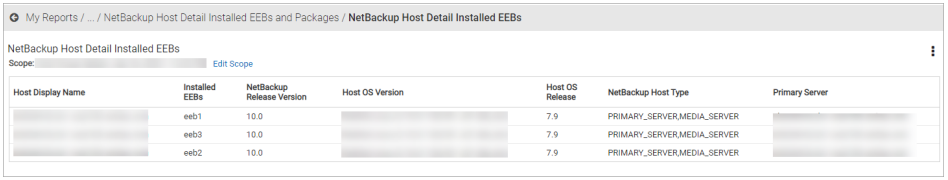
The scope selector of the NetBackup Host Detail Installed EEBs and Packages report searches for EEBs and Packages that match the search criteria specified in both **EEB** and **Package** fields. The search criteria applies AND condition between **EEB** and **Package** fields while generating the report. The search string **%** is a wildcard for all characters. You can further refine the scope for **Installed** and **Not Installed** EEBs and Packages. The resulting report contains details that match the search criteria defined in both **EEB** and **Package** fields.





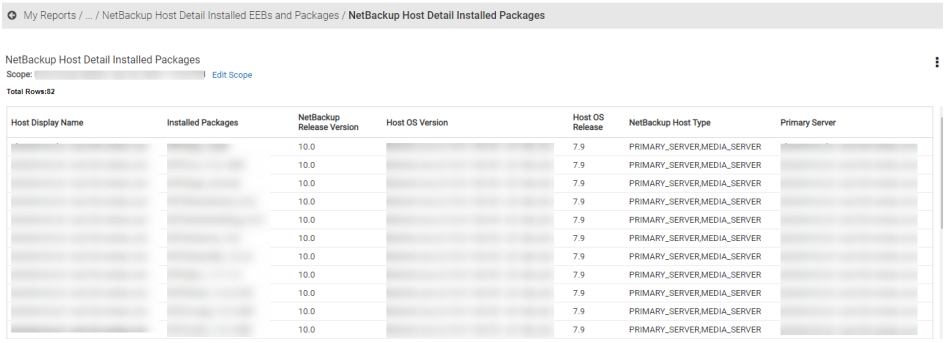
NetBackup Host Detail Installed EEBs

NetBackup Host Detail Installed EEBs displays all Installed EEBs for the selected host, including Host version information. Click the drill down hyperlink on the main report NetBackup Host Detail Installed EEBs and Packages to view this report.



NetBackup Host Detail Installed Packages

NetBackup Host Detail Installed Packages displays all Installed Packages for the selected host, including Host version information. Click the drill down hyperlink on the main report NetBackup Host Detail Installed EEBs and Packages to view this report.



Chapter 24

Backup Media Management Reports

This chapter includes the following topics:

- Current Media Summary
- Data Domain VTL Errors
- Data Domain VTL Media Summary
- Data Domain VTL Performance
- Tape Media Summary Report
- Tape Media Detail Report
- Commvault Simpana Tape Media Detail
- HP Data Protector Tape Media Table
- NetBackup Tape Media Table
- NetWorker Tape Media Table
- TSM Tape Media Detail
- TSM Tape Media Detail Table

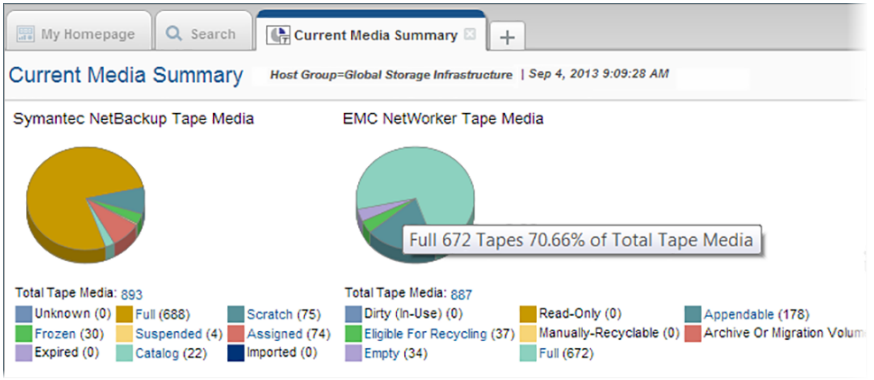
Current Media Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Media Management Reports > Current Media Summary

This report is applicable to only NetBackup, Networker, TSM, Commvault Simpana, and HP Data Protector.



The Current Media Summary is a real-time snapshot of all the media utilized by the backup solutions for servers within the scope of the report. The status types vary for each pie chart because they are directly related to the specific backup solution. So, for example, NetBackup reports expired tapes, while NetWorker indicates which tapes can be recycled. Refer to the backup solution's reference guides for details about media status types.

- Hold your mouse over any of the pie chart sections and a pop-up will display details about the media of the corresponding status. Similar information is reflected in the chart legend.
- Click any section of the pie chart and/or legend labels to view the Tape Media Summary report.
- Click the **Total Tape Media** link to view the Tape Media Summary report.

Data Domain VTL Errors

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Media Management Reports > Data Domain VTL Errors

View a summary of error statistics for all VTLs in a Data Domain system.

Table 24-1 Data Domain VTL Errors

Column name	Description
System	The Data Domain system
Library	Virtual Tape Library name
Model	Model of the library
# Drives	Number of drives in the VTL
# Slots	The number of slots in the library; note that the default is 20, with a max of 20,000 for all VTLs
Cartridge Access Ports	The CAPs (cartridge access ports) in the system
# Soft Errors	The number of errors from which the VTL recovered. Typically, no actions are required for these types of errors, although excessive errors could impact performance.
# Hard Errors	The number of errors from which the VTL was unable to recover. Refer to Data Domain logs to determine if corrective actions need to be taken.
Last Error On	Date and time of the last error

Data Domain VTL Media Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Media Management Reports > Data Domain VTL Media Summary

Table 24-2 Data Domain VTL Media Summary

Column name	Description
System	The Data Domain system
Media ID	Media ID
Library	The tape library name See " Data Domain VTL Media Summary " on page 873.
Pool	The configured tape pool
Location	Location of the media
State	Tape state (RW/RO/WO) is collected only if data collection is configured to use SNMP.
Total	Tape size
Available	Amount of available space on the tape
% Available	Percentage of the tape that is still available

For NetBackup, some environments freeze or suspend their media and do not allow their tapes to fill to capacity. In this situation, a NetBackup IT Analytics database table can be updated to override the standard estimated capacity calculation with values for Volume Pool ID, Media Type, and Estimated Capacity. When values are supplied in this table, the NetBackup estimated capacity calculations will first determine the average size of full tapes--by tape type and volume pool. If this value is less than the value provided in the table, the value in the table will be used for the estimated capacity.

Refer to the *System Administrator Guide* for details.

Data Domain VTL Performance

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

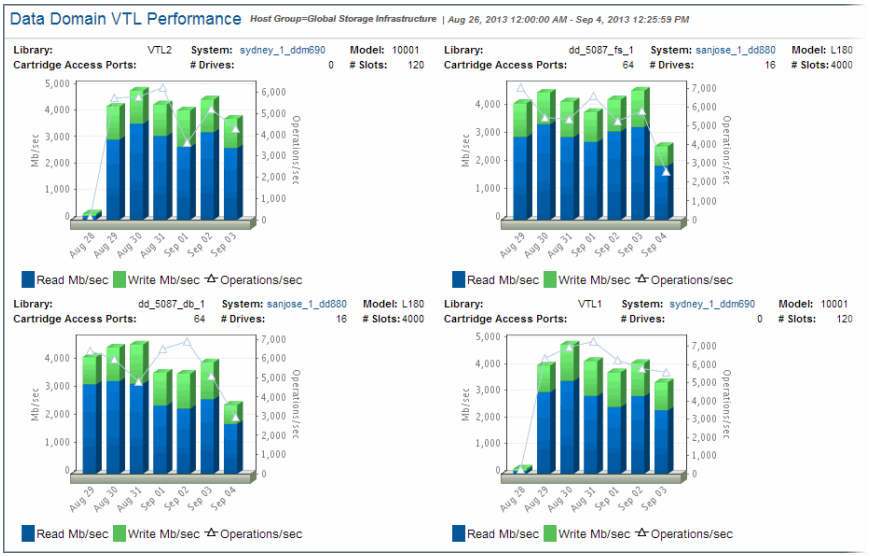
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Media Management Reports > Data Domain VTL Performance

The graphs in this report represent statistics for all the drives in a VTL: Reads Mb/sec, Writes Mb/sec, and Operations/sec.

Note: Click a green portion of a bar to view the status.

See " Data Domain Filesystem Status " on page 520.



Tape Media Summary Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Media Management Reports > Tape Media Summary

A more detailed version of this report is also available from the Operations Dashboard by clicking the Total Tape Media link in the Current Media Summary.

- See " Operations Dashboard " on page 862.
  See " Current Media Summary " on page 871.
- See " Tape Media Summary Report " on page 875.

This report is applicable to only NetBackup, Networker, IBM Spectrum Protect (TSM), Commvault Simpana, and HP Data Protector.

Depending on your environment, this report listing can get rather long, so you may want to narrow the report's scope. When you generate the report, you can limit the output using the following parameter choices:

- Scope: Select specific hosts and/or host groups
- Event Status: Select All Media, Media Inside Tape Library, or Media Outside Tape Library
- Media Status: All, Expired, Available, In Use, or Full. Note that the Expired status is not a valid option for IBM Spectrum Protect (TSM) and HP Data Protector, as those vendors do not support this status for tapes.

When your Tape Media Summary report contains multiple listings, the sortable headings enable you to find the tape media you want quickly. Click on any of the table headings to sort the table by that heading in ascending order. Click the heading a second time to sort in descending order.

Table 24-3 Tape Media Summary Report

Column name	Description
Media ID	The tape ID links to the Tape Media Detail
Server	The management Server links to the Host Details
Product	The backup product associated with the tape: Commvault Simpana, HP Data Protector, NetBackup, NetWorker, IBM Spectrum Protect (TSM)
Status	The tape media status is reported as: Assigned, Catalog, Cleaning, Filling, Frozen, Full, Scratch, Unavailable, or Vendor Specific (which indicates the disposition specific to the backup product). Place your mouse pointer on top of Vendor Specific to view the message.
Type	Type of tape device
Library	Library to which the tape belongs
Slot #	Physical position in the library
Slot	Commvault Simpana only. Name of the slot for the media.
Location	Commvault Simpana only. Location of the slot.
Volume Pool	NetBackup volume pool
Volume Group	NetBackup volume group

Table 24-3 Tape Media Summary Report *(continued)*

Column name	Description
Estimated Capacity (MB)	See "Estimated Capacity Notes" on page 883.
MBytes	Amount of media in use
Available (MB)	Amount of media storage still available
% Available	Percentage of the media that is still available
Expiration	Media expiration date
Slot #	Physical position in the library
Last Read	Date and time of last read
Last Written	Date and time of last write
Last Exported	Commvault Simpana only. Data & time of last export
Export Location	Commvault Simpana only. Destination of the export.

Tape Media Detail Report

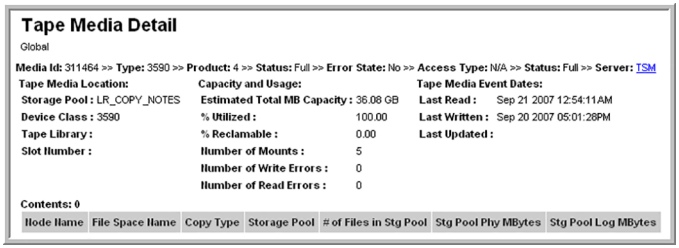
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Media Management Reports > Tape Media Summary

Click the Media or Volume ID.



The details displayed in a Tape Media Detail report are specific to the backup vendor, although some fields are common to all backup products. At the bottom of the Tape Media Detail report, a table lists additional backup event details by node name, session name, or save set--depending on the structure and terminology of the corresponding backup product. These tables are listed by backup product in this section.

Table 24-4 Tape Media Detail Report

Column name	Description
Media ID/Volume ID	Either the Media ID or Volume ID
Label	The tape label
Type/Media Type	The type of tape media, such as DLT or LTO-Ultrium
Product	The backup product using the tape media
Status	Tape status, such as Scratch, Good, Full, Filling, as reported by the backup software
Policy	HP Data Protector policy
Usage Allocation	The type of usage allocation
Error State	The tape media error state, as reported by the backup software
Access Type	Access type, such as Read Write or Read Only, as reported by the backup software
Server	The server managing the job that's transferring the data to tape
Volume Group	NetBackup volume group
Volume Pool	NetBackup volume pool
Tape Library	The name of the library in which the media resides

Table 24-4 Tape Media Detail Report *(continued)*

Column name	Description
Slot Number	The media's location in the library
Estimated Total MB Capacity	See "Estimated Capacity Notes" on page 883.
Written MBytes	The amount of data written to the media. This field is not displayed for TSM.
Available MBytes/Estimated MBytes Available	The amount of available space on the media. This field is not displayed for TSM. See "Estimated Capacity Notes" on page 883.
Number of Mounts	Number of times the media was mounted
Number of Write Errors	Number of media write errors
Number of Read Errors	Number of media read errors
Age Limit (months)	HP Data Protector age limit
Initialized Date	HP Data Protector initialized date
Last Read	Date and time of last read
Last Written	Date and time of last write
Last Updated	Date and time of last update
Last Overwrite Date	Date and time of last overwrite

Commvault Simpana Tape Media Detail

The following table is displayed for reports that have a Commvault Simpana host group scope.

Table 24-5 Commvault Simpana Tape Media Detail

Column name	Description
Media ID	The media's name.
Barcode	Barcode identification for the media.

Table 24-5 Commvault Simpana Tape Media Detail *(continued)*

Column name	Description
Type	Type of tape media, such as cleaning tape.
Status	Tape status, such as Scratch, Good, Full, Filling, as reported by the backup software
Server	The server managing the job that's transferring the data to tape.
Tape Media Location	Location for the tape media.
Location	Media's physical location.
Media State	Indicates the disposition of the media, such as "P" for Courier Retrieve.
Library	Library to which the tape belongs. See " Tape Media Summary Report " on page 875.
Slot	Physical position in the library.
Export Location	Media export location.
Capacity and Usage	
Capacity (MB)	Total capacity of the media.
Written (MB)	The amount of data written to the media.
Available (MB)	Amount of media storage still available.
# Backups	Number of backups that have occurred.
# Overwrites	Number of overwrites that have occurred.
# Errors	Number of media read/write errors.
Tape Media Event Dates	
Initialized	Date and time tape media was created.
Last Backup	Date and time tape media was last backed up.
Last Restore	Date and time tape media was last restored.
Last Exported	Date and time tape media was last exported.

HP Data Protector Tape Media Table

This table is displayed for reports that have an HP Data Protector host group scope. For a description of the fields in this table refer to the following.

See " HP Data Protector Session Summary " on page 741.

NetBackup Tape Media Table

The following table is displayed for reports that have a NetBackup host group scope.

Table 24-6 NetBackup Tape Media Table

Column name	Description
Backup ID	The NetBackup ID
Client	Client from which the data was copied
Backup Type	Type of backup to tape, such as Full or Partial
Start Time	Date and time the backup to tape began
Finish Time	Date and time the backup to tape ended
Expiration	Expiration date, if any
Written MBytes	Amount of data written to tape for this session
Exit Code	Exit code associated with the backup job

NetWorker Tape Media Table

The following table is displayed for reports that have a NetWorker host group scope.

Table 24-7 NetWorker Tape Media Table

Column name	Description
Save Set	NetWorker Save Set name
Client	Client from which the data was copied
Level	Type of tape media storage, which could be a level number or text, such as Full
Start Time	Start of Save Set write

Table 24-7 NetWorker Tape Media Table *(continued)*

Column name	Description
Finish Time	End of Save Set write
Expiration	Expiration date, if any
Written MBytes	Amount of data written to the tape
State	Tape media state, such as Browsable or Recoverable

TSM Tape Media Detail

The following table is displayed for reports that have a TSM host group scope.

Table 24-8 TSM Tape Media Detail

Column name	Description
Location	The media's physical location
Storage Pool	Lists the associated Storage Pool
Device Class	Lists the device class of the tape media
Tape Library	The name of the library in which the media resides
Slot Number	The media's location in the library
Estimated Total MB Capacity	Every time tape data is captured, the system recalculates the estimated tape capacity for the tape media type and volume pool combination. It takes into consideration the written KBs for the media type and volume pool for full tapes and then stores the average as the estimated tape capacity.
% Utilized	The percentage of the media that is in use.
% Reclaimable	The percentage of the media that may be reclaimable.
Number of Mounts	Number of times the media was mounted
Number of Write Errors	Number of media write errors
Number of Read Errors	Number of media read errors
Media State	Indicates the disposition of the media, such as "P" for Courier Retrieve.

Table 24-8 TSM Tape Media Detail *(continued)*

Column name	Description
Last Read	Date and time of last read
Last Written	Date and time of last write
Last Updated	Date and time of last update

TSM Tape Media Detail Table

The following table is displayed for reports that have a TSM host group scope.

Table 24-9 TSM Tape Media Detail Table

Column name	Description
Node Name	Name of the TSM node/server
File Space Name	Name of the TSM file space that was copied to tape
Copy Type	Type of transfer to tape
Storage Pool	The TSM Storage Pool name
# of Files in Stg Pool	Number of files in the TSM storage pool for this node
Stg Pool Phy MBytes	Amount of data in the TSM storage pool for this node
Stg Pool Log MBytes	Size of the TSM storage pool log file

Estimated Capacity Notes

Every time tape data is captured, NetBackup IT Analytics recalculates the estimated tape capacity for the tape media type and volume pool combination. It takes into consideration the written KBs for the media type and volume pool for full tapes and then stores the average as the estimated tape capacity. Initially, when there is not enough data captured, this value may look low compared to the capacity stated by the vendor. Over time, the estimated capacity improves to the actual number of KBs that is being written to the tapes. Note that the amount of data that fits on the tapes differs based on the compression algorithm used and the type of data that is being backed up, which results in different compression ratios.

Chapter 25

Backup Service Level Agreement (SLA) Reports

This chapter includes the following topics:

- Available Service Level Agreement (SLA) Reports
- Configure a Backup SLA Policy
- Backup Start Time SLA
- Determining and Improving Backup Start Time Performance
- Backup Status SLA
- Changing the Success Percentage
- Backup Primary Server Overview
- Determining and Improving Backup Success Performance
- Backup Duration SLA
- Determining and Improving Backup Duration Performance
- SLA Job Summary

Available Service Level Agreement (SLA) Reports

Whether you're an Administrator in an IT department or an Administrator at a Data Center, you commit a specific level of service to your end-users or customers.

See "Configure a Backup SLA Policy" on page 885.

Use default SLA reports or create custom reports to provide information that determines if you are satisfying the commitments you made to your users or

customers. With access to the Reporting Database, you can configure the metrics upon which the reports are based. Service Level Agreements (SLAs) are typically based on the following metrics:

- Backup Start Time Performance
- Backup Status Performance
- Backup Duration Performance

The reports can provide data to fine-tune your SLA and to determine if you're meeting your commitments. The more detailed your SLAs, the easier it is to cut costs, in the case of an IT department, or provide competitive SLA pricing, in the case of Data Centers.

Refer to the following SLA reports:

- See " Backup Start Time SLA " on page 885.
- See "Backup Status SLA" on page 888.
- See " Backup Duration SLA " on page 890.

Configure a Backup SLA Policy

The Service Level Agreement (SLA) reports chart data against an expected successful backup goal, represented by a line in the bar chart. This goal can be modified to match your expectations.

Backup Start Time SLA

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

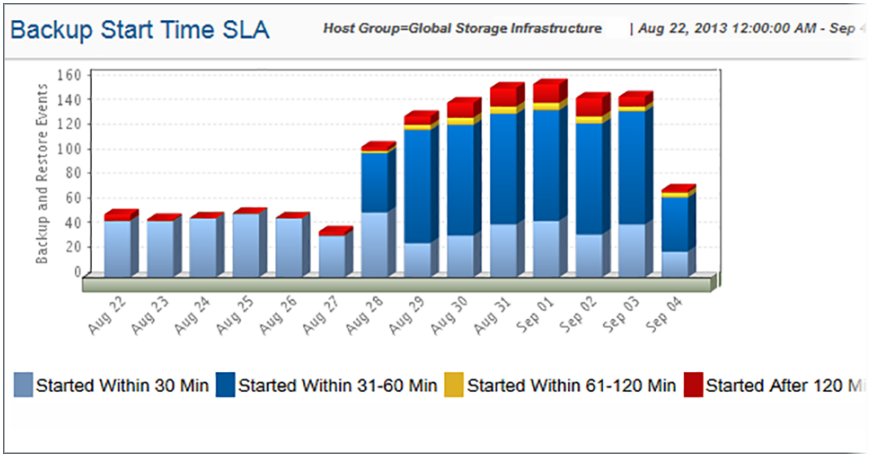
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > SLA Reports > Backup Start Time SLA

The Backup Start Time SLA report displays the total number of backups that began within a specific time range of their scheduled start time.

