IT Analytics User Guide

Release 11.8

COHESITY

IT Analytics User Guide

Last updated: 2026-10-01

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

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.

   Note: 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.

Contents

Chapter 1 Introduction 25

Overview 25
Licensed modules 25
IT Analytics components 27

Chapter 2 Understand the Portal 29

Overview 30
Portal access 30
Portal overview 33
Navigating with search 35
Valid entries for search 37
Search results view: Narrowing the scope 38
Filtering data in management grids 39
Advanced filtering 40
Customizing advanced filter logic 42
Saving advanced filters (Inventory) 42
Advanced filtering examples 44
Advanced filter operators 46
Clear advanced filters 47
Delete saved advanced filters 48
Select columns on management pages 48
Detect alerts in the inventory 51
About the Inventory tab 52
About the Reports tab 52
Reports tab keyboard shortcuts 53
Reports tab basics 53
Templates, reports, and dashboards overview 54
About the Alerts tab 55
About the Admin tab 55
- Users 55
- Domains 56
- Chargeback 56
- Solutions 56
- Data Collection 56
- Reports 57

Advanced 57

Chapter 3 Explore your inventory 59

Exploring your inventory 60
Inventory privileges 61
Getting started with the Inventory navigator 61
Hierarchy toolbar to organize your data 63
Attribute value limitations 65
Selection combinations 66
Inventory object limits 66
Show objects 66
Backup Servers in the Inventory 