See "Available Service Level Agreement (SLA) Reports" on page 884.

This report is applicable to only NetBackup, Avamar, and TSM.



Use this report to determine what proportion of backups started at their scheduled time, or were delayed and by how long. If jobs are consistently missing their targeted start times, you may need to investigate the feasibility of your current schedule. A delay could be the result of other jobs taking too long or possibly because there are no drives available. To anticipate delays as a result of drive availability, consider using the

To determine and improve backup start time performance,

See " Determining and Improving Backup Start Time Performance " on page 886.

Determining and Improving Backup Start Time Performance

Use the Backup Start Time SLA report to gauge your backup efficiency. The Backup Start Time SLA report consists of a bar chart that provides a visual indicator of how long backups take to complete. Backup completion time is measured in days.

After you use this chart to identify the backups that represent poor performance, you can drill down to determine the causes so that you can prevent inefficiencies in the future. Your SLA determines what values represent poor performance vs. excellent performance. However, if you don't currently have an SLA implemented, strive for **Good** (31-60 minutes) performance by improving backup jobs with **Average** to **Poor** performance.

See "Available Service Level Agreement (SLA) Reports" on page 884.

To determine and improve backup start time performance

1. Go to the Backup Start Time SLA.

   See " Backup Start Time SLA " on page 885.

2. For each day, observe the number of backups in the bar graph's y-axis or hold your cursor over the color-coded segments to display the number of backups. Each segment represents 100% of the overall jobs for the day.

   Each bar represents a specific day. If no bar appears for a specific day, then either no backup jobs ran that day or no data is available because the Data Collector and the Reporting Database were not installed at that time.

3. Identify the poorly performing jobs and the reasons for the poor performance:

   Click on the color-coded bar segments that represent poor performance as defined by your SLA, drilling down into the actual jobs. As a general guideline, the Backup Start Time SLA report provides a few metrics:

   Metric

   Performance

   Started Within 30 minutes

   Excellent

   Started Within 31-60 minutes

   Good

   Started Within 61-120 minutes

   Average

   Started After 120 minutes

   Poor

   Identify any trends between backup jobs. For example, are the poorly performing backup jobs associated with a particular server?

4. Implement solutions to increase backup start time performance, such as:

   - Reduce the number of tape mounts.
   - Increase network speed.
   - Add additional processors or upgrade existing processors with faster processors.
   - Increase hard drive speed and size.

Backup Status SLA

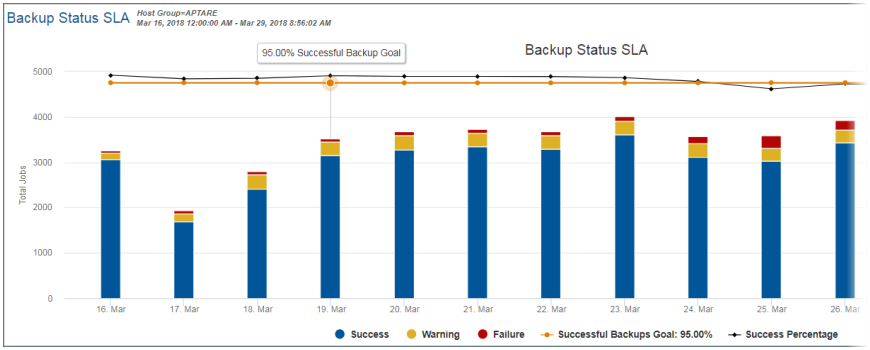
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > SLA Reports > Backup Status SLA

Use the Backup Status SLA report to determine backup success performance and to identify backup failures. A typical backup success percentage might be 85%, as indicated by the Successful Backups Goal line shown in the bar chart.

See “Available Service Level Agreement (SLA) Reports” on page 884.



Changing the Success Percentage

If you want to change the success percentage metric, refer to Defining Report Metrics in the *System Administrator Guide*. To change the percentage, you must have permission to write to the Reporting Database.

To determine and improve backup success performance, refer to the following.

See “ Determining and Improving Backup Success Performance ” on page 890.

See “Configure a Backup SLA Policy” on page 885.

Backup Primary Server Overview

For a view of relevant backup information, go to the Inventory and select **Backup Servers** or individual backup servers in the default hierarchy view (object type and

subsystem). Once in the Inventory, the Overview tab presents several reports in dashboard format.

Inventory > **Default Hierarchy View** > **Backup Servers** > **Overview** tab

The following reports are listed, prefaced by a Backup Primary Server Overview for the time period.

- See “Backup Status SLA” on page 888.
- See “Job Type Count” on page 790.
- See “ Backup Executive Summary ” on page 727.
- See “Anomalies Summary” on page 724.

Table 25-1 Backup Primary Server Overview

Column name	Description
#Primaries	The number of backup Primary Servers in the scope of this report.
Total Bytes	The total bytes backed up by the Primary Servers in the selected scope and time period. This value is represented in GiB.
# Files	The total number of files backed up to both disk and tape, for the selected scope and time period.
# Clients	The number of backup Clients in the scope of this report.
Disk Bytes	The number of bytes backed up to disk, for the selected scope and time period. This value is represented in GiB.
Disk # Files	The number of files backed up to disk, for the selected scope and time period.
Product(s)	A comma-separated list of the backup vendors and products in the selected scope.
Tape Bytes	The number of bytes of data backed up to tape, for the selected scope and time period. This value is represented in GiB.
Tape Files	The number of files backed up to tape, for the selected scope and time period.

Determining and Improving Backup Success Performance

The Backup Status SLA report provides you the ability to drill down into any failed job so that you can identify the cause of the failure. In addition to walking you through this identification process, this procedure provides a checklist to help you improve your backup success performance.

To determine and improve backup success performance

1. Navigate to the Backup Status SLA.

   See "Backup Status SLA" on page 888.

2. For each day, week, or month, hold your cursor over the triangle icon at the top of each bar to display the success rate. Compare this rate against your SLA.

   Each bar represents a specific day. If no bar appears for a specific day, then no backup jobs ran that day.

3. Use the chart to identify any jobs that failed.
4. Identify any trends between failed jobs. For example, are the failed jobs associated with a particular tape drive or disk drive?
5. See "Configure a Backup SLA Policy" on page 885.

Backup Duration SLA

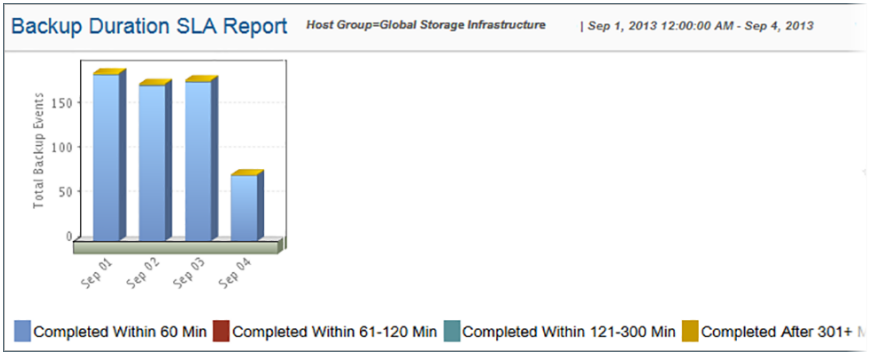
Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > SLA Reports > Backup Duration SLA

The Backup Duration SLA report displays the total number of backups that completed within the specific time range.

See "Available Service Level Agreement (SLA) Reports" on page 884.



As you are generating your Backup Duration SLA report, consider the following best practices:

- Use a consistent duration so that you can develop and compare data across the same interval.
- Create a duration report for each business unit or for each customer. Depending on how your company stores data, the Administrator can set up host groups for each business unit
- If you have different SLAs with different groups of users or customers, create a unique report for each SLA so that the data doesn't become skewed. Even if you have the same SLA for all users or customers, consider generating a unique report so that you can share backup duration performance data that is specific to their jobs.

To determine and improve backup duration performance, go to the following.

See " Determining and Improving Backup Duration Performance " on page 891.

Determining and Improving Backup Duration Performance

Use the Backup Duration SLA report to gauge your backup environment's efficiency. The Backup Duration SLA report consists of a bar chart. Backup completion time is measured in days.

Many backups run for a long time simply because they work with large amounts of data, not because of poor backup duration performance. However, sometimes a job with poor performance might be the result of a faulty tape drive that causes slow write speeds.

After you use this Backup Duration SLA chart to identify the backups that represent poor duration performance, you can drill down to determine if there is a problem.

Your service level agreement determines what values represent poor performance. However, if you don't currently have an SLA implemented, strive for **Good** (61-120 minutes) performance by improving backup jobs with **Average** to **Poor** performance.

The Backup Duration SLA report provides the following metrics:

Table 25-2 Table 2 Performance Metrics

Metric	Performance
Completed Within 60 minutes	Excellent
Completed Within 61-120 minutes	Good
Completed Within 121-300 minutes	Average
Completed After 301+ minutes	Poor

To determine and improve backup duration performance

1. Go to the Backup Duration SLA.

   See " Backup Duration SLA " on page 890.

2. For each day, observe the number of backups in the bar graph's y-axis or hold your cursor over the color-coded segments to display the number of backups.

   Each bar represents a specific day. If no bar appears for a specific day, then either no backup jobs ran that day or no data is available because the Data Collector and the Reporting Database were not installed at that time.

3. Identify poorly performing jobs:

   Click on the color-coded bar segments that represent poor performance, as defined by your SLA, drilling down into the actual jobs. Each segment represents a 100% of the overall jobs for the day.

   Identify any trends between backup jobs. For example, are the poorly performing backup jobs taking place at a particular time?

4. Implement solutions to increase backup performance, such as:

   - Back up during a less busy time.
   - Increase network speed.
   - Add additional processors or upgrade existing processors.

SLA Job Summary

This report can be accessed as a drilldown from the Backup Start Time SLA.

- See “ Backup Start Time SLA ” on page 885.
- See “ Backup Duration SLA ” on page 890.
- See “ Job Summary Report ” on page 750.

Table 25-3 SLA Job Summary

Column name	Description
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drill down to the Host Details report. For other backup sources, the object name is displayed with a drill down to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
Parent/Child	Displays Parent or Child for those jobs that have a hierarchical relationship. The field is blank if the job is not categorized as a parent or child job.
Server	The management server name.
Product	Vendor-specific backup software.
Type	The type of backup job. For a list of the types that are relevant for each backup product, click the **Advanced** button in the Job Summary Report Designer.
Start Date	Date and time of the beginning of the backup job
Finish Date	Click on this link to launch the Job Details report. For HP Data Protector jobs, multiple items will be reported, representing the HP Data Protector session, which is a collection of clients and their mount points. For these jobs, the finish date is derived from one of the jobs in a session.
Duration	Backup job duration
MBytes	Amount written to the backup media
MBytes/Sec	Backup throughput
Exit Code	The exit code of the backup job. For failed jobs, place your mouse on top of the red exit code to view the exit code details in a pop-up tool tip.

Table 25-3 SLA Job Summary *(continued)*

Column name	Description
# of Files	The number of files written to the backup media.
Tapes	Click on a tape link to launch a Tape Details Report.

Chapter 26

Backup Storage Utilization Reports

This chapter includes the following topics:

- Using Storage Utilization Reports
- Data Domain System Summary - Backup
- Data Domain System Details
- Data Domain Disk Summary
- Disk Usage and Performance Report
- Drive Performance Summary Report
- NetBackup Drive Performance Summary
- Drive Performance Detail
- Drive Utilization & Performance Report
- NetBackup Deduplication Effect
- NetBackup Deduplication Savings Trend
- NetBackup Storage Unit Summary
- NetBackup Storage Unit Usage
- Real Time Library and Drive Status
- Storage Unit Detail Report
- Storage Unit Images

- Storage Unit Event Details
- Disk Volume Details
- Tape Library and Drive Utilization
- IBM Spectrum Protect (TSM) Database Utilization
- IBM Spectrum Protect (TSM) Recovery Log Utilization

Using Storage Utilization Reports

To optimize storage utilization and performance, you can:

- Identify duplicate backups.
- Identify free capacity.
- Ensure maximum disk utilization throughput.
- Ensure that jobs start on time.
- Forecast future storage needs.

A number of practices contribute to storage consumption issues, such as over-provisioning storage and neglecting to deallocate resources. Your data protection strategy should include proactive procedures for reclaiming unused storage. The Portal provides reports to assist you in optimizing storage utilization and performance.

Data Domain System Summary - Backup

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Data Domain System Summary - Backup

View a summary of the Data Domain systems.

Use this report to monitor Data Domain system capacities to determine the amount available for data storage. The capacity and used values are derived from Data Domain file system data, while virtual tape library (VTL) values represent the totals for all VTLs and tapes in the system.

Table 26-1 Data Domain System Summary - Backup

Column name	Description
System	The Data Domain system.
OS Version	The operating system of the host. Valid values: Windows, Linux, Unknown.
Last Updated	The date and time the host was last updated in the NetBackup IT Analytics database.
Filesystem	
Active Capacity	Total capacity for the host's filesystems.
Used Active	The percentage of the file system capacity that is used.
Archive Capacity	Archive total capacity, in kilobytes, for the host's filesystems. This is the sum of the archive storage capacity for all filesystems for the host.
Used Archive	The percentage of the archive capacity for the host's filesystems that is used.
# of MTrees	Number of MTrees associated with this system. Click the link to view the Data Domain MTree Summary.
# of Disks	Number of physical disks. See " Data Domain Disk Summary " on page 900.
Raw Capacity	The data storage capacity of the disk when used in a Data Domain system. When calculating disk space, Data Domain defines one gigabyte as 230 bytes, which results in a disk capacity that is different from the disk manufacturer's rating.
Virtual Tape Library	
# VTLs	Number of VTLs.
Capacity	The capacity for all VTLs in this system.
Used	The amount of the VTL capacity that is in use.
Storage System ID	Storage system ID, for internal use by the Portal database.
# Tapes in VTLs	Number of VTL tapes.

Data Domain System Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Storage Utilization Reports > Data Domain System Summary - Backup

This full set of Data Domain System Details is available as a drill down from other Data Domain reports that contain a System link.

Table 26-2 Data Domain System Details

Column name	Description
System	
Internal System ID	Data Domain's internal ID links to the Host Detail report.
System	The Data Domain system name.
Model	The Data Domain system model.
OS Version	The Data Domain operating system version running on this system.
Serial #	Serial number of the Data Domain system.
Filesys Encrypt	Indicates if file system encryption is enabled.
Last Poll Date	Last time the Data Domain system was updated.
# of Disks	Number of physical disks. Click to drill down to Data Domain Disk Summary report.
Capacity	
Raw Capacity	Total of all disk capacities.
Active Capacity	Amount of capacity that is RAID-configured (used + free).

Table 26-2 Data Domain System Details *(continued)*

Column name	Description
Used Capacity	Derived from the file system information.
Usage	Mouse over this thermometer to view the active capacity usage %.
Archive Capacity	Amount of archived file system capacity.
Used Archive	Amount of the archived file system capacity that is in use.
Usage	Mouse over the thermometer to view the archive capacity usage %
Active Metadata capacity	Total allocated storage capacity of the active tier in GiB.
Used Metadata capacity	Amount of storage that is currently occupied by data and metadata on the active tier in GiB.
Usage	Mouse over the thermometer to view the active metadata usage in percentage.
Archive Metadata capacity	Total allocated storage capacity for archive metadata within the system in GiB.
Used Metadata capacity	Amount of storage currently occupied by archive metadata in GiB.
Usage	Mouse over the thermometer to view the archive metadata usage in percentage.
LSU Summary	
# LSUs	Number of configured logical storage units in this Data Domain system; links to the Storage Unit Detail.
# Files	Total number of files in the LSUs.
Compression Factor	De-duplication ratio (original bytes divided by locally compressed).
Original Bytes	Amount of original data before compression and de-duplication.
Globally Compressed	Indicates de-duplication efficiency.
Locally Compressed	File compression.
Metadata	Amount of LSU storage consumed by metadata; this is only collected if SNMP is configured as your data collection method.
MTree Summary	
# MTrees	The number of MTrees associated with this Data Domain system. This number links to the Data Domain MTree Summary report.

Table 26-2 Data Domain System Details *(continued)*

Column name	Description
Last 24 hrs Pre Compression	Data written before compression within the last 24 hours. Represents the space used and compression achieved for files and directories in the file system.
Last 24 hrs Post Compression	Storage used after compression within the last 24 hours.
Last 24 hrs Compression Factor	The factor used by Data Domain for compression.
VTL Summary	
# VTLs	Number of VTLs available to this Data Domain system.
# Tapes	Number of tapes in the VTLs.
Total Size	Total storage capacity in the VTLs.
Space Used	Amount of VTLs in use.
% Free	% of VTL storage still available.
Usage	Mouse over this thermometer to view the usage %.
Average Comp	Average compression factor.
Overall Compression Factor	Overall compression factor is derived from pre-compression size (data written before compression) divided by file system used.

Data Domain Disk Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Data Domain System Summary - Backup

1. Click a **System** link to access **Data Domain System Details.**
2. Click **# Disks** on the **Data Domain System Details** report.

Table 26-3 Data Domain Disk Summary

Column name	Description
System	The Data Domain system links to details. See " Data Domain System Details" on page 246.
Disk	Disk number in this system.
Manufacturer/Model	The disk's manufacturer and model number.
Firmware	Disk's firmware version identifier.
Serial #	Disk's serial number.
Capacity	Maximum capacity of the disk.
Read Sects/s	Read throughput averaged over time since the disk's last performance statistics reset or since the last system power-up cycle.
Write Sects/s	Write throughput averaged over time since the disk's last performance statistics reset or since the last system power-up cycle.
Cumulative MB/s	Cumulative throughput averaged over time since the disk's last performance statistics reset or since the last system power-up cycle.
Busy	% of time the disk was busy, based on the last time this data was collected.

Disk Usage and Performance Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Disk Usage & Performance

Note: This report is a NetBackup-only report.

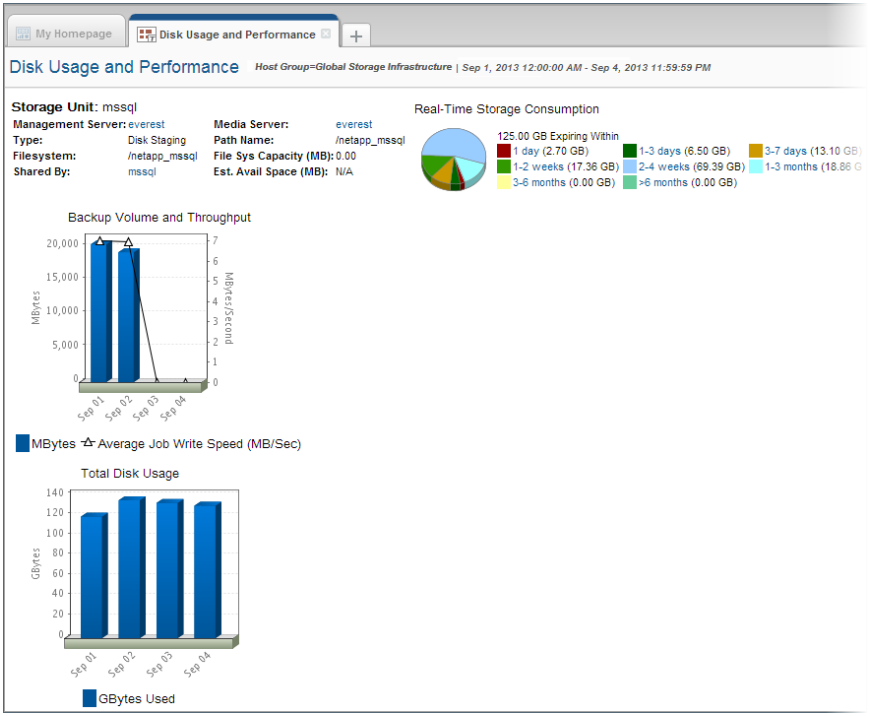


Table 26-4 Disk Usage and Performance Report

Column name	Description
Storage Unit Details	The name of the storage unit.
Mgmt Server Details	Click on the Management Server name link to go to the corresponding Host Details Report.
Media Server Details	Click on the Media Server name link to go to the corresponding Host Details page.

Table 26-4 Disk Usage and Performance Report *(continued)*

Column name	Description
Storage Unit Filesystem Info	Each storage unit has an associated path name that determines where the backups are stored. The report shows both the path name and the actual filesystem on which the storage unit resides. If more than one storage unit shares the same filesystem, the report lists the names of the other storage units. It's important to know if other storage units share the same filesystem, since any physical free space on the filesystem will be shared as one common pool by all the storage units for that filesystem. The Filesystem Capacity (MB) field is populated only if Discovery has been enabled.
Real-Time Storage Consumption	Shows the amount of data on the Storage Unit for the selected scope grouped by expiration period. Mouse-over the pie slices to view the total number of backups associated with the amount of allocated storage. Click a pie slice to go to the Storage Unit Detail which displays a list of all the active backups that fall within the expiration period and are within the report scope. See " Storage Unit Detail Report " on page 914.
Backup Volume & Throughput	Each bar in this graph represents the total number of GBytes written to the Storage Unit for backups jobs that completed during the period represented by the bar. Click a bar to display a list of all the backups represented by the bar. See " Storage Unit Detail Report " on page 914. The dotted line plotted against the secondary y-axis represents the Average Throughput of the backups for the period represented by each bar. The calculation only includes those jobs within the current scope. The throughput is calculated as follows: backup volume/(sum of all backup durations in seconds)
Total Disk Usage	Each bar in this graph represents the total amount of space in GBytes occupied by unexpired backups.

Drive Performance Summary Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Drive Performance Summary

This report also can be generated via a link in the following report.

See “Drive Utilization & Performance Report ” on page 908.

This report provides high-level drive details for the reporting period that you specified in the scope.

Note: If you are accessing this report from the navigation pane, be sure to define a host group scope that is unique by product type--for example, only NetBackup servers. Also, this report is not supported for Veritas Backup Exec or IBM Spectrum Protect (TSM) data because data collection does not capture the association between backup jobs and tapes.

Table 26-5 Drive Performance Summary Report

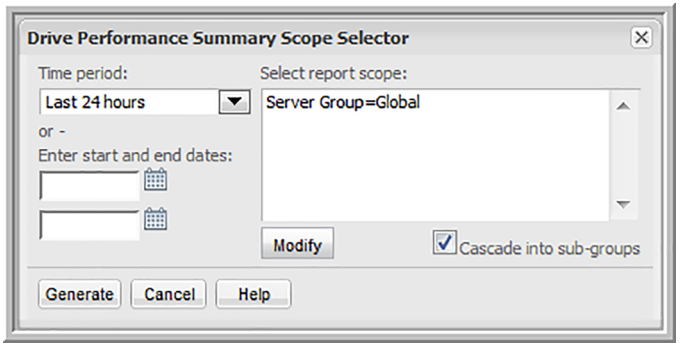
Column name	Description
Drive Name	Links to the details. See “Drive Performance Detail” on page 907.
Drive #	Drive index.
Drive Type	The drive type label.
Drive Status	The current status of the drive.
Library Name	The name of the library.
Controlling Server	Links to the controlling server’s details in the Host Details report. See “ Host Details ” on page 838.
Total MBytes	Total MBytes used. A total for all drives is shown at the bottom of this column.
Number of Jobs	The number of backups that used this drive.

Table 26-5 Drive Performance Summary Report *(continued)*

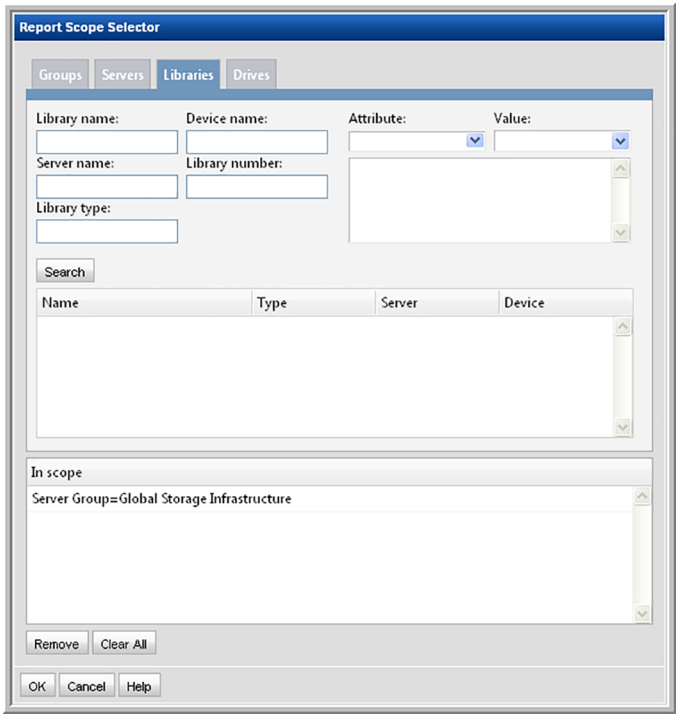
Column name	Description
Min MBytes/Sec Max MBytes/Sec Avg MBytes/Sec	These three values provide indicators of the drive's performance. The Maximum rate represents the fastest speed that a job ran during the selected time frame. Likewise, the Minimum rate represents the slowest speed that a job ran during the same time frame. See " NetBackup Drive Performance Summary " on page 906.

As an alternative to Host Group scope, you can select Tape Libraries as the scope:

1. In the Drive Performance Summary Scope Selector window, click **Modify** to change the scope.



2. In the Client Selection window, first click **Clear All** to remove the host groups from the "In scope" pane.



3. Click the **Tape Library** tab to display the list of possible tape libraries in your environment.
4. Double-click the tape libraries to include them in the scope and then click **OK**.

NetBackup Drive Performance Summary

For NetBackup jobs, the drive performance algorithm takes into account the use of multiple tapes, as some backups require the dismounting and re-mounting of several tapes. The Kilobytes/second rate is calculated for each piece of tape media used by the backup job.

```
mediaKilobytes / ((finishReadWrite - startReadWrite) * SECS_IN_A_DAY
 * (mediaKilobytes/jobKilobytes))
```

Where:

- mediaKilobytes is the total kilobytes written to an individual piece of tape media
- jobKilobytes is the total size of the backup job, as reported by bpdbjobs
- startReadWrite is the date/time that the job started writing to the tape

- finishReadWrite is the date/time that the job finished writing to the tape

The Drive Performance Summary report displays the minimum, average, and maximum overall drive performance for all jobs for the reporting period. When you drill down to an individual tape drive, the Drive Performance Detail Report uses this same algorithm to display the performance of each job. If there are multiple rows for a single job, such as when multiple tapes are used by the job, the performance metric is averaged across all the rows returned from the above algorithm.

The details from which drive performance is derived can be viewed via drill down links. In the Drive Performance Detail report, click the Finish Date for one of the backup jobs to view the Job Details. At the bottom of the report the backup attempts (try records) are displayed.

See "Drive Performance Detail" on page 907.

See " Job Details Report " on page 770.

Drive Performance Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Storage Utilization Reports > Drive Performance Summary

Click a Drive Name.

Table 26-6 Drive Performance Detail

Column name	Description
Drive Name	The drive name--this was the link to this report from the Drive Performance Summary report.
Library Name	The library name links to the Tape Media Summary report. See " Tape Media Summary Report " on page 875.

Table 26-6 Drive Performance Detail *(continued)*

Column name	Description
Client	Links to the Host Detail for the client. See " Host Details " on page 838.
Start Time	When the backup job started.
Finish Time	Links to the Job Details report. See " Job Details Report " on page 770.
Duration	Duration of the job.
Status	Status of the job.
# of Files	The number of files backed up to this drive.
MBytes	Amount of data backed up to this drive.
MBytes/Sec	For NetBackup, this value is calculated using the job write start time and job write end time (not the job start time and job finish time). For example, a job can start, but then it can wait in a queue before the job begins to write to a tape/disk. That queue time is eliminated in this calculation.

For NetBackup specifics, refer to the following.

See " NetBackup Drive Performance Summary " on page 906.

Drive Utilization & Performance Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Drive Utilization & Performance

The Drive Utilization and Performance report enables you to view and compare utilization and performance of the tape drives in your environment. By viewing the total aggregate throughput per drive, you can use this information to help you tune your backup policies.

This report is applicable to only Veritas NetBackup, HP Data Protector, Commvault Simpana, IBM Spectrum Protect (TSM).

Avg KBytes/Sec and **Max KBytes/Sec** is displayed for every bar during which an event finish time took place. These calculations are available only for periods of time when an event finished. So, if 20 events finish in one of the bar periods, the calculation is based on the average write speed for the total time for all 20 jobs / total bytes written. Likewise, using this example, Max speed would be the maximum of the averages of all 20 jobs.

The Portal does not tie into the hardware directly and monitor the actual drive write speeds. Instead, the reports use the data that the Data Collector retrieved from the backup software--including the amount of data written, job duration time, and job completion time--for each job.

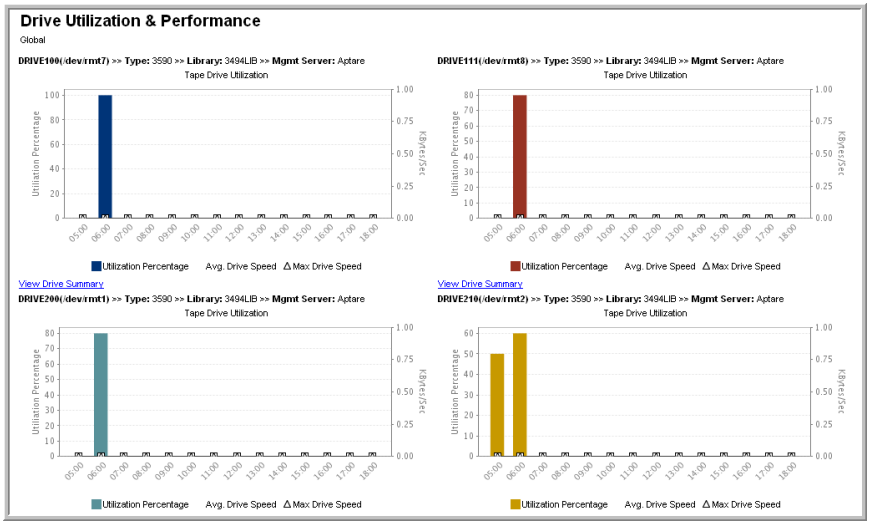


Table 26-7 Drive Utilization & Performance Report

Column name	Description
Drive Details	For each drive, listed at the top of each Drive Utilization and Performance Chart: ■ Drive Name ■ Drive Type ■ Tape Library Name ■ Backup Management Server Click on the Backup Management Server name to view its Host Details page.

Table 26-7 Drive Utilization & Performance Report *(continued)*

Column name	Description
Pop-up Details	Mouse-over any chart bar or line symbol to view the corresponding values.
Drive Summary	Click on the list at the bottom of each chart to go to the report, which lists all the backup events that took place on that drive for the time period that you specified in the scope. See “ Drive Performance Summary Report ” on page 904. Additionally, the min, max, and avg KBytes/Sec for the time period are provided at the bottom of the report.
Avg & Max KBytes/Sec	The Avg KBytes/Sec and Max KBytes/Sec is displayed for every bar during which an event completed. If the report shows a bar with utilization but does not have accompanying line symbols, either no backups or restores completed during that time period or the drive was being utilized for another purpose, such as copying a tape to be vaulted.
Local Start and End Times	Each chart displays the utilization and performance of one drive in the system. The start and end time for each chart is localized to the time zone where that drive’s management server is located. For example, if Drive A’s management server is located in the Pacific Standard Time Zone and Drive B’s management server is located in the Eastern Standard Time Zone (3 hours ahead of Drive A’s management server) and you ran the Drive Utilization and Performance Report for the last 12 hours, the Drive A chart’s start and end times would be 3 hours earlier than the chart for Drive B. This start/end time localization enables real time performance comparison between drives, even if they are located in different parts of the world.

NetBackup Deduplication Effect

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > NetBackup Deduplication Effect

This report enables you to see the benefits of deduplication.

NetBackup Deduplication Savings Trend

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > NetBackup Deduplication Savings Trend

Use this report to quickly validate how much storage has been saved because of deduplication. The graph displays a graphical representation of the deduplication savings.

NetBackup Storage Unit Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > NetBackup Storage Unit Summary

This report is NetBackup-specific.

Note: Only the first four tape libraries are displayed.

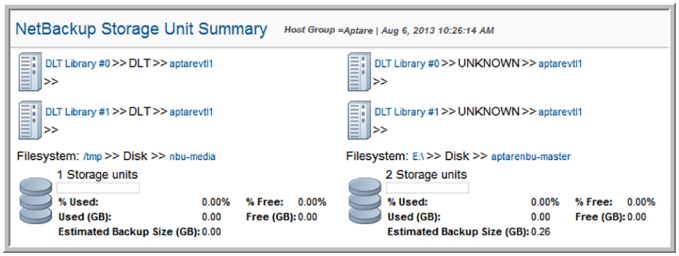


Table 26-8 NetBackup Storage Unit Summary

Column name	Description
Library Details	Click on the library name or server to view the utilization and performance. See "Drive Utilization & Performance Report " on page 908.
Filesystem Details	Click on the server name to view the details. See " Host Details " on page 838.

NetBackup Storage Unit Usage

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > NetBackup Storage Unit Usage

Use this report to view the backup storage unit utilization by primary and media server for Veritas NetBackup and Veritas NetBackup Appliances for a selected time frame.

Table 26-9 NetBackup Storage Unit Usage

Column name	Description
Storage Unit	Storage unit name as assigned either by Veritas NetBackup or the user.
Storage Unit Type	Lists the disk volume disk type such as PureDisk, AdvancedDisk etc.

Table 26-9 NetBackup Storage Unit Usage *(continued)*

Column name	Description
Primary Server	Primary host name to display in Portal reports.
Media Server Name	Media Server name to display in Portal reports.
Total Job Attempt Count	The number of attempts performed by jobs before succeeding.
Job Duration	Duration of all the jobs for the particular storage unit.
Backup Image Size	The size of the backup image.
MBytes/Sec	Backup throughput
Utilization % for a Duration	Utilization percentage of the storage unit for a selected duration.
Snapshot Configured	Yes indicates the disk pool is configured to perform snapshot backups.

Real Time Library and Drive Status

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Real Time Library and Drive Status

Use this report to monitor the health of each library. This report displays current high-level information for each tape library in your environment. The data will only be as current as the time of the most recent poll of your management server(s).

This report is applicable to only NetBackup and NetWorker.

Table 26-10 Real Time Library and Drive Status

Column name	Description
Library	Library name.
Type	The type of library, such as Optical or HCART.
Management Server	Click this link to view Host Details.
Slots	The number of slots in the library.
Drives	The number of drives in the library.
In Use	Number of drives currently in use.
Drives Available	Drives - In Use = Drives Available
Down	The number of disabled drives. When this number becomes equal to or greater than 25% of the library's drives, the library's row of information will be highlighted in red.
Avail %	The percentage of drives in this library that are available.

Storage Unit Detail Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Storage Unit Detail

This report is applicable to only NetBackup.

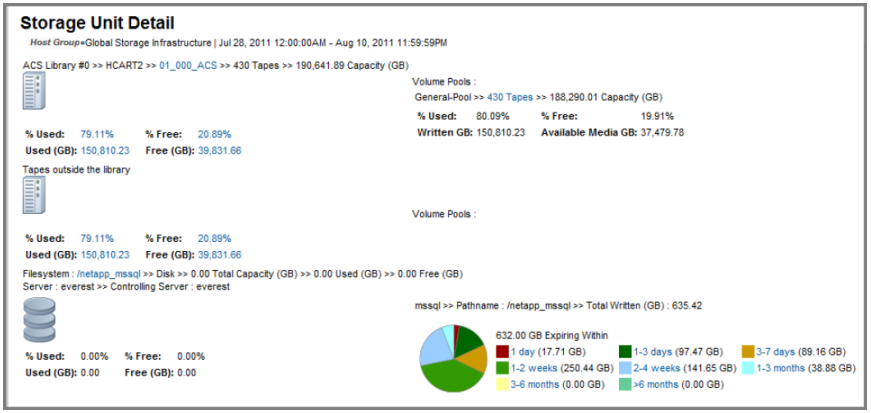


Figure 22.1 Example of the Storage Unit Detail Report

Table 26-11 Storage Unit Detail Report

Column name	Description
Tape Library	Click on a Library name link or the icon to go to the summary for all media inside the library. In addition, various usage links launch the Tape Media Summary. See “ Tape Media Summary Report ” on page 875. Each library is accompanied by a media alert thermometer whose color is an indicator of the percentage of GByte availability for all media in the library: ■ Green: More than 50% available ■ Yellow: Less than 50% available ■ Red: Less than 25% available
Volume Group	Displays each library or group of tapes outside the library, and each has its own volume group. The media availability for the volume group is represented by a color coded thermometer so it is easy to tell whether there is a healthy amount of media GBytes available for future backups.
Scratch Pool Media Usage	If there is a designated scratch pool for a Tape Library or Group of Tapes Outside the Library, the report display the scratch pool as a feeder pool to the Volume Pools. Each scratch pool is accompanied by a color-coded scratch pool thermometer that indicates the amount of scratch media available. The value represented is a percentage: the capacity of the scratch media compared to the capacity all media in the library.

Table 26-11 Storage Unit Detail Report *(continued)*

Column name	Description
Volume Pools	When there is a designated scratch pool, this displays all volume pools in that library that are fed by the scratch pool. Otherwise, the general pool is shown. Click the number of tapes to view the summary. See " Tape Media Summary Report " on page 875.
Volume Pool Usage	Displays one or more volume pools, depending on how many exist in the library. Each volume pool has a pie chart that shows the % GBytes Free compared to the % GBytes Used. Click on any part of the pie chart or the Free and Used values to go to the summary for that tape library. See " Tape Media Summary Report " on page 875.
Filesystem	Displays details about the disk space of the media server's filesystem on which the disk storage units reside. The operating system reports this data. Values of **Unknown** indicate that the Discovery module is not configured to probe the media servers' filesystems for their physical characteristics. The management server and controlling server also are shown with the filesystem details.
Storage Unit Usage	Shows the active backups on the Storage Unit grouped by expiration period. Click a pie slice to go to the details which displays a list of all the active backups within the particular expiration period. See " Storage Unit Event Details " on page 918.
NetBackup Storage Units	
Storage Unit	Indicates the NetBackup disk storage unit; links to the Storage Unit Images report.
Disk Pool	Identifies the disk pool.
Storage Server	The storage server for this storage unit.
Status	Indicates the state of the disk pool, for example: Up.
Flags	A list of the various flags that represent disk pool attributes.
High Water Mark	This threshold triggers actions associated with a full volume: failed backup jobs, new jobs will not be assigned to a storage unit since the disk pool is full, and image cleanup is initiated. Default is 98%.
Low Water Mark	This value has no effect for PureDisk volumes.
Num Read Mounts	The number of times the volume was mounted for reads.

Table 26-11 Storage Unit Detail Report *(continued)*

Column name	Description
Num Write Mounts	The number of times the volume was mounted for writes.
Current Read Streams	Indicates the current number of read streams.
Current Write Streams	Indicates the current number of write streams.
Num Repl Sources	Indicates the number of replication source deduplication pools.
Num Repl Targets	Indicates the number of targets for the duplication operations.
Mounts	Lists the filesystem mounts.

Storage Unit Images

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Storage Utilization Reports > Storage Unit Images

This report is a drilldown from the Storage Unit Details report.

This report provides an image list for the disk storage units that are connected to the Media Server. Use this report to identify possible data loss when disks have gone bad or a media server crashes.

Table 26-12 Storage Unit Images

Column name	Description
Backup ID	Links to the relevant Job Details report.
Client	The client that is being backed up; links to the Host Details report.

Table 26-12 Storage Unit Images *(continued)*

Column name	Description
Image Created	Creation date and time of the storage unit image.
Expires	Date and time when the storage unit image expires.
Size (MB)	Size of the storage unit image.
Schedule	The name of the NetBackup schedule that drives the frequency of the backups.
Type	Type of image, such as Full Backup or Differential-Incremental Backup.
Policy	The backup policy associated with this storage unit image.

Storage Unit Event Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Storage Utilization Reports > Storage Unit Detail

In the file system section of the report, click a pie slice or pie chart legend link.

Table 26-13 Storage Unit Event Details

Column name	Description
Backup ID	Links to the Job Details.
Client	Links to the Host Details.
Start Time	When the backup job started.
Finish Time	Links to the Job Details.
Expiration	Date and time of the tape media expiration.

Table 26-13 Storage Unit Event Details *(continued)*

Column name	Description
MBytes	Amount of data written to the Storage Unit.
# of Files	Number of files written.
Exit Code	Status of the backup job.

Disk Volume Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Storage Utilization Reports > Storage Unit Detail

Click the numeric **Disk Volumes** link.

This NetBackup-specific report provides additional details about the volumes associated with a disk pool.

Table 26-14 Disk Volume Details

Column name	Description
Disk Pool	Identifies the disk pool.
Disk Volume	Indicates the disk volume included in the pool.
Status	Indicates the state of the disk pool, for example: Up.
Flags	A list of the various flags that represent disk pool attributes.
Total Capacity (GB)	Total raw size for the disk pool.
Free (GB)	Amount of usable disk pool space.

Table 26-14 Disk Volume Details *(continued)*

Column name	Description
Num Read Mounts	The number of times the volume was mounted for reads.
Num Write Mounts	The number of times the volume was mounted for writes.
Current Read Streams	Indicates the current number of read streams.
Current Write Streams	Indicates the current number of write streams.
Num Repl Sources	Indicates the number of replication source deduplication pools.
Num Repl Targets	Indicates the number of targets for the duplication operations.

Tape Library and Drive Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > Tape Library & Drive Utilization

The Tape Library and Drive Utilization report displays the average drive utilization against the actual drive utilization for any given time period. Often, this report is used to justify VTL purchases. Identify how many tape drives are writing and for how long. In addition, the Drive Utilization and Performance report provides insight into the throughput, with drilldowns to drive performance.

This report does not display data from HP Data Protector or Veritas Backup Exec.

See “Drive Utilization & Performance Report ” on page 908.

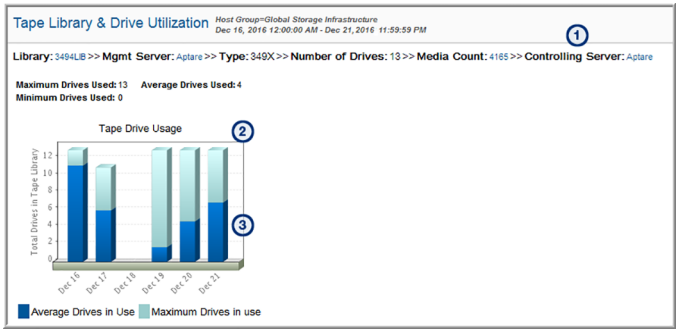


Table 26-15 Tape Library and Drive Utilization

Column name	Description
1. Library Utilization	At the top of each chart, a summary of the library utilization provides a context for the bar chart.
2. Bar Chart	Each bar shows the average amount of time the library drive(s) are being utilized stacked against the actual utilization time.
3. Tooltip Details	Mouse-over any of the bars on the chart to launch pop-ups of details.

IBM Spectrum Protect (TSM) Database Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > IBM Spectrum Protect (TSM) Database Utilization

The IBM Spectrum Protect (TSM) Database Utilization report provides usage and capacity trending of the database. Using this report, you can determine how much of the allocated database space has been used over the selected report period.

Table 26-16 IBM Spectrum Protect (TSM) Database Utilization

Column name	Description
Drillable Stacked Bar Chart	Each bar shows the average amount of time the library drive(s) are being utilized stacked against the actual amount of utilization time. Click on any bar section will to view the Backup Events Report.
Pop-up Details	Mouse-over any of the bars on the chart to display pop-ups including the details of that particular bar section.
Usage & Availability Trending	Shows how much of the allocated database space has been utilized within the TSM database over the time period that you specified in the scope.

IBM Spectrum Protect (TSM) Recovery Log Utilization

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Storage Utilization Reports > IBM Spectrum Protect (TSM) Recovery Log Utilization

This report displays Recovery Log usage--both the maximum used and the available capacity. The Recovery Log contains transactions prior to TSM committing them to the database. Use this report to monitor the size of the log to prevent degradation on the TSM server.

The data can be displayed either as a bar chart or a tabular report. To select a report type, click **Advanced** in the scope selector.

Chapter 27

Backup Manager Forecasting Reports

This chapter includes the following topics:

- Forecasting Overview
- NetBackup Library Capacity Forecast
- Media Availability Forecast
- Media Consumption Forecast
- Media Usage Details
- Scratch Pool Forecast Report
- Tape Drive Usage and Forecast Report

Forecasting Overview

To ensure accurate forecasts, the reports make the following decisions:

- To provide a good statistical sample, when you select a display range, the report uses historical data that is twice that display range.
- If the full range of historical data is unavailable, the data that is available is used. For example, if you want to view data for January of the previous year to December of the current year, the Portal looks for four years worth of historical data.
- Although the report forecasts the current period, the report ignores the current period when it derives statistical calculations. For example, if the current month is July and you request to view the report by months, the bar for July of this year

will be forecasted and data for July of this year will be ignored for statistical calculations.

NetBackup Library Capacity Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

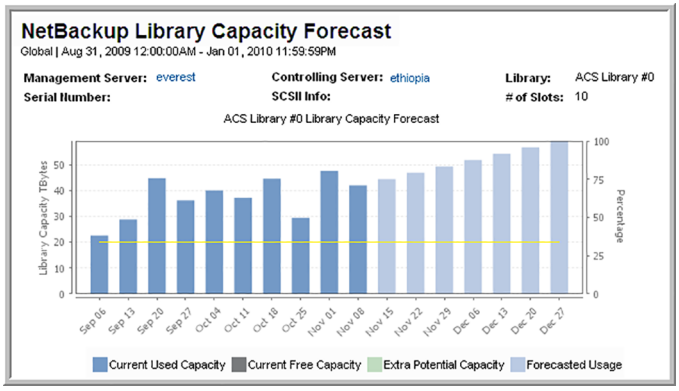
Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Forecasting & Capacity Planning > Library Capacity Forecast

This report can display either a bar chart view or a tabular report of library capacity, current usage, and forecasted usage. In the scope selector, select display as Chart or display as Table.

Note that in the bar chart, the yellow line represents estimated library capacity.

See " Tape Media Detail Report " on page 877.



- Current Used Capacity - Space occupied on full and filling tapes now in slots
- Current Free Capacity - Space remaining on empty or filling tapes now in slots
- Extra Potential Capacity - Average measured capacity of one full tape times the number of open slots
- Forecasted Usage - The anticipated tape usage based on historical data

To view this same data in tabular format, select Display as Table in the scope selector.

The following table lists fields displayed in either the bar chart or the table.

Table 27-1 NetBackup Library Capacity Forecast

Column name	Description
Management Server	This is the NetBackup Primary Server.
Controlling Server	This is the NetBackup media server that controls the library. This links to Host Details.
Library	The name of the tape library
# of Slots	The number of slots in the library
Serial Number	Serial number of the tape library
SCSII Info	The device's SCII details
Date	The date that represents the time period start date
Current Usage	Average usage, as of the associated date
Forecasted Usage	Forecasted usage, based on the reports start and end dates

Media Availability Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Forecasting & Capacity Planning > Media Availability Forecast

This report is applicable to only NetBackup.

The Media Availability Forecast plots forecasted consumption against forecasted availability.

The system calculates the forecasted consumption by the close of each day, using the following algorithm:

(previous day's availability + forecasted media expiration - forecasted consumption)

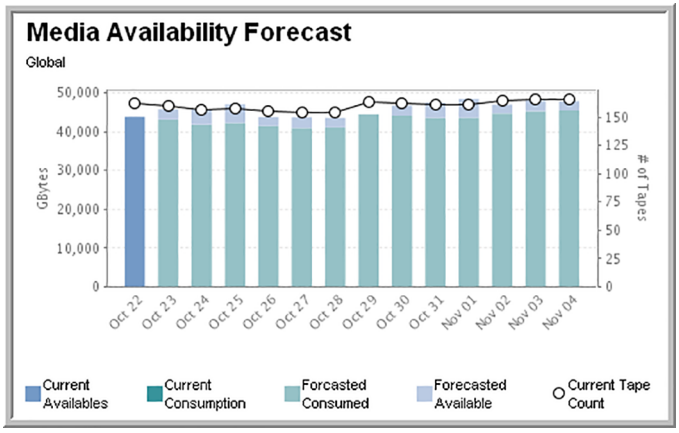
In order to forecast the 7 days of usage, the system averages day-of-the-week consumption data for a historical period specified in the “Historical Data Period” parameter on the Media Forecasting Dashboard scope selector.

When forecasting the number of tapes consumed on each day, the system tallies the GByte consumption for all backups written to tape per media type in each volume group. The system then divides the total GByte consumption for that volume group / media type by the estimated capacity of that media type in that volume group. This process is repeated for each volume group / media type instance within the scope of the report and the aggregated tape consumption results are displayed via the circular line symbols in the bar chart.

The circular line symbol tracks the # of tapes available for each day displayed on the chart.

Place your mouse over each colored portion of a bar to see the following details for the associated time period:

- Currently Available Storage Media
- Current Consumption
- Forecasted Consumption
- Forecasted Available
- Forecasted Tape Availability (mouse over connected circles at the top of the bars)



Media Consumption Forecast

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Forecasting & Capacity Planning >Media Consumption Forecast

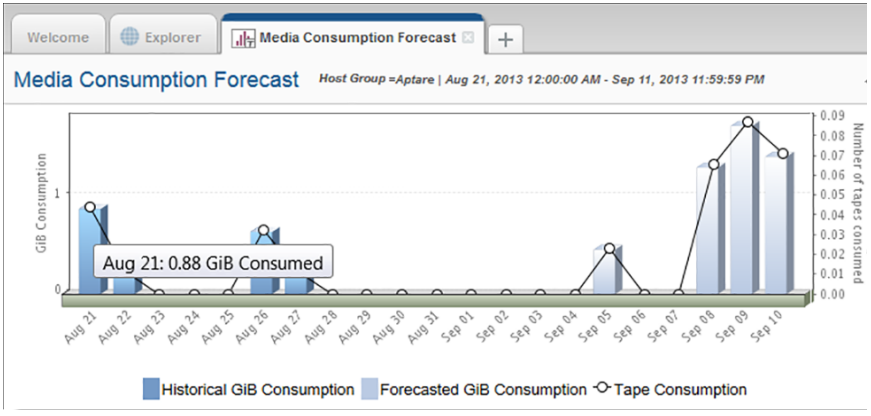
This report is applicable to only NetBackup.

A bar chart represents historical tape usage, with a projection of future needs. You can select report periods ranging from one week to eight weeks of data. In addition, you can narrow the report scope by choosing specific host groups or even clients within a host group.

In order to forecast the seven days of usage, the report averages day-of-the-week consumption data for a historical period specified in the Time Period.

To forecast the number of tapes consumed on each day, the report tallies the GByte consumption for all backups written to tape per media type in each volume group. Then the total GByte consumption for that volume group/media type is divided by the estimated capacity of that media type in that volume group. This process is repeated for each volume group/media type instance within the scope of the report and the aggregated tape consumption results are displayed via the circular line symbols in the bar chart.

Note: The current day is treated as the future for the purpose of forecasting.



Media Usage Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Forecasting & Capacity Planning > Media Consumption Forecast

Click a Media Consumed bar in the bar chart.

Table 27-2 Media Usage Details

Column name	Description
Client	Click this link to view the details. See “ Host Details ” on page 838.
Backup ID	Click this link to launch the details. See “ Job Details Report ” on page 770.
Mgmt Server	The server that is executing the backup policy. Click this link to view the details. See “ Host Details ” on page 838.
Volume Pool	The name of the volume pool to which the media belongs.
Finish Time	Click on this link to launch the details. See “ Job Details Report ” on page 770.
Weekday	The day of the week on which the backup occurred.
Total MBytes	Total amount written to the media.
Total Files	Total number of files written to the media.
Estimated Tapes	Forecasted tape need for the backup job.

Table 27-2 Media Usage Details *(continued)*

Column name	Description
Tape ID	Click on these links to launch the details. See " Tape Media Detail Report " on page 877.

Scratch Pool Forecast Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Forecasting & Capacity Planning > Scratch Pool Forecast

This report is applicable to only NetBackup, Networker, TSM, and Commvault Simpana.

A scratch pool is a repository of media that is available to backup servers. There are a few things that affect the availability of media in a scratch pool:

- If your backup servers run out of media, they automatically take and use media from the scratch pool.
- If your backup servers delete data, making media available, they automatically add the recycled media to the scratch pool.
- When you clean tapes or expire tapes, you temporarily or permanently affect the pool of media.

Note: The current day is treated as the future for the purpose of forecasting.

Because backup success is tied directly to scratch pools, it's important that you monitor and forecast media usage in the scratch pools.

The Scratch Tape Usage and Forecast Report provides the following information:

Table 27-3 Scratch Pool Forecast Report

Column name	Description
# of Tapes drives	The number of tape drives in all tape libraries for the report scope that you specified.

Table 27-3 Scratch Pool Forecast Report *(continued)*

Column name	Description
Usage/Forecast	Mouse-over each bar to retrieve the following information: ■ Average scratch tapes used previously ■ Maximum scratch tape used previously ■ Average scratch tapes forecasted ■ Maximum scratch tapes forecasted

Tape Drive Usage and Forecast Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Forecasting & Capacity Planning > Tape Drive Usage & Forecast

This report is applicable to only NetBackup, Networker, TSM, and Commvault Simpana.

Use this report to determine if and when you need to add more tape drives to your libraries. When you customize this report, if you enter a date range in the past, the report will include historical usage information. If your date range includes a finish date beyond the current day, the report will include forecast estimates.

Note: The current day is treated as the future for the purpose of forecasting.

Note: Tip: When entering the date range in the scope selector, use a finish date that is a few weeks or months into the future. This will ensure that a sufficient forecast is produced. Also, to produce a realistic forecast, group the data by Weeks, Months, or Quarters. For this type of report, the time units of Hours and Days are too small to provide meaningful charts.

The report derives the forecast by comparing usage with your existing tape inventory. To be safe, use the highest forecast when determining how many tape drives you need to add.

The Tape Drive Usage and Forecast Report provides the following information:

Table 27-4 Tape Drive Usage and Forecast Report

Column name	Description
# of Tapes drives	The number of tape drives in all tape libraries for the report scope that you specified.
Usage/Forecast	Mouse-over each color of a bar to retrieve the following information: ■ Average tape drives used previously ■ Maximum tape drives used previously ■ Average tape drives forecasted ■ Maximum tape drives forecasted

Chapter 28

Backup Billing and Usage Reports

This chapter includes the following topics:

- Available Billing and Usage Reports
- Determining Media Chargebacks
- Determining Backup Chargebacks
- Billing and Chargeback Summary
- NetBackup Front End Size by Duration and Workload
- NetBackup Front End Size by Workload
- Server Consumption Summary
- NetBackup Front End Size Usage Summary
- NetBackup Client Front End Size Report
- NetBackup Policy Front End Size Report
- NetBackup Images Report

Available Billing and Usage Reports

The Billing and Usage reports can be used to determine which business units are using your resources. In addition, you can use these reports to calculate chargeback costs.

If you're in the process of rolling out chargeback policies, consider using these reports as planning tools--to see what impact your policies might have on your business units or customers.

However, you need not charge your business units in order to find the Billing and Usage Reports useful. You simply can use the reports to gauge which business units are consuming the majority of your resources.

The following Billing and Usage reports are provided:

- See " Billing and Chargeback Summary " on page 934.
- See " Server Consumption Summary " on page 938.

Determining Media Chargebacks

The Billing and Chargeback Summary report provides a line item of each client's tape consumption. The Billing Summary Report calculates the chargeback based on billing and usage policies.

The report displays data based on the hosts/servers. This procedure assumes that you know which business units own specific hosts/servers.

There are a couple of limitations to the billing and chargeback reports:

- The report does not consider shared tape usage. If Client A wrote to tapes 1, 2 and 3, the report shows charges for the client for those three tapes. If Client B also wrote to tapes 1 and 2, the report lists charges for that client for those two tapes. Therefore, if both Client A and Client B belong to the same business unit, they are billed twice for two of the three tapes.
- The report always uses the current cost of a tape. If the cost changes and you generate a new report, the report reflects the new cost. However, if you run an old report from a year ago, it would also pick up the new cost and therefore not be the same as the original report. Therefore, before you change the unit cost, generate and export previous reports for comparison.

To determine chargebacks for a business unit or customer

1. Go to the Billing and Chargeback Summary report.

   See " Billing and Chargeback Summary " on page 934.

   Each client server row contains subtotals for GB and tape consumption.
2. To limit the data to a single business unit or department, modify the scope of the Billing and Chargeback Summary report.

3. Check the value in the **Total Cost** field, which is displayed at the right of the table. This value is the total charge to the business unit or customer.
4. If you want to itemize your chargebacks for your business unit/customer, consider the individual costs that comprise the total.

If you feel that the chargebacks are not accurate, do the following:

- Verify that the billing and chargeback policies are accurate.
- Ensure that the host group does not contain clients that belong to another business unit or customer.

Determining Backup Chargebacks

The Server Consumption Summary report displays backup consumption and the number of backup files captured for all hosts/servers that you specify. The report includes backups that were successful and those that have warnings. It does not include failed backups.

To determine server consumption chargebacks

1. Go to the Server Consumption Summary report. Hosts that do not have any consumption or backups for the period that you specified, do not appear in the report.

   See " Server Consumption Summary " on page 938.
2. Drill down into the clients, management servers, or backups to get more detail.
3. Export the data to your ERP software if applicable.

Billing and Chargeback Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Billing and Usage Reports > Billing & Chargeback Summary

---

Note: Tip: To limit the data to a single business unit or department, modify the scope of this report.

Note: Column aggregations are displayed in green on a separate line with the label **Aggregation**. Rollover the value to display the aggregation type: Avg., Min., Max., or Sum. Column filtering is not supported on reports with aggregated columns.

Note: If the scope of this report is for NetBackup-only host groups, three additional columns will be included: Total Dup (GB), Cost Per Dup'd GB, and Dup'd Usage Cost GB, which are set in the Billing and Usage Policies.

To view or modify billing and chargeback items, go to **Admin > Chargeback > Billing and Usage**.

Table 28-1 Billing and Chargeback Summary

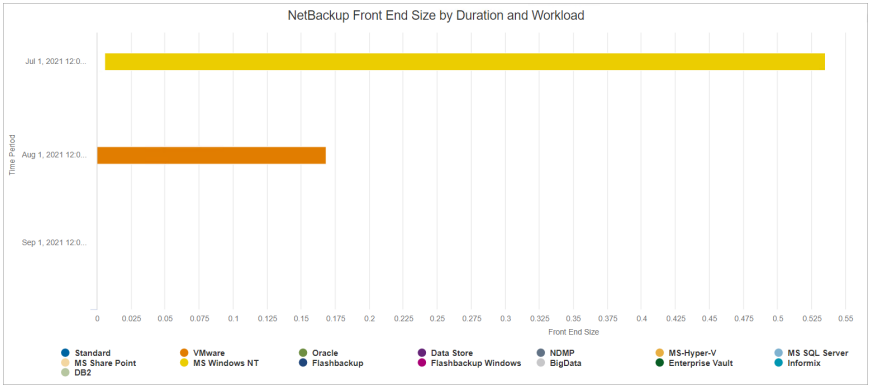
Column name	Description
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drilldown to the Host Details report. For other backup sources, the object name is displayed with a drilldown to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
Backed Up	The amount of storage consumed by the client server during the specified time period. Click this link to view a Job Summary report.
Cost Per GB	Amount the client is charged per GB of consumption.
GB Usage Cost	Backed Up * Cost per GB = GB Usage Cost
# of Tapes	The # of tapes consumed by the client during the specified time period. Click this link to view a Job Summary report.
Cost Per Tape	Amount the client is charged per tape consumed.
Total Tape Cost	# of Tapes * Cost per Tape = Total Tape Cost
Total Dup (GB)	Total duplicated storage (NetBackup-only scope)
Cost per Dup'd GB	The charge that has been configured for duplicated storage (NetBackup-only scope)
Dup'd Usage Cost	Total Dup * Cost per Dup'd = Dup'd Usage Cost (NetBackup-only scope)
Total Cost	GB Usage Cost + Total Tape Cost = Total Cost

NetBackup Front End Size by Duration and Workload

This report is a stacked bar chart report that plots the FETB collection size for each NetBackup workload or client against the collection time period. Since the bar chart is horizontal the bars appear adjacent to each other. Color codes are used to differentiate the collection size of each workload, as indicated by the legend.

This report is located at:

Backup Manager > **Billing and Usage Reports** > **NetBackup Front End Size by Duration and Workload**



NetBackup Front End Size Summary

This report appears as a drill down report of the NetBackup Front End Size by Duration and Workload report and for NetBackup Front End Size by Workload. This tabular report can also be converted to a chart using the **Actions** menu. The report displays the host and policy specific storage consumption.

The report also provides a further drill down view specific to every Primary Server, Client, Policy, and Job Start Date.

Table 28-2 NetBackup Front End Size Summary

Column name	Description
Primary Server	Name of the primary server.
Client	Name of the client backed up by the primary server.
Policy	Name of the policy used for backup.

Table 28-2 NetBackup Front End Size Summary *(continued)*

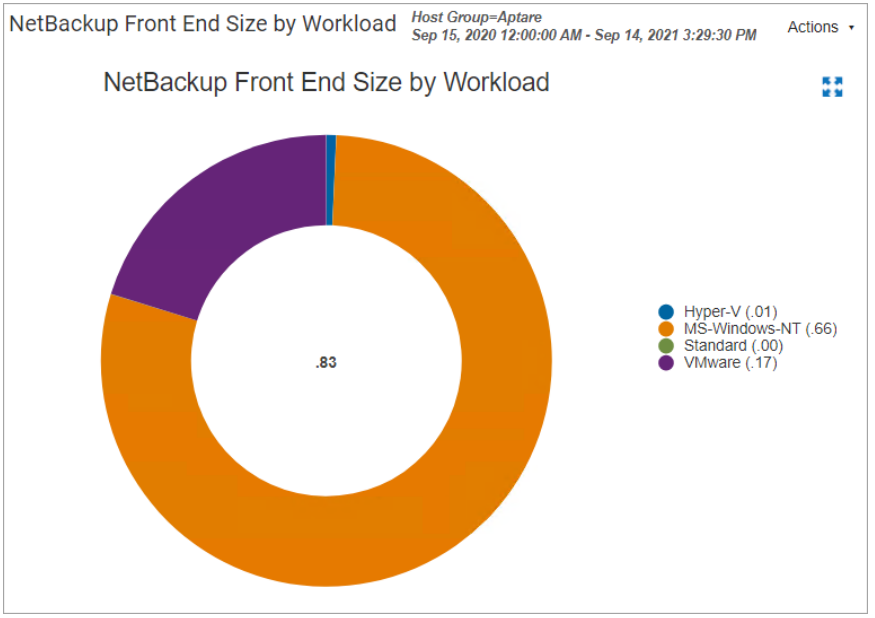
Column name	Description
Policy Type	Indicates the type of data being backed up. For example, Oracle indicates data being backed up is from Oracle Database.
Job Type	The type of backup job. For example, differential, full, incremental, synthetic, and so on.
Job Start Date	Date when the job was run.
Consumption	Storage consumption in TB.

NetBackup Front End Size by Workload

This report displays a donut view of the Front End Terabyte (FETB) distribution based on the workload type. You can click the donut portions to drill down to the Front End Size Summary report specific to the respective workload.

This report is located at:

Backup Manager > **Billing and Usage Reports** > **NetBackup Front End Size by Workload**.



Server Consumption Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Billing and Usage Reports > Server Consumption Summary

The consumption details are listed in a table, which you can sort by column headers. If you are in a large enterprise, this table can get unwieldy, so consider modifying the scope to filter the data by host group.

By default, the report lists backup sources in descending consumption order, with the heaviest consumer at the top. The following algorithm determines total consumption:

total disk backup + total tape backup = total consumption

Note: Tip: Sort the data by Tape Backup or Disk Backup to identify heaviest storage device consumers.

Table 28-3 Server Consumption Summary

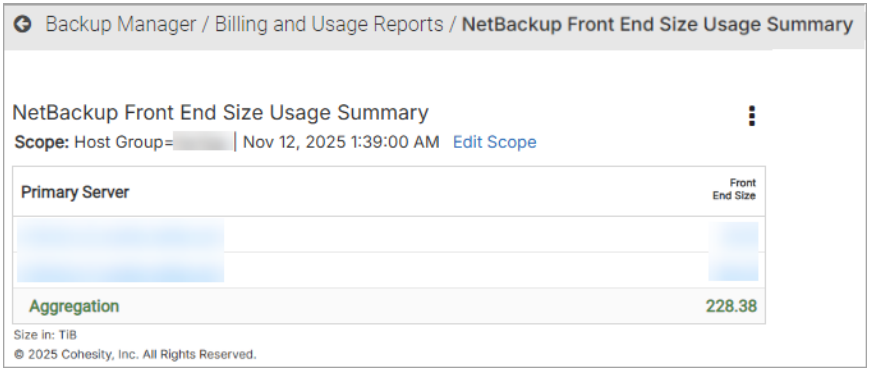
Column name	Description
Source	Name of backup source. If the backup source is a known host, host name is displayed with a drilldown to the Host Details report. For other backup sources, the object name is displayed with a drilldown to the Object Details report.
Source Type	Type of backup source, such as datastore, folder, database, virtual machine, and so on.
IP Address	IP address of the client.
Server	Backup server. Drill down to the Host Details report.
Backup Product	Backup product used to backup the client data.
Last Backup	Date and time of the last backup event.
Disk Backup MBytes	Amount of data backed up to disk.
Disk Backup Files	Number of files backed up to disk.
Tape Backup MBytes	Amount of data backed up to tape.
Tape Backup Files	Number of files backed up to tape.
Number of Tapes	Number of tapes used for the backup.

NetBackup Front End Size Usage Summary

This report displays summarizing Front End Terabyte (FETB) consumption of each NetBackup primary server collected in NetBackup IT Analytics.

This report is located at:

Backup Manager >> **Billing and Usage Reports** >> **NetBackup Front End Size Usage Summary**.



This is a drill down report. Click appropriate hyperlinks in the columns and NetBackup IT Analytics displays detailed report.

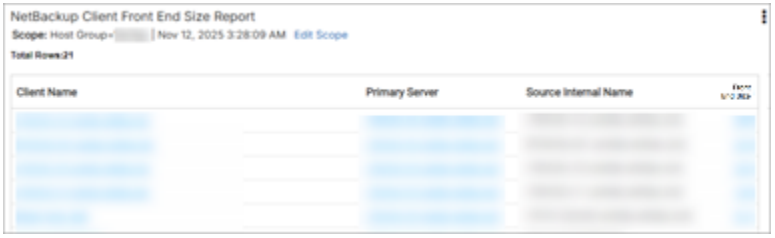
Click Front End Size value to view images backup wise Front End Size usage.

NetBackup Client Front End Size Report

This report displays summarizing Front End Terabyte (FETB) consumption of each Client of a NetBackup Primary Server.

This report is located at:

Backup Manager >> **Billing and Usage Reports** >> **NetBackup Client Front End Size Report**.



This is a drill down report. Click appropriate hyperlinks in the columns and NetBackup IT Analytics displays detailed report.

Click Client Name value to view the client information.

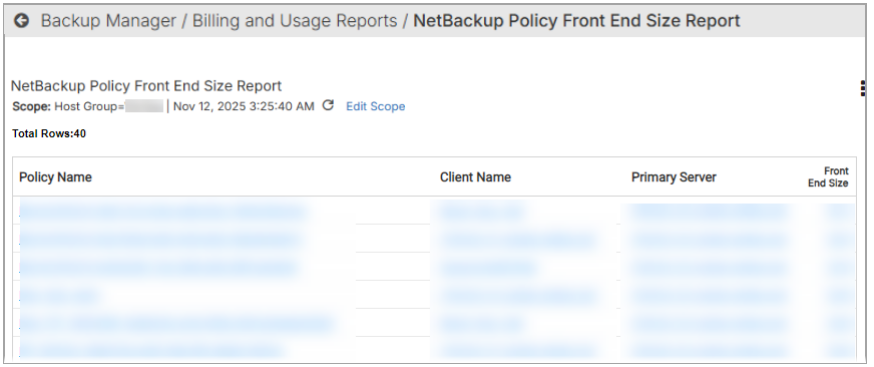
Click Front End Size value to view images backup wise Front End Size usage.

NetBackup Policy Front End Size Report

This report displays the summarized Front End Terabyte (FETB) consumption of each Policy of a NetBackup Primary Server.

This report is located at:

Backup Manager >> **Billing and Usage Reports** >> **NetBackup Policy Front End Size Report**.



This is a drill down report. Click appropriate hyperlinks in the columns and NetBackup IT Analytics displays detailed report.

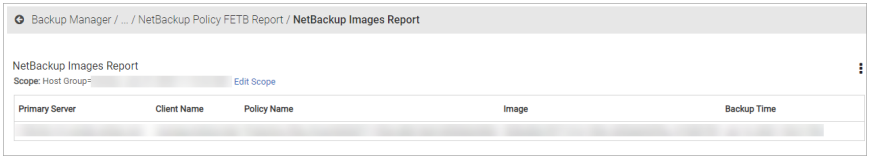
Click Policy Name values to view the policy information.

Click Front End Size value to view images backup wise Front End Size usage.

NetBackup Images Report

This is the drill down report from NetBackup Policy FETB report.

Click the appropriate values displayed on the FETB column. The image details for a FETB is displayed.



Chapter 29

Backup Policies Reports

This chapter includes the following topics:

- Commvault Simpana Policy Summary Report
- EMC Avamar Groups
- EMC NetWorker Groups
- HP Data Protector Backup Specification Summary
- IBM Spectrum Protect (TSM) Policy Domains
- NetBackup Job Status Summary by Policy Type
- Veritas NetBackup Policy
- Veritas Netbackup protection plan report

Commvault Simpana Policy Summary Report

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Backup Policies > Commvault Simpana Policies

---

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

A Commvault Policy provides the mechanism for applying a set of backup rules to a collection of clients.

Table 29-1 Commvault Simpana Policy Summary Report

Column name	Description
CommServe	Click to display details. See " Host Details " on page 838.
Policy Name	The Commvault Simpana storage policy.
Backup Copy	The Commvault Simpana storage policy copy name.
Default Copy	Indicates if the policy listed is the primary copy. Only one copy in a policy will be the default copy.
Legal Hold	Indicates if the policy listed is enabled for legal holds.
Comp	Indicates if hardware compression is enabled for the storage used by this backup copy.
Encrypted	Indicates if hardware encryption is enabled for the storage used by this backup copy.
Max Streams	Maximum number of simultaneous streams of data usable by the policy
(Retention) Days	Number of days that backups are retained by this copy. A hyphen indicates unlimited retention. Both the number of days and full cycles must be exceeded before a backup is aged out.
(Retention) Cycles	Number of full cycles that backups will be retained by this copy. A hyphen indicates unlimited retention. Both the number of days and full cycles must be exceeded before a backup is aged out. A cycle is a group of backup jobs that starts with each successful Full Backup.
(Retention) Data Aging	Indicates if data aging is enabled. Data Aging ages data based on the retention rules of a storage policy copy, and then removes that data based on the media recycling rules of the associated media.
(Retention) Managed Space	Indicates if managed disk space is in use. Managed Disk Space provides a way to prune data according to disk capacity in addition to the existing retention criteria.
(Libraries) Drive Pool	Name of the drive pool associated with the Library.

Table 29-1 Commvault Simpana Policy Summary Report *(continued)*

Column name	Description
(Libraries) Library	Backup location for the data.
(Dedup) Type	Indicates if data deduplication (if enabled) will be performed on the client or the server.
(Dedup) Level	Indicates if block-level or object-level deduplication will be performed. Commvault Simpana 9.0 only supports block-level deduplication.
(Dedup) Agent	The MediaAgent where deduplicated data is stored.
(Dedup) Store	The file path where deduplicated data is stored on the deduplication agent.
Clients	The number of subclients associated with this policy. The number of associated sub-clients will only be shown for the default copy.

Commvault Simpana Policy Backup Copy Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Commvault Simpana Policies

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

1. Generate the Commvault Simpana Policy Summary Report.
2. Click the **Policy Name**. The Policy Backup Copy Details are displayed.

Table 29-2 Commvault Simpana Policy Backup Copy Details

Column name	Description
Policy Name	The Commvault Simpana storage policy and links to the details.
Backup Copy	The Commvault Simpana storage policy copy name.
CommServe	Click to display Host Details report.
Default Copy	Indicates if the policy listed is the primary copy. Only one copy in a policy will be the default copy.
Library	Backup location for the data.
Drive Pool	Name of the drive pool associated with the Library.
Max Streams	Maximum number of simultaneous streams of data usable by the policy
Retention Days	Number of days to retain data as specified by the policy.
Retention Cycles	Number of full cycles that backups will be retained by this copy. A hyphen indicates unlimited retention. Both the number of days and full cycles must be exceeded before a backup is aged out. A cycle is a group of backup jobs that starts with each successful Full Backup.
Dedup Agent	The MediaAgent where deduplicated data is stored.
Dedup Store	The file path where deduplicated data is stored on the deduplication agent.
Enabled for Legal Holds	Indicates if the policy listed is enabled for legal holds.
Hardware Compression	Indicates if hardware compression is enabled for the storage used by this backup copy.
Hardware Encryption	Indicates if hardware encryption is enabled for the storage used by this backup copy.
Data Aging	Indicates if data aging is enabled. Data Aging ages data based on the retention rules of a storage policy copy, and then removes that data based on the media recycling rules of the associated media.
Managed Space	Indicates if managed disk space is in use. Managed Disk Space provides a way to prune data according to disk capacity in addition to the existing retention criteria.
Dedup Type	Indicates if data deduplication (if enabled) will be performed on the client or the server.

Table 29-2 Commvault Simpana Policy Backup Copy Details *(continued)*

Column name	Description
Dedup Level	Indicates if block-level or object-level deduplication will be performed.

EMC Avamar Groups

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Backup Policies > Avamar Groups

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

An Avamar Group provides the mechanism for applying a set of backup rules to a collection of clients. Backups are scheduled to run automatically, using the group specifications. Groups contain a Dataset, a Schedule, and a Retention Policy which comprise the backup policy for all members (clients) of the group. Each client that is to be backed up must be a member of at least one group; therefore, a Default Group provides a policy when a client has not been assigned to a more specific policy.

This report lists the groups defined on the selected Avamar servers.

Table 29-3 EMC Avamar Groups

Column name	Description
GSAN System Name	GSAN (Global Storage Area Network): The name of the Avamar Data Server subsystem. Links to the Avamar Server Details
Group Name	Links to the Avamar Group Detail
# of Clients	Indicates the number of clients that are members of the group. Links to the Group Client Summary.

Table 29-3 EMC Avamar Groups *(continued)*

Column name	Description
Dataset	Identifies the content of the backup; for example, filesystems, directories, or files. The Default Dataset is inherited from the Avamar Group, but can be overridden with a specific dataset. Click on the link to view the Dataset Details.
Schedule	The name of the Schedule that contains the scheduling details. Links to the Schedule Details.
Retention Policy	Indicates how long the backup will be kept. Backups older than what is specified in the policy are removed from the system. Links to the Retention Policy Details.
Active	Indicates if the group is enabled (active)

Avamar Server Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **GSAN System Name**.

Table 29-4 Avamar Server Detail

Column name	Description
Axion System ID	The unique Axion system ID
GSAN System Name	GSAN (Global Storage Area Network) Name: The user-assigned name of the Avamar Data Server. Links to the Avamar Server Details.
GSAN System ID	The Avamar server ID

Table 29-4 Avamar Server Detail *(continued)*

Column name	Description
HFS Address	The location of the hash filesystem where client backups are stored
HFS Port	Port address of the server
Local HFS Address	Location of hashes for the utility node
MCS Port	Management Console Server port number
Last Updated	Last date and time the details were updated on the Avamar server

Avamar Group Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **Group Name**.

Table 29-5 Avamar Group Detail

Column name	Description
Group Name	The name of the Avamar Group.
Dataset	The name of the Dataset covered by the Group backup.
Schedule	The schedule that contains the automated backup details.
Retention Policy	The retention policy name.
Enabled	Indicates if the group is active/enabled.
Date Created	The date and time the group was created.

Table 29-5 Avamar Group Detail *(continued)*

Column name	Description
Halt on Failure	Indicates if the group backups should cease if there is a failed backup.
Date Last Modified	The last time the group was modified.
Priority	Indicates the group's priority.
Read-only	Indicates if the group can be modified.
# of Retries	The number of retries attempted.
Run Once Only	Indicates if the group is running only one backup.
Skip Next	Indicates if the next scheduled backup is to be skipped.
Target DPN	The Avamar DPN (Data Protection Network) server that is to be used for the group.
Timeout in Minutes	Time out in minutes.
Exists in Catalog	Indicates if this retention policy is included in the Avamar Group catalog

Clients for Group

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **# of Clients** link.

Table 29-6 Clients for Group

Column name	Description
Client Name	The client that is backed up as part of the group policy

Table 29-6 Clients for Group *(continued)*

Column name	Description
Group Name	The group name
Enabled	Indicates if the client is assigned to a group
Restore Only	Indicates if this client can do only restores
Use Client Dataset	Indicates if the client is using its own dataset, overriding the group dataset
Use Client Retention Policy	Indicates if the client is using its own retention policy, overriding the group retention policy
Use Client Retry	Indicates if the client is using its own retry count, overriding the group retry count
Use Client Timeout	Indicates if the client is using its own timeout, overriding the group timeout setting

Avamar Client Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **# of Clients** link and in the Clients for Group window, click a **Client Name**.

Each client known to the Avamar Administrator server displays the following details.

Table 29-7 Avamar Client Detail

Column name	Description
Client Summary	

Table 29-7 Avamar Client Detail *(continued)*

Column name	Description
Client Name	The Avamar backup client (the source)
Domain	The specific zone on the Avamar server
Operating System	The Avamar server's operating system
Agent Version	The version of the client agent, which communicates with the Avamar Management Console Server (MCS)
Client ID	The Client ID
Created	The creation date for this Avamar client
Last Modified	The last date and time that the client information was updated via the Avamar Management Console
Contact	
Name	Name of the person to contact
Email	Email of the person to contact
Location	Location of the contact
Notes	Notes that may be required for client management
Phone	Phone number of the contact
Authorization Properties	
User-initiated Backups	Indicates if this client is configured to allow user-initiated backups
User-Initiated Backup File Selections	Indicates if this client is configured to allow file selection on user-initiated backups
Client to Override Group Dataset	Indicates if the client can override the group dataset
Allow Client Overtime	Indicates if the client can ignore the scheduling window end time
Override of Group Retry Count	Indicates if the client can override the group retry count
Override of Group Retention Policy	Indicates if this client can override the standard retention policy for user-initiated backups

Table 29-7 Avamar Client Detail *(continued)*

Column name	Description
Allow Group Overtime	ALWAYS allow a scheduled group backup to exceed the schedule duration setting; NEXT allows only the next scheduled group backup to exceed the schedule duration setting; NEXT_SUCCESS allows scheduled group backups to exceed the schedule duration setting until a successful backup completes; NEVER allow a scheduled group backup to exceed the schedule duration setting
Management Console Server Access	
MCS Can Page this Client	Indicates if the Management Console Server is configured to page this client
Page Address	The IP address that the Avamar server uses to communicate with this client
Page Address Locked	A setting of Yes indicates that the page IP address cannot be updated automatically by the Management Console Server.
Page Port	The data port that the Avamar server uses to contact this client
Last Check-in Date	The timestamp of the last time this client checked in with the Avamar server
Date Registered	The date and time that this client was added to the Avamar system
Backup Details	
Has Backups	Indicates if this client has backups
Last Backup Date	The timestamp of this client's last backup
Plugin for Backup	Lists the plug-in that was used for the last backup
Restore Only	Indicates if this client can do only restores
Retry Count	The client's connection retry count
Overrides Retention Policy	Indicates if the client retention policy, instead of the standard group retention policy, is being used for this user-initiated backup.
Timeout	Connection timeout

Table 29-7 Avamar Client Detail *(continued)*

Column name	Description
Allows Override Timeout	Indicates if the client can override the group's timeout setting
Enabled	Indicates if the client is enabled in a group, enabling it to be backed up by the Avamar server
Registered	Backup is registered

Avamar Dataset Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **Dataset** link.

Files, directories, and file systems comprise an Avamar Dataset, providing a policy that can be associated with multiple Avamar Groups.

Table 29-8 Avamar Dataset Detail

Column name	Description
Dataset	The dataset name.
Domain	The specific zone on the Avamar server that is associated with this dataset.
All Data Saved	Indicates if the dataset will save all data.
Is Link	Indicates if this dataset is a pointer to another dataset.
Parent Dataset Name	If this dataset is linked to another dataset, this field identifies the parent dataset.

Table 29-8 Avamar Dataset Detail *(continued)*

Column name	Description
Link Name	The name of the dataset to which this dataset links.
Read-only	Indicates if the dataset can be modified.
Exists in Catalog	Indicates if this retention policy is included in the Avamar Dataset catalog.

Avamar Schedule Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **Schedule** link.

Table 29-9 Avamar Schedule Detail

Column name	Description
Schedule	The backup schedule name.
Domain	The specific zone on the Avamar server; the Avamar domain to which this schedule applies.
Description	Schedule description.
End Policy	The termination setting: Never end, Run n times, or End on a date.
End Date	The date the backup is scheduled to terminate.
Enabled	Indicates if the schedule is active.
End Recurrence	In conjunction with the End Policy, this is either a specific date or the number of times a schedule should run.

Table 29-9 Avamar Schedule Detail *(continued)*

Column name	Description
First Start Date	First start data of the schedule.
Is Fixed Interval	Indicates if the scheduling interval is fixed.
Interval Mod	This is the modifier associated with recurring schedules: day, every, first, fourth, hour, last, second, third.
Is Link	Indicates if this schedule is linked to another schedule.
Last Check Date	Last check timestamp.
Last Start Date	The last time this schedule started.
Can Override End Time	Indicates if the end time can be overridden.
Read-only	Indicates if this schedule can be modified.
Recurrence Counter	The number of times the schedule has run.
Recurrence Interval	The period set for recurrence: daily, hourly, weekly, monthly, or yearly.
Start Duration	The duration of the start scheduling window.
Start Time	The scheduling window's start date and time.
Time Zone	The time zone of the schedule.
Total Duration	The scheduling window's total duration.
Exists in Catalog	Indicates if this is a schedule provided in the Avamar Schedule Catalog.
Last Updated	The date and time that the data collector last collected details.

Avamar Retention Policy Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > Avamar Groups

Click a **Retention Policy**link.

An Avamar Retention Policy dictates how long a backup will remain on an Avamar server before being automatically deleted.

Table 29-10 Avamar Retention Policy Detail

Column name	Description
Domain	The specific zone on the Avamar server; the Avamar domain to which this retention policy applies.
Retention Policy	The name of the retention policy.
Daily	Indicates how many daily backups will be retained.
Weekly	Indicates how many weekly backups will be retained.
Monthly	Indicates how many monthly backups will be retained.
Yearly	Indicates how many yearly backups will be retained.
Duration	The duration of the backup retention.
Enabled	Indicates if the retention policy is enabled.
Expiration Date	The date on which backups will be deleted from the Avamar server.
Is Link	If this is set to Yes, this retention policy refers to another retention policy.
Parent Name	The name of the parent retention policy to which this policy refers; only exists if the link is set to Yes.
Allow Override	If set to Yes, an advanced policy is used for scheduled backups.
Policy	Undefined, Compute expiration date, Static expiration date, No expiration date.
Read-only	Indicates if the retention policy can be modified.
Exists in Catalog	Indicates if this retention policy is included in the Avamar Retention Policy catalog.

Table 29-10 Avamar Retention Policy Detail *(continued)*

Column name	Description
Last Updated	The date and time that the data collector last collected details.

EMC NetWorker Groups

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Backup Policies > NetWorker Groups

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

EMC NetWorker provides a Group resource to automate backups on a group-by-group basis. By setting up groups, you can manage the load on backup servers.

Table 29-11 EMC NetWorker Groups

Column name	Description
Server	The EMC NetWorker server that is managing the backups for the associated group. Click on this link to access Host Details.
Group Name	A group of clients to be backed up. Click this name to view the list of client instances in the details. See " Group Detail " on page 958.
Start Time	If Autostart is Enabled, the group will begin being backed up at this time.
Interval	The frequency with which autostart will initiate backups.

Table 29-11 EMC NetWorker Groups *(continued)*

Column name	Description
Autostart	This attribute indicates if the group will be backed up automatically for each time period specified as the interval. Values include: E = Enabled, which means clients will be backed up automatically at the time specified in Start Time; D = Disabled; S = Start now, which means that clients are backed up immediately, overriding the start time.
Schedule Override	NetWorker allows you to override a scheduled backup by selecting a default group and initiating a Start Now autostart.
Status	Backup Status for the group can be one of the following: Running, Idle, Finished, Not Finished (backup exited without completing), Preview Run (test a backup configuration).

Group Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > NetWorker Groups

Click a Group Name.

A NetWorker Group provides the mechanism for establishing which clients will be backed up together and then scheduled for their backups. This report displays the NetWorker Group Control details.

The following NetWorker-specific details are provided in this report.

Table 29-12 Group Detail

Column name	Description
Group Name	A group of clients to be backed up.

Table 29-12 Group Detail *(continued)*

Column name	Description
Schedule Override	NetWorker allows you to override a scheduled backup by selecting a default group and initiating a Start Now autostart.
Start Time	When the group backup is scheduled to begin.
Override Level	Override levels can be configured for a specific date in a NetWorker backup schedule, enabling you to override the backup level setting for every client in that group.
Auto-Start	When enabled, this attribute triggers regularly scheduled backups.
Interval	The frequency with which autostart will initiate backups.
Status	Backup Status for the group can be one of the following: Running, Idle, Finished, Not Finished (backup exited without completing), Preview Run (test a backup configuration)
Force Incremental	When this field is set to Yes, if two full backups start within a 24-hour period, NetWorker forces the second one to be an incremental backup. (for an interval attribute of less than 24 hours)
Last Start/End	The last time the group started/ended a save operation.
Duration	Duration of the last group save.
Clones & Clone Pool	NetWorker uses a cloning process to make copies of saved data. In addition to copying the data, NetWorker keeps track of the details of the cloning operation in a client file index and in the media database. When the Clones attribute is set to Yes, NetWorker automatically clones a save set whenever it is backed up. The Clone Pool indicates where the clones are sent.
Auto Restart	This restart attribute indicates if a group should be automatically restarted after an incomplete save operation.
Success Threshold	Indicates the success or failure of the groups save set operations.
Client Retries	You can configure NetWorker to be persistent in cases where a client may be hung. Based on the Client Retries value, NetWorker will make multiple attempts at backing up a client before it moves on to other clients.
Inactivity Time-out	The inactivity (in minutes) before NetWorker determines that the client is hung.

Table 29-12 Group Detail *(continued)*

Column name	Description
Save Group Parallelism	The NetWorker server can perform simultaneous backups in either of the following ways: ■ multiple save sets from multiple clients ■ multiple save sets from a single client.
Client Instance Name	Click this link to access the details. See " Client Instance Detail " on page 960. report.
Scheduled Backup	Indicates if there is a backup scheduled.
Schedule	Schedule description.
Last Successful Backup	Click this link to access the details. See " Job Details Report " on page 770.
Comment	Comments related to the schedule.

Client Instance Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > NetWorker Groups

Click a **Group Name** and then click a **Client Instance Name**.

HP Data Protector Backup Specification Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Backup Policies > HP DP Backup Specification

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

The HP Data Protector Backup Specification report lists the Backup Specification, with the objects that are backed up, the servers and the next time the backup will be performed. Note that the objects can be entire disks or volumes, or sub-sets of these, such as directories or the Windows Registry.

Table 29-13 HP Data Protector Backup Specification Summary

Column name	Description
Specification	Links to the details. See " HP Data Protector Backup Specification Detail " on page 961.
Server	Backup server.
Object Type	The type of data for which the backup specification is configured--database entity or a disk image (raw disk)--for example, filesystem.
# of Clients	Number of clients backed up by this backup specification.
Next Backup Mode	Full, Incr, Incr 1, Incr 2, Incr 3, Incr 4, Incr 5, Incr 6, Incr 7, Incr 8, Incr 9.
Next Backup Date	Date and time of the next backup.

HP Data Protector Backup Specification Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports

cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > HP DP Backup Specification

Click a **Specification** link.

The HP Data Protector Backup Specification Detail report lists the details of the specification, such as mount points and clients. Each of the media must be in a media pool.

Table 29-14 HP Data Protector Backup Specification Detail

Column name	Description
Specification	Specification Name.
Object Type	The type of data for which the backup specification is configured--database entity or a disk image (raw disk)--for example, filesystem.
Next Backup Mode	Full, Incr, Incr 1, Incr 2, Incr 3, Incr 4, Incr 5, Incr 6, Incr 7, Incr 8, Incr 9.
Next Backup Date	Date and time stamp of the next backup.
Pre Exec Script Name	Pre-execution script name.
Post Exec Script Name	Post-execution script name.
Status	The current status of the specification.

Object's Latest Backup and Not Configured Filesystems

A table lists the mount points, with the scheduled backup type, media pool, and last full & incremental backup dates.

Table 29-15 HP Data Protector Backup Specification Detail

Column name	Description
Mount Point	A mount point provides the access to the directory structure for a disk or logical volume on the client, to locate the backup object.
Client	The client links to the Mount Point Details.
Next Backup Mode	Type of backup that is scheduled next, for example, Full.
Media Pool	The media pool in which this data was stored.
Last Full Backup	Date and time of the last full backup for this mount point.
Last Incr Backup	Date and time of the last incremental backup for this mount point.

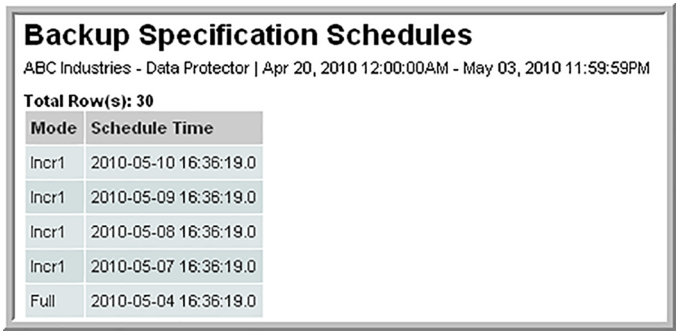
Not Configured Filesystems

Table 29-16 HP Data Protector Backup Specification Detail

Column name	Description
Mount Point	Mount points listed in the Not Configured Filesystems table indicate mount points that are not backed up.
External Name	The name of the server.

Backup Specification Schedules

A table lists the recent backup specification schedules, if any: Mode and Schedule Time.



Backup Specification Schedules

ABC Industries - Data Protector | Apr 20, 2010 12:00:00AM - May 03, 2010 11:59:59PM

Total Row(s): 30

Mode	Schedule Time
Incr1	2010-05-10 16:36:19.0
Incr1	2010-05-09 16:36:19.0
Incr1	2010-05-08 16:36:19.0
Incr1	2010-05-07 16:36:19.0
Full	2010-05-04 16:36:19.0

IBM Spectrum Protect (TSM) Policy Domains

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Backup Policies > IBM Spectrum Protect (TSM) Policy Domains

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

This report is specific to IBM Spectrum Protect (TSM).

Table 29-17 IBM Spectrum Protect (TSM) Policy Domains

Column name	Description
Instance Name	The policy instance.
Domain Name	A backup group of filesystems and drives. Click this link to access the details. See " Policy Domain Details " on page 965.
Description	The domain's description.
Activated Management Class	A management class is a policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. A copy group determines how the TSM server manages backup versions or archived copies of the file. The space management attributes determine whether the file is eligible for migration by the space manager client to server storage, and under what conditions the file will be migrated.
Backup Retention Days	If a TSM management class is deleted, files that were associated with it need to rebind to another management class. TSM will keep the backup files for a specified number of days until the server is able to rebind the files to an appropriate management class. After this period, the files will no longer be retained. Retention details are defined in the TSM copy group and default retention grace periods for the domain.

Table 29-17 IBM Spectrum Protect (TSM) Policy Domains *(continued)*

Column name	Description
Archive Retention Days	If a TSM management class is deleted, files that were associated with it need to rebind to another management class. TSM will keep the archived files for a specified number of days until the server is able to rebind the files to an appropriate management class. After this period, the files will no longer be retained.
Number of Nodes	The number of file servers registered with the TSM server.
Active	When the instance became active.

Policy Domain Details

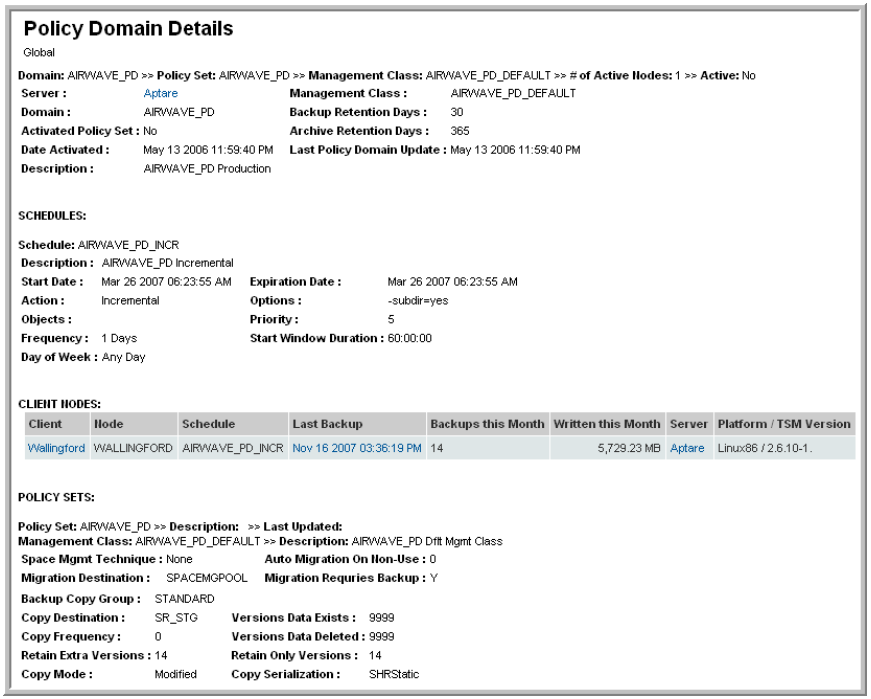
Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > IBM Spectrum Protect (TSM) Policy Domains

Click a **Domain Name**.



Policy Domain Details

Global

Domain: AIRWAVE_PD >> **Policy Set:** AIRWAVE_PD >> **Management Class:** AIRWAVE_PD_DEFAULT >> **# of Active Nodes:** 1 >> **Active:** No

Server :	Aptare	**Management Class :**	AIRWAVE_PD_DEFAULT
Domain :	AIRWAVE_PD	**Backup Retention Days :**	30
Activated Policy Set :	No	**Archive Retention Days :**	365
Date Activated :	May 13 2006 11:59:40 PM	**Last Policy Domain Update :**	May 13 2006 11:59:40 PM
Description :	AIRWAVE_PD Production		

SCHEDULES:

Schedule: AIRWAVE_PD_INCR

Description :	AIRWAVE_PD Incremental		
Start Date :	Mar 26 2007 06:23:55 AM	**Expiration Date :**	Mar 26 2007 06:23:55 AM
Action :	Incremental	**Options :**	-subdir=yes
Objects :		**Priority :**	5
Frequency :	1 Days	**Start Window Duration :**	60:00:00
Day of Week :	Any Day		

CLIENT NODES:

Client	Node	Schedule	Last Backup	Backups this Month	Written this Month	Server	Platform / TSM Version
Wallingford	WALLINGFORD	AIRWAVE_PD_INCR	Nov 16 2007 03:36:19 PM	14	5,729.23 MB	Aptare	Linux86 / 2.6.10-1.

POLICY SETS:

Policy Set: AIRWAVE_PD >> **Description:** >> **Last Updated:**
Management Class: AIRWAVE_PD_DEFAULT >> **Description:** AIRWAVE_PD Dflt Mgmt Class

Space Mgmt Technique :	None	**Auto Migration On Non-Use :**	0
Migration Destination :	SPACEMGPOOL	**Migration Requries Backup :**	Y
Backup Copy Group :	STANDARD		
Copy Destination :	SR_STG	**Versions Data Exists :**	9999
Copy Frequency :	0	**Versions Data Deleted :**	9999
Retain Extra Versions :	14	**Retain Only Versions :**	14
Copy Mode :	Modified	**Copy Serialization :**	SHRStatic

A Policy Set is an object that contains a group of management class definitions that exist for a policy domain. Several policy sets can exist within a policy domain, but only one policy set can be active at a time.

NetBackup Job Status Summary by Policy Type

This report is located here:

Backup Manager > **Backup Policies** > **NetBackup Job Status Summary by Policy Type**

This report shows a status summary for NetBackup jobs by Policy Type for the selected clients in scope. The stacked bar chart shows totals for successful, partially successful, and failed jobs for the selected clients selected in the report scope. Hover your mouse pointer over the bars to view the job counts. Run this report on specific groups of clients by modifying the scope selector. This report is equivalent to the following NetBackup OpsCenter reports:

- Job Success rate by policy type (Graphical)
- Job Success rate by policy type (Tabular)

Drill-down reports

Click the bar to see the drill-down view containing the selected job details. The job's summary status (Successful/Partial/Failed) is used to drill-down into the NetBackup Job Summary Status by Policy Type details report. The drill-down reports provide details of all jobs for a selected client and policy type. The drill-down reports use run-time parameters that are passed from the main report. The tables in the drill-down reports provide details such as policy and schedule names, job start and end time, and job duration. The table for failed jobs also includes a column for the job exit status.

Three drill-down reports can be viewed from this report:

- Number of Successful Jobs
- Number of Partially Successful Jobs
- Number of Failed Jobs

Veritas NetBackup Policy

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Backup Manager > Backup Policies > NetBackup Policies

Note: The Backup Policies folder is available only if one of the following subsystems is licensed for data collection: EMC Avamar, Commvault Simpana, IBM Spectrum Protect (TSM), HP Data Protector, or Veritas NetBackup.

This report provides the details of the underlying NetBackup Policies that have been defined for your environment.

Table 29-18 Veritas NetBackup Policy

Column name	Description
Management Server	The server that is executing the backup policy. See “ Host Details ” on page 838.

Table 29-18 Veritas NetBackup Policy *(continued)*

Column name	Description
Policy Name	Name assigned to this policy. See " Backup Policy Details " on page 968.
Type	Indicates the specific components it is managing--for example, MS-SQL server, Oracle, NDMP, MS-Windows.
Residence	Indicates where the Storage Lifecycle Policy (SLP) resides.
Volume Pool	The name of the volume pool to which the media belongs.
# of Clients	# of clients being backed up.
Status	The backup status code returned by NetBackup.
Deleted	Policy has been deleted and does not exist in the NetBackup catalog.

Backup Policy Details

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

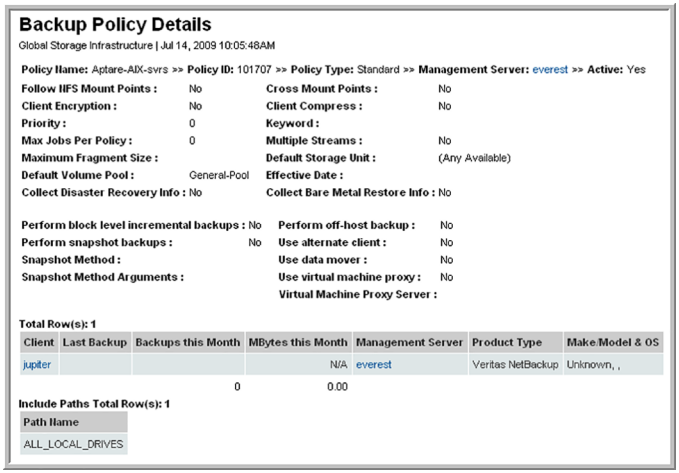
Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Backup Manager > Backup Policies > NetBackup Policies

Click the **Policy Name**.

The Policy Details include not only a list of clients covered by the policy, but also the scheduling details. Often the list of clients is long, so you'll need to scroll to the bottom of the report to see the scheduling details.



Backup Policy Details

Global Storage Infrastructure | Jul 14, 2009 10:05:48AM

Policy Name: Aptare-AIX-svrs >> **Policy ID:** 101707 >> **Policy Type:** Standard >> **Management Server:** everest >> **Active:** Yes

Follow NFS Mount Points :	No	**Cross Mount Points :**	No
Client Encryption :	No	**Client Compress :**	No
Priority :	0	**Keyword :**	
Max Jobs Per Policy :	0	**Multiple Streams :**	No
Maximum Fragment Size :		**Default Storage Unit :**	(Any Available)
Default Volume Pool :	General-Pool	**Effective Date :**	
Collect Disaster Recovery Info :	No	**Collect Bare Metal Restore Info :**	No

Perform block level incremental backups :	No	**Perform off-host backup :**	No
Perform snapshot backups :	No	**Use alternate client :**	No
Snapshot Method :		**Use data mover :**	No
Snapshot Method Arguments :		**Use virtual machine proxy :**	No
		Virtual Machine Proxy Server :	

Total Row(s): 1

Client	Last Backup	Backups this Month	MBytes this Month	Management Server	Product Type	Make/Model & OS
jupiter			N/A	everest	Veritas NetBackup	Unknown, ,
		0	0.00			

Include Paths Total Row(s): 1

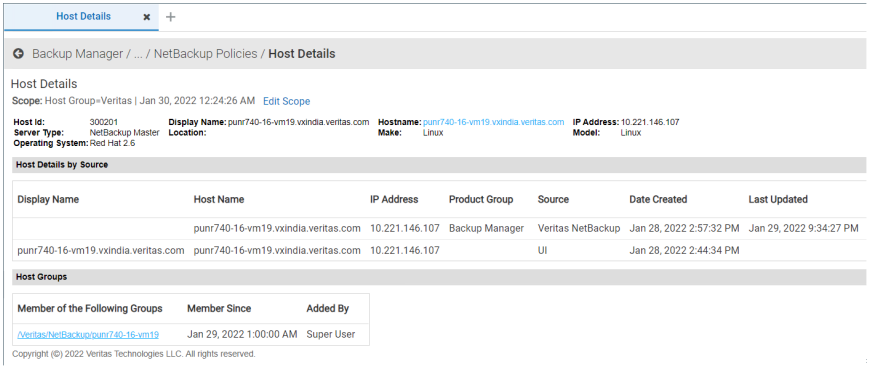
Path Name
ALL_LOCAL_DRIVES

Click **Client** hypelink and NetBackup IT Analytics displays relevant information of that client.

In case of VMware Intelligent policy, the NetBackup IT Analytics retrieves the following details of the clients:

- VM UUID
- VM Display Name
- VM DNS Name
- VM Host Name
- VM Instance Name

Note: The **Client** displayed in the report is the display name.



Host Details

Backup Manager / ... / NetBackup Policies / **Host Details**

Host Details

Scope: Host Group=Veritas | Jan 30, 2022 12:24:26 AM Edit Scope

Host Id: 300201 **Server Type:** NetBackup Master **Operating System:** Red Hat 2.6 **Display Name:** punr740-16-vm19.vxindia.veritas.com **Location:** **Hostname:** punr740-16-vm19.vxindia.veritas.com **Make:** Linux **IP Address:** 10.221.146.107 **Model:** Linux

Host Details by Source

Display Name	Host Name	IP Address	Product Group	Source	Date Created	Last Updated
	punr740-16-vm19.vxindia.veritas.com	10.221.146.107	Backup Manager	Veritas NetBackup	Jan 28, 2022 2:57:32 PM	Jan 29, 2022 9:34:27 PM
punr740-16-vm19.vxindia.veritas.com	punr740-16-vm19.vxindia.veritas.com	10.221.146.107		UI	Jan 28, 2022 2:44:34 PM	

Host Groups

Member of the Following Groups	Member Since	Added By
/Veritas/NetBackup/punr740-16-vm19	Jan 29, 2022 1:00:00 AM	Super User

Copyright (©) 2022 Veritas Technologies LLC. All rights reserved.



Table 29-19 Backup Policy Details

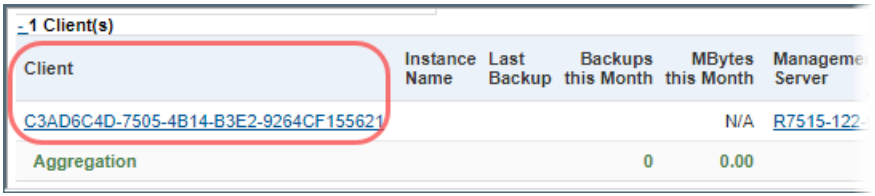
Column name	Description
Policy Details	This section includes details that are found in the EMC Avamar Groups page. In addition, the following Yes/No indicators reflect the current policy settings: ■ Follow NFS Mount Points ■ Cross Mount Point ■ Client Encryption ■ Client Compress ■ Priority ■ Keyword ■ Max Jobs Per Policy ■ Multiple Streams See “ EMC Avamar Groups ” on page 946.

Table 29-19 Backup Policy Details *(continued)*

Column name	Description
Client Details	■ Client - The Client name. ■ Last Backup - Date of the last backup job. ■ Backups this Month- Number of backup jobs for the current month. ■ MBytes this Month- Amount of data backed up for the current month. ■ Management Server - The Server responsible for backups. ■ Product Type - Backup solution product name. ■ Make/Model & OS- Client's details, if known.

Known Limitation for HyperV Policy

In case of HyperV Intelligent policy, the client name returned by Veritas NetBackup is always VM UUID regardless of the Primary VM Identifier of the HyperV Virtual Machine. This behavior may lead to NetBackup IT Analytics reporting some/all such clients as unprotected.



Veritas Netbackup protection plan report

Overview

NetBackup Protection Plan report has been incorporated to display Protection plan details. It lists down all the protection plans currently created under NetBackup along with number of assets protected by the plan.

It also displays the protection plan details and protected client details in drill down summary reports.

Accessing the report

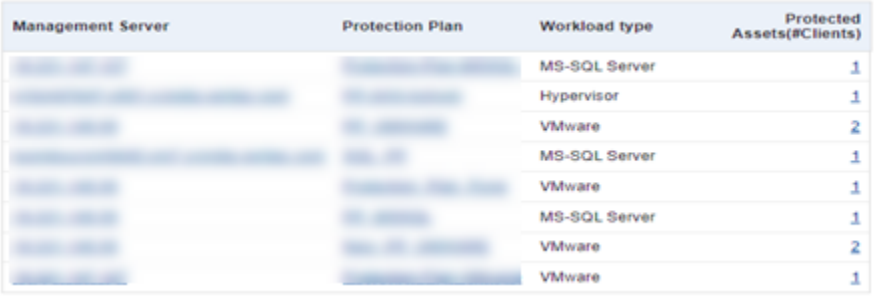
The NetBackup Protection Plan report can be accessed in NetBackup IT Analytics.

Navigate to **Backup Manager** > **Backup Policies** > **NetBackup Protection Plan** report.

The report displays the following columns for further drill down detailed report.

Table 29-20 NetBackup Protection Plan overview report

Column	Description
Management Server	NetBackup Primary Server name.
Protection Plan	NetBackup Protection Plan Name which is created at the NetBackup side and under which clients are protected along with the policies. Policies are created dynamically and its name holds protection plan name.
Workload Type	NetBackup Policy Type Name.
Protected Assets (#Clients)	Number of clients which are getting protected under policy.

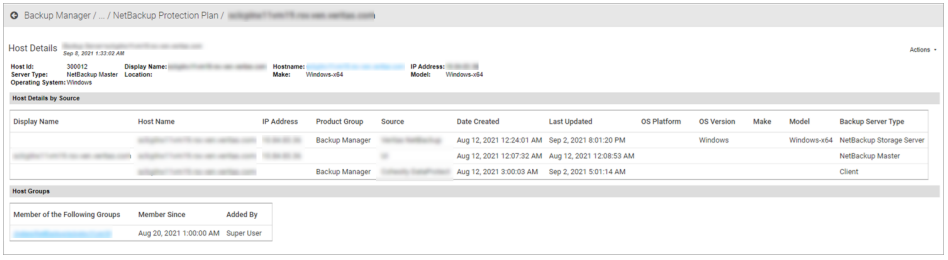


Management Server	Protection Plan	Workload type	Protected Assets(#Clients)
[redacted]	[redacted]	MS-SQL Server	1
[redacted]	[redacted]	Hypervisor	1
[redacted]	[redacted]	VMware	2
[redacted]	[redacted]	MS-SQL Server	1
[redacted]	[redacted]	VMware	1
[redacted]	[redacted]	MS-SQL Server	1
[redacted]	[redacted]	VMware	2
[redacted]	[redacted]	VMware	1

File Analytics further supports drill down reports for more details.

Management Server drill down report

Click required value displayed in the **Management Server** column of the overview report. NetBackup IT Analytics displays details of the server.



Protection Plan drill down report

Click required value displayed in the **Protection Plan** column of the overview report. NetBackup IT Analytics displays **NetBackup Protection Plan Policies** report.

The report displays lists of policies created under the protect plan with other relevant information of the policy such as policy type, schedule name, retention days and status.

Protected Assets drill down report

Click required value displayed in the **Protection Plan** column of the overview report. NetBackup IT Analytics displays **NetBackup Protection Plan Clients** report.

The report displays lists of clients protected under the protection plan with relevant details of clients such as last backed up date and number of backups happened in last month.

Chapter 30

Public Cloud Reports

This chapter includes the following topics:

- About Public Cloud Reports
- AWS Reports
- Microsoft Azure Reports

About Public Cloud Reports

Public Cloud Reports provide a consolidated view of all cloud reports to help you optimize costs of cloud resources and migration use cases. The reports in this category provide:

- View of health and operational status of public cloud workloads
- Actionable analytics
- Correlated insights with on-prem IT infrastructure environment
- Monitoring and management of cloud assets

AWS Reports

AWS reports list:

- AWS Account Billing Detail
- AWS Account Listing
- AWS Billing Summary Over Time
- AWS Product Summary
- AWS S3 Bucket Summary

- AWS EC2 Instance Summary
- AWS EC2 Snapshot
- AWS EC2 Snapshot Detail

AWS Account Billing Detail

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > AWS Account Billing Detail

View Amazon Web Services (AWS) billing costs associated with an account.

Details for each discovered product include usage as well as blended and unblended rates and costs.

Use the Custom Filter in the scope selector to isolate specific types of AWS data.

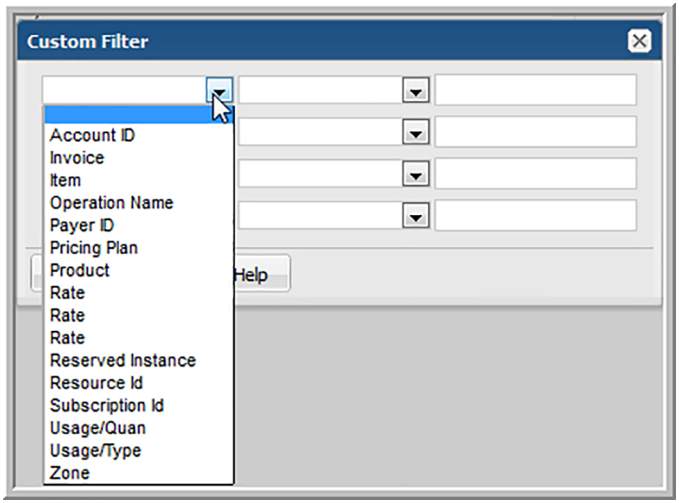


Table 30-1 AWS Account Billing Detail

Column name	Description
Invoice	The 8-digit identifier for the AWS invoice. This value maps to the invoice number on the PDF file for the corresponding month. If the data is in the current month this field value is Estimated.
Payer ID	The AWS account ID, a 12-digit number.
Account ID	The AWS account ID, a 12-digit number.
Product	Long name of the AWS product.
Item	Description of the type of usage for this line item. Includes the average price of any usage that is based on a volume pricing tier. Prices shown are rounded to three decimal places.
Usage/End	The end of the applicable charge period. All usage charges and recurring fees between the period start and end are included in your monthly invoice.
Resource ID	The resource ID or name (e.g., EC2 instance ID, S3 bucket name). May be an arbitrary ARN.
Zone	Specifies the Availability Zone in which an AWS object resides or a charged usage occurred.
Subscription ID	Unique identifier for the AWS product described by the line item.
Pricing Plan	Numeric ID for the rate listed in the AWS pricing plan for the corresponding service.
Rate	Numeric rate ID that maps to the Item description. The value is assigned by AWS for the Item description, and identifies the billing rate for computation purposes.
Usage/Type	Describes the specific usage of the line item.
Operation Name	Describes the operational details of the usage line item. For example, a value of RunInstances indicates the operation of an Amazon EC2 instance.
Reserved Instance	Indicates whether the line item was calculated at a Reserved Instance rate. Values are Y for Yes and N for No.
Usage/Quan	Computing resources used, in the units specified by each service.
Blended Rate	For consolidated billing accounts, the effective rate for the line item, calculated as an average of the cost of identical Amazon EC2 instances operating in that hour in the same Availability Zone.

Table 30-1 AWS Account Billing Detail *(continued)*

Column name	Description
Blended Cost	Cost for the operation listed in the line item based on the blended rate.
UnBlended Rate	Published rate per hour of operation for the usage type specified in the Item Description field.
UnBlended Cost	Cost for the operation listed in the line item based on the unblended rate.

AWS Account Listing

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > AWS Account Listing

View a list of Amazon Web Services (AWS) accounts with month-to-date costs.

Access the relevant AWS Product Summary report for additional details.

Table 30-2 AWS Account Listing

Column name	Description
Account ID	The AWS account ID, a 12-digit number.
Payer ID	The AWS account ID, a 12-digit number.
Zone	Specifies the Availability Zone in which an AWS object resides or a charged usage occurred.
Alias	The account alias.
Number of Products	Long name of the AWS product. Click the link to access the AWS Product Summary.
Blended Cost	Cost for the operation listed in the line item based on the blended rate. Click the link to access the AWS Account Billing Detail.
Tags	Tags associated with the account.
Last Invoice Date	Date of the latest Usage Start Date that was collected for this month.

Table 30-2 AWS Account Listing *(continued)*

Column name	Description
Last Updated	Last updated date.

AWS Billing Summary Over Time

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > AWS Billing Summary Over Time

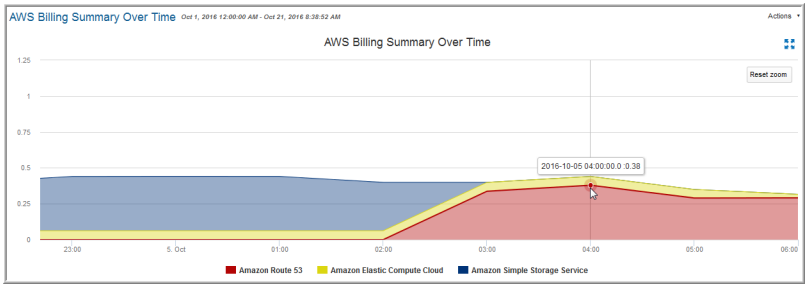
View a graphical representation of Amazon Web Services (AWS) total costs over time.

This chart breaks out the various cloud services. Click an area of the chart to access the AWS Account Billing Details. Hover your mouse over data points on the line to view costs for a specific date and time.

The report template uses following data to render the chart.

Table 30-3 AWS Billing Summary Over Time

Column name	Description
Product	Long name of the AWS product.
Blended Cost	Cost for the operation listed in the line item based on the blended rate.
Usage End Date	The end of the applicable charge period. All usage charges and recurring fees between the period start and end are included in your monthly invoice.



AWS Product Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > AWS Product Summary

View a list of all discovered Amazon Web Services (AWS) products with associated month-to-date costs.

Access the relevant AWS Account Billing Detail report for additional details.

Use the Custom Filter in the scope selector to isolate specific types of AWS data.

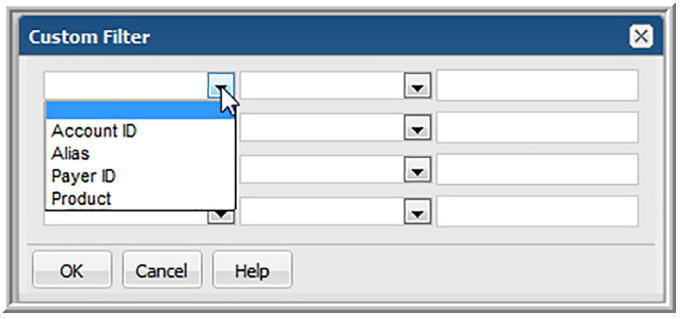


Table 30-4 AWS Product Summary

Column name	Description
Product	Long name of the AWS product.
Account ID	The AWS account ID, a 12-digit number.

Table 30-4 AWS Product Summary *(continued)*

Column name	Description
Payer ID	The AWS account ID, a 12-digit number.
Alias	The account alias.
Blended Cost	Cost for the operation listed in the line item based on the blended rate. Click the link to access the AWS Account Billing Detail.
Last Invoice Date	Date of the latest Usage Start Date that was collected for this month.
Last Updated	Last Updated Date

AWS S3 Bucket Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > AWS S3 Bucket Summary

Access an at-a-glance view of Amazon Web Services (AWS) Simple Storage Service (S3) buckets and a summary of the costs associated with an account.

This summary of S3 buckets includes data such as used capacity, month-to-date cost, and the trend of unblended costs associated with a bucket.

Use the Custom Filter in the scope selector to isolate specific types of AWS data.

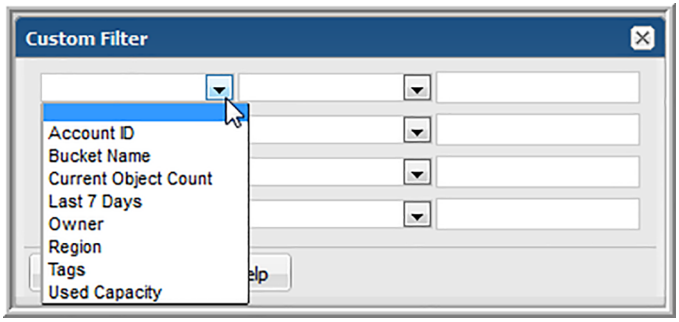


Table 30-5 AWS S3 Bucket Summary

Column name	Description
Bucket Name	S3 bucket name.
Account ID	The AWS account ID, a 12-digit number.
Region	Specifies the region in which an AWS object exists.
Zone	Specifies the Availability Zone in which an AWS object resides or a charged usage occurred.
Owner	S3 bucket owner name.
Current Object Count	Number of current objects in the S3 bucket.
Used Capacity	Used capacity of current objects in the S3 bucket.
Last Updated	Last Updated Date
Tags	Tags associated with the account.
Blended Cost	Cost for the operation listed in the line item based on the blended rate.
Last 7 Days	The sparkline chart represents the cost over the seven days prior to the end date of the report's scope.

AWS EC2 Snapshot

This report displays Amazon Web Services (AWS) EC2 instances with the status of snapshot backups. The main report is located here:

Public Cloud Reports > AWS > AWS EC2 Snapshot

Each line item displays the following information in a tabular format:

Table 30-6 AWS EC2 Snapshot

Column name	Description
Instance name:	AWS Instance name.
Last Snap Status:	Status of last snapshot. Red icon indicates no snapshot is available or more than one day has passed since a snapshot was created. Green icon is successful snapshot creation.

Table 30-6 AWS EC2 Snapshot *(continued)*

Column name	Description
Days Since Last Snap:	Total number of days since last snapshot occurred.
Last Snap Date:	Date and time of last snapshot creation for the instance.
Last Snap Size:	Size of the last snapshot recorded.
Total Snaps:	Total number of snapshots recorded.
Total Size:	Total size of all the snapshots recorded.

AWS EC2 Snapshot

EC2 Instance=All EC2 Instances
05-Feb-2021 11:24:54

Total Rows:45

Instance Name	Last Snap Status	Days Since Last Snap	Last Snap Date	Last Snap Size	Total Snaps	Total Size
OraLinInstallTestQA1				N/A		N/A
OraLinInstallTestQA2				N/A		N/A
QAInstallWin2016d				N/A		N/A
QAInstallWin2016g				N/A		N/A
QaSuse12InstallTest				N/A		N/A
scprod_clone				N/A		N/A
svnprod_				N/A		N/A
tmpFATest2				N/A		N/A
tmpOraWinBuild				N/A		N/A
TmpTestHostname				N/A		N/A
tmpWikiMigration				N/A		N/A
Win2012OraPatchTest				N/A		N/A
WinDevTest				N/A		N/A
WOW104				N/A		N/A
NBU 8.1.4 Beta1 (nbu80)		1,243	11-Sep-2017 07:23:30	100.00 GiB	1	100.00 GiB
QAPerfTesting		797	30-Nov-2018 23:58:47	110.00 GiB	2	110.00 GiB
Win2012_120_8GB		700	07-Mar-2019 23:30:30	120.00 GiB	1	120.00 GiB
PenTestCollectorLin		283	28-Apr-2020 07:06:26	40.00 GiB	2	80.00 GiB
PenTestCollectorWin		283	28-Apr-2020 07:05:30	80.00 GiB	2	160.00 GiB
PenTestPortal		283	28-Apr-2020 07:05:58	120.00 GiB	3	360.00 GiB
aptarefin3		1	04-Feb-2021 05:02:43	160.00 GiB	15	2.34 TiB
AWSCloud		1	04-Feb-2021 05:02:44	100.00 GiB	30	1.46 TiB
bugzillaAWS		1	04-Feb-2021 05:02:42	16.00 GiB	15	240.00 GiB
DemoAWS		1	04-Feb-2021 05:02:43	100.00 GiB	16	1.56 TiB
jira-prod		1	04-Feb-2021 05:02:42	60.00 GiB	23	1.35 TiB
wowAWS		1	04-Feb-2021 05:02:45	240.00 GiB	20	4.47 TiB

AWS EC2 Snapshot Detail

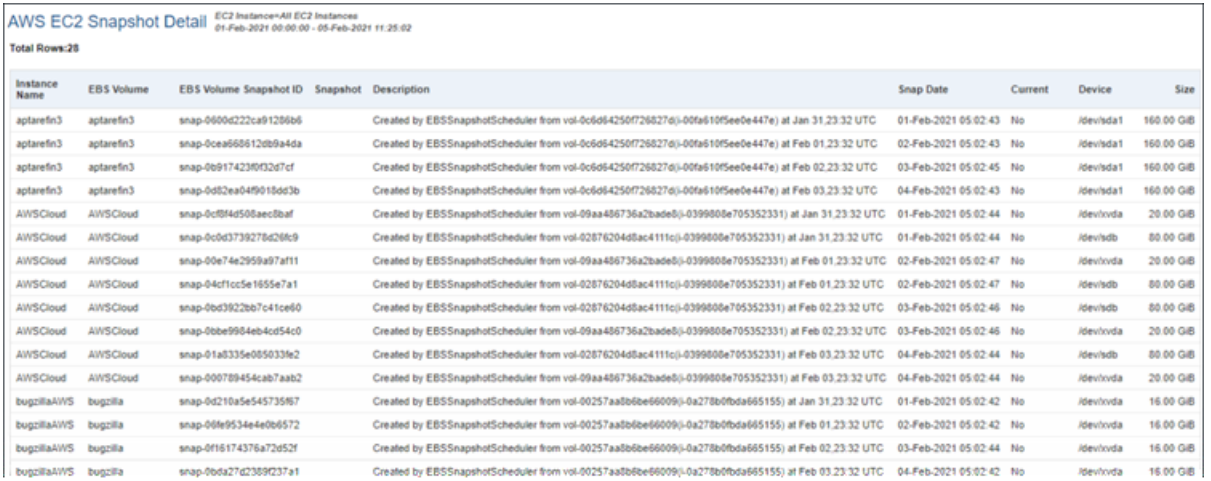
This report displays Amazon Web Services (AWS) EC2 instance snapshot backup details including volume info, snapshots dates and sizes. The main report is located here:

Public Cloud Reports > AWS > AWS EC2 Snapshot Detail

Each line item displays the following information in a tabular format:

Table 30-7 AWS EC2 Snapshot Detail

Column name	Description
Instance Name:	AWS Instance Name.
EBS Volume:	AWS Elastic Block Storage (EBS) volume name.
EBS Volume Snapshot ID:	AWS Elastic Block Storage (EBS) volume snapshot ID.
Snapshot:	AWS Snapshot Name.
Description:	Description of the AWS snapshot.
Snap Date:	Date and time when the snapshot creation was started.
Current:	Indicates if the EBS volume is currently at this snapshot.
Device:	Device Name.
Size:	Capacity of AWS Elastic Block Storage (EBS) volume.



AWS EC2 Snapshot Detail *EC2 Instance=All EC2 Instances 01-Feb-2021 00:00:00 - 05-Feb-2021 11:25:02*

Total Rows:28

Instance Name	EBS Volume	EBS Volume Snapshot ID	Snapshot	Description	Snap Date	Current	Device	Size
aptarefin3	aptarefin3	snap-0600d222ca91286b6		Created by EBSSnapshotScheduler from vol-0c6d64250f726827d(i-00fa610f5ee0e447e) at Jan 31,23:32 UTC	01-Feb-2021 05:02:43	No	/dev/sda1	160.00 GiB
aptarefin3	aptarefin3	snap-0cea668612db9a4da		Created by EBSSnapshotScheduler from vol-0c6d64250f726827d(i-00fa610f5ee0e447e) at Feb 01,23:32 UTC	02-Feb-2021 05:02:43	No	/dev/sda1	160.00 GiB
aptarefin3	aptarefin3	snap-0b917423f0f32d7cf		Created by EBSSnapshotScheduler from vol-0c6d64250f726827d(i-00fa610f5ee0e447e) at Feb 02,23:32 UTC	03-Feb-2021 05:02:45	No	/dev/sda1	160.00 GiB
aptarefin3	aptarefin3	snap-0d82ea04f9018dd3b		Created by EBSSnapshotScheduler from vol-0c6d64250f726827d(i-00fa610f5ee0e447e) at Feb 03,23:32 UTC	04-Feb-2021 05:02:43	No	/dev/sda1	160.00 GiB
AWSCloud	AWSCloud	snap-0cf8f4d508aec8baf		Created by EBSSnapshotScheduler from vol-09aa486736a2bade8(i-0399808e705352331) at Jan 31,23:32 UTC	01-Feb-2021 05:02:44	No	/dev/xvda	20.00 GiB
AWSCloud	AWSCloud	snap-0c0d3739278d26fc9		Created by EBSSnapshotScheduler from vol-02876204d8ac4111c(i-0399808e705352331) at Jan 31,23:32 UTC	01-Feb-2021 05:02:44	No	/dev/sdb	80.00 GiB
AWSCloud	AWSCloud	snap-00e74e2959a97af11		Created by EBSSnapshotScheduler from vol-09aa486736a2bade8(i-0399808e705352331) at Feb 01,23:32 UTC	02-Feb-2021 05:02:47	No	/dev/xvda	20.00 GiB
AWSCloud	AWSCloud	snap-04cf1cc5e1655e7a1		Created by EBSSnapshotScheduler from vol-02876204d8ac4111c(i-0399808e705352331) at Feb 01,23:32 UTC	02-Feb-2021 05:02:47	No	/dev/sdb	80.00 GiB
AWSCloud	AWSCloud	snap-0bd3922bb7c41ce60		Created by EBSSnapshotScheduler from vol-02876204d8ac4111c(i-0399808e705352331) at Feb 02,23:32 UTC	03-Feb-2021 05:02:46	No	/dev/sdb	80.00 GiB
AWSCloud	AWSCloud	snap-0bbe9984eb4cd54c0		Created by EBSSnapshotScheduler from vol-09aa486736a2bade8(i-0399808e705352331) at Feb 02,23:32 UTC	03-Feb-2021 05:02:46	No	/dev/xvda	20.00 GiB
AWSCloud	AWSCloud	snap-01a8335e0850338e2		Created by EBSSnapshotScheduler from vol-02876204d8ac4111c(i-0399808e705352331) at Feb 03,23:32 UTC	04-Feb-2021 05:02:44	No	/dev/sdb	80.00 GiB
AWSCloud	AWSCloud	snap-000789454cab7aab2		Created by EBSSnapshotScheduler from vol-09aa486736a2bade8(i-0399808e705352331) at Feb 03,23:32 UTC	04-Feb-2021 05:02:44	No	/dev/xvda	20.00 GiB
bugzillaAWS	bugzilla	snap-0d210a5e545735f67		Created by EBSSnapshotScheduler from vol-00257aa8b6be66009(i-0a278b0fbda665155) at Jan 31,23:32 UTC	01-Feb-2021 05:02:42	No	/dev/xvda	16.00 GiB
bugzillaAWS	bugzilla	snap-06fe9534e4e0b6572		Created by EBSSnapshotScheduler from vol-00257aa8b6be66009(i-0a278b0fbda665155) at Feb 01,23:32 UTC	02-Feb-2021 05:02:42	No	/dev/xvda	16.00 GiB
bugzillaAWS	bugzilla	snap-0f16174376a72d52f		Created by EBSSnapshotScheduler from vol-00257aa8b6be66009(i-0a278b0fbda665155) at Feb 02,23:32 UTC	03-Feb-2021 05:02:44	No	/dev/xvda	16.00 GiB
bugzillaAWS	bugzilla	snap-0bda27d2389f237a1		Created by EBSSnapshotScheduler from vol-00257aa8b6be66009(i-0a278b0fbda665155) at Feb 03,23:32 UTC	04-Feb-2021 05:02:42	No	/dev/xvda	16.00 GiB

AWS EC2 Instance Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Public Cloud Reports> AWS > AWS EC2 Instance Summary

Access an at-a-glance view of each Amazon Web Services (AWS) Elastic Cloud Compute (EC2) instance.

Each instance includes relevant data, such as Elastic Block Store (EBS) usage and the current month-to-date cost.

Use the Custom Filter in the scope selector to isolate specific types of AWS data.

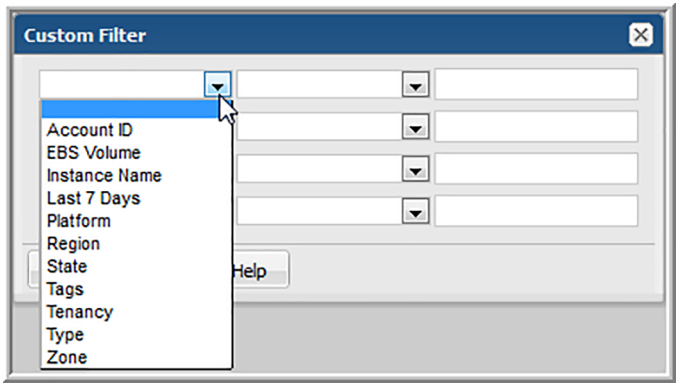


Table 30-8 AWS EC2 Instance Summary

Column name	Description
Instance Name	AWS Instance Name.
Host Name	Name of the host to display in Portal reports.
Account ID	The AWS account ID, a 12-digit number.
Region	Specifies the region in which an AWS object exists.
Zone	Specifies the Availability Zone in which an AWS object resides or a charged usage occurred.
Type	Type of instance (determines the instance CPU, capacity, memory, and storage).
Platform	Platform of the guest (blank if not Windows).
Tenancy	The EC2 instance tenancy type, such as dedicated or shared.
EBS Volume	EBS Volume
Tags	Tags associated with the instance.
Launched	When the instance was launched.
Last Updated	Last updated date.
State	Instance state, such as stopping or running.
Blended Cost	Cost for the operation listed in the line item based on the blended rate.
Last 7 Days	The sparkline chart represents the cost over the seven days prior to the end date of the report's scope.

Microsoft Azure Reports

Microsoft Azure reports list:

- Azure Cloud Cost Spend Dashboard
- Cost History By Region - Top N
- Cost History by Service (Meter Name/Category) - Top N
- Cost History by Subscription - Top N
- Cost History by Tags
- Microsoft Azure Virtual Machine Summary
- Microsoft Azure Virtual Machine Detail
- Microsoft Azure Billing Summary
- Microsoft Azure Resource Group Summary
- Microsoft Azure Subscription Summary

Azure Cloud Cost Spend Dashboard

This dashboard shows the cost spent on Azure cloud subscription. This dashboard is located here:

Public Cloud Reports > **Azure** > **Azure Cloud Cost Spend Dashboard**

The dashboard displays the spends as follows:

- Spend - Last Month: Overall cost spent in the previous month for all the subscriptions which are configured under Azure Data collection policy.
- Spend so far - Current Month: Cost spent from the first day of the current month up to the current date for all the subscriptions which are configured under Azure Data collection policy.
- Projected spend - Current Month: Projected cost for the current month and is calculated as (Cost spent in the current month till date/last collection date)*(Total no of days in the month)).

Apart form the calculated values above the dashboard also provides a single-window view of the following reports:

- Cost History By Region - Top N: This report shows the week-wise cost history of the top 10 high-spending regions for the last 90 days by default.
- Cost History by Service (Meter Name/Category) - Top N: This report shows the week-wise cost history of the top 10 high-spending services for the last 90 days by default..

- Cost History by Subscription - Top N: This report shows the week-wise cost history of the top 10 high-spending subscriptions for the last 90 days by default.
- Cost History by Tags: This report shows week-wise cost history of values under the tags for the last 90 days by default.

Note: If the exported report or dashboard contains a generic placeholder chart image, the chart export feature on your system has been disabled. See the Remove "PhantomJS" from NetBackup IT Analytics Portal tech note for further details.

Azure Daily Usage Future Forecast by Org Level

The Azure Daily Usage Future Forecast by Org Level with history and trend sparklines is a usage report that displays cost-based trending and statistical analysis grouped by organization level. Before you run this report, you must define the values of the organization level attributes Org_level1, Org_level2, and Org_level3 as they are used by this report. See *Load Subscription attributes and values* section in the *NetBackup IT Analytics System Administrator Guide* for the configuration information.

You can click the number of subscriptions value to drill down to a details report about the subscriptions of the specific organizational level that you select.

The report can be accessed from **Public Cloud Reports** > **Azure** > **Azure Daily Usage Future Forecast by Org Level**.

Table 30-9 Azure Daily Usage Future Forecast by Org Level

Column name	Description
Org Level1	The top level of your organizational structure. This could be a person's name or the parent division
Org Level2	The second level of your organizational structure.
Org Level3	The third level of your organizational structure.
#Subscriptions	The number of subscriptions associated with these three levels. This is a drill-down link to the subscription details report.
Usage History	A sparkline depiction of the daily storage cost values.
Linear trend	A sparkline depiction of the linear slope of the daily storage cost.
PTD Cost	The sum total of storage cost for the time period selected.

Table 30-9 Azure Daily Usage Future Forecast by Org Level *(continued)*

Column name	Description
Largest Day	The largest single day storage cost within the time selected period.
Avg Day	The average daily storage cost within the time period selected.
#Days	The number of days returned based on the date range selected. **Note:** #Days could be less than that depending on when the data collection started.
First Date	The first date found within the time period selected.
First Value	What your daily cost was on the first date of the period.
Last Date	The most recent date found within the time period selected.
Last Value	What your daily cost was on the last date of the period.
Daily Change Linear	The amount of cost change on a daily basis on a linear regression algorithm. (The slope is based on the best fit sum of squares.)
%Change Linear	The percentage of daily change
Cost Change	A verbal description of the cost change between the first date and the last date. This value can be used to filter, such as to show only where the cost "Increased".
Trend	A visual depiction of how the cost is trending. Represents the up, down, or flat trend.
Probability	The coefficient of determination (R^2) how well a statistical model predicts an outcome.
30 Day Forecast	Forecast of the PTD cost in the next 30 days based on the Daily Change Linear value.
6 Month Forecast	Forecast of the PTD cost in the next 6 months based on the Daily Change Linear value.

Table 30-9 Azure Daily Usage Future Forecast by Org Level *(continued)*

Column name	Description
1 Yr Forecast	Forecast of the PTD cost in the next year based on the Daily Change Linear value.

Azure Daily Usage Forecast by Subscription Details

Azure Daily Usage Forecast by Subscription Details plots the future cost per subscription. This report appears as drill down report of Azure Daily Usage Future Forecast by Org Level, when you click the value in the #Subscriptions column.

Table 30-10 Azure Daily Usage Forecast by Subscription Details

Column name	Description
#Subscriptions	The number of subscriptions associated with these three levels. This is a drill-down link to the subscription details report.
Usage History	A sparkline depiction of the daily storage cost values.
Linear trend	A sparkline depiction of the linear slope of the daily storage cost.
PTD Cost	The sum total of storage cost for the time period selected.
Largest Day	The largest single day storage cost within the time selected period.
Avg Day	The average daily storage cost within the time period selected.
#Days	The number of days returned based on the date range selected. **Note:** #Days could be less than that depending on when the data collection started.
First Date	The first date found within the time period selected.
First Value	What your daily cost was on the first date of the period.
Last Date	The most recent date found within the time period selected.

Table 30-10 Azure Daily Usage Forecast by Subscription Details *(continued)*

Column name	Description
Last Value	What your daily cost was on the last date of the period.
Daily Change Linear	The amount of cost change on a daily basis on a linear regression algorithm. (The slope is based on the best fit sum of squares.)
%Change Linear	The percentage of daily change
Cost Change	A verbal description of the cost change between the first date and the last date. This value can be used to filter, such as to show only where the cost "Increased".
Trend	A visual depiction of how the cost is trending. Represents the up, down, or flat trend.
Probability	The coefficient of determination (R^2) how well a statistical model predicts an outcome.
30 Day Forecast	Forecast of the PTD cost in the next 30 days based on the Daily Change Linear value.
6 Month Forecast	Forecast of the PTD cost in the next 6 months based on the Daily Change Linear value.
1 Yr Forecast	Forecast of the PTD cost in the next year based on the Daily Change Linear value.

Cost History By Region - Top N

This report shows the week-wise cost history in the form of a bar chart of the top 10 high-spending regions for the last 90 days by default. Its scope can be tweaked to obtain the required view. Each bar further drills down to its respective Azure Region Summary report. Use this report to obtain the region-wise spends on your Azure subscriptions. This report is located here:

Public Cloud Reports > **Azure** > **Cost History By Region - Top N**

Cost History by Service (Meter Name/Category) - Top N

This shows the week-wise cost history as a bar chart of the top 10 high-spending services for the last 90 days by default. Its scope can be tweaked to obtain the required view. Each bar further drills down to its respective Azure Service Summary

report. Use this report to obtain the service-wise spends on your Azure subscriptions. This report is located here:

Public Cloud Reports > **Azure** > **Cost History by Service (Meter Name/Category) - Top N**

Cost History by Subscription - Top N

This shows the week-wise cost history as a bar chart of the top 10 high-spending subscriptions for the last 90 days by default. Its scope can be tweaked to obtain the required view. Each bar further drills down to display a summary report of the respective subscription. Use this report to view the spends on each of your Azure subscriptions.

This report is located here:

Public Cloud Reports > **Azure** > **Cost History by Subscription - Top N**

Cost History by Tags

This report shows week-wise cost history as a bar chart for values under the tags for the last 90 days. The default value of the tag for this report under this dashboard is always the first value in the drop-down list. This report can be empty if the resource usage data is not available for the default tag.

This report is located here:

Public Cloud Reports > **Azure** > **Cost History by Tags**

Note: This cost report does not take into consideration the discounts that Microsoft might be offering for a given tenant/subscription.

Microsoft Azure Virtual Machine Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Virtualization Manager > Administration Reports > Microsoft Azure Virtual Machine Summary

View a list of all discovered Microsoft Azure virtual machines (VMs) with associated month-to-date costs.

This VM overview provides subscriptions with related VM location, status, and last backup information. Access the Microsoft Virtual Machine Detail report for this report.

Table 30-11 Microsoft Azure Virtual Machine Summary

Column name	Description
VM Name	Virtual Machine name.
Virtual Machine ID	Unique Azure identifier of the virtual machine.
Account Administrator Email	The email address for the administrator of the Microsoft Azure Cloud account.
Subscription Name	Subscription display name.
Resource Group Name	The Resource Group name.
Operating System	The operating system profile for the VM.
State	Virtual Machine provisioning state.Possible values are Succeeded,Failed.
Location	Virtual Machine location.
VM Size Profile	The virtual machine size name. For example, Basic_A0, Basic_A1,Standard_A3.
# of Disks	Number of associated VM Disks. Click this link to view the Microsoft Azure Virtual Machine Disk Detail.
Total Disk Size	The initial disk size, for blank data disks, and the new desired size for resizing existing OS and data disks.
IP Address	The private IP address.
Last Backup Date	Last Backup Date.
Last Updated	Last Updated Date.

Microsoft Azure Virtual Machine Detail

Detail reports are related to a specific enterprise object, such as a backup job. You can only access detail reports through a link presented in the context of a main report, providing additional information that augments the main report. Detail reports cannot be generated, customized, or saved, as they are specific to the report from which they were derived. They are not available in search results.

Explore your data center using NetBackup IT Analytics customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the NetBackup IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. The main report is located here:

Virtualization Manager > Administration Reports > Microsoft Azure Virtual Machine Summary

Click the **# of VM Disks** link.

View Azure Virtual Machine details for a storage accounts and resource groups. Use this report to identify associated disk and storage details.

Table 30-12 Microsoft Azure Virtual Machine Detail

Column name	Description
Virtual Machine	Virtual Machine name.
Subscription	Subscription display name.
Resource Group	The Resource Group name.
Storage Account	Name of the Storage Account.
Disk	The VM Disk name
Disk Size	The initial disk size, for blank data disks, and the new desired size for resizing existing OS and data disks.
LUN	Logical unit number.
Creation Option	Virtual Machine creation option. Possible values: fromImage, empty, attach.
Caching	Caching type. Possible values: None, ReadOnly, ReadWrite.

Microsoft Azure Billing Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Microsoft Azure Billing Summary

View Microsoft Azure billing costs associated with an account.

This summary of accounts and subscriptions includes associated meters, usage, and applicable rates and costs.

Table 30-13 Microsoft Azure Billing Summary

Column name	Description
Account	The Account Administrator Email Address of Microsoft Azure Cloud account.
Subscription Name	The subscription display name.
Meter Name	The name of the meter, within the given meter category.
Meter Category	The category of the meter, e.g., Cloud services, Networking, etc..
Meter Subcategory	The subcategory of the meter, e.g., A6 Cloud services, ExpressRoute (IXP), etc..
Unit	The unit in which the meter consumption is charged, e.g., Hours, GB, etc.
Region	The region in which the Azure service is available.
Service	The unique identifier of the resource
Service Type	The resource type - Microsoft.Compute or Microsoft.Storage.
Usage End	UTC end time for the usage bucket to which this usage aggregate belongs.
Quantity	The amount of the resource consumption that occurred in this time frame.

Microsoft Azure Resource Group Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Microsoft Azure Resource Group Summary

View a list of all discovered Microsoft Azure resource groups and products.

The list of resources includes types such as, storage accounts, IP addresses, VMs, NICs, Vaults, and Security Groups.

Table 30-14 Microsoft Azure Resource Group Summary

Column name	Description
Account	The Account Administrator Email Address of the Microsoft Azure Cloud account.
Subscription Name	The subscription display name.
Subscription ID	The subscription ID of Microsoft Azure Cloud account.
Group Name	The resource group name.
Resource ID	The resource group URI.
Resource Name	The unique identifier of the resource.
Resource Type	The resource type - Microsoft.Compute or Microsoft.Storage.
Resource Location	Resource Group location.
Resource Tags	Tags.
Last Updated	Last Updated Date.

Microsoft Azure Subscription Summary

Explore your data center using customizable report templates or by using parts of your IT infrastructure as entry points. Use Search to find reports, templates and dashboards across the portal.

Use the Reports tab to examine the catalog of templates, dashboards and reports - organized by products along with user-created, and system folders. This report is located here:

Capacity Manager > Chargeback and Billing > Microsoft Azure Subscription Summary

View a list of all discovered Microsoft Azure accounts with associated services and month-to-date costs.

A summary of allocated cores, cloud services, storage accounts, and network sites provides an overview of Azure subscriptions.

Table 30-15 Microsoft Azure Subscription Summary

Column name	Description
Account	The Account Administrator Email Address of Microsoft Azure Cloud account.
Subscription ID	The subscription ID of Microsoft Azure Cloud account.
Name	The subscription display name.
Status	The subscription state. Possible values are Enabled, Warned, PastDue, Disabled, and Deleted.
Admin	The Account Administrator ID of Microsoft Azure Cloud account.
Current Storage Accounts	Specifies the number of currently allocated storage accounts.
Current Virtual Machines	Specifies the number of currently allocated virtual machines.

Microsoft Azure Unattached Disks Summary

This report displays the summarized view of all the unattached disks and its corresponding subscription, creation date, and cost implications.

Navigate to *Public Cloud Reports* > *Azure* then click *Microsoft Azure Unattached Disks Summary*.

Table 30-16 Microsoft Azure Unattached Disks Summary

Column name	Description
Subscription name	Name of the Subscription under which the corresponding VM Storage is present.
Subscription ID	ID of the subscription.
Unattached Disk Name	Name of the unattached VM storage disk.
Storage type	Type of the VM storage. **Note:** Default is **Premium SSD LRS**, if not assigned.
Creation date	Creation date of the VM storage disk.
Cost implication	Cost implication. **Note:** Calculated in Dollar currency value on the basis of the price of the resource usage started under a given date range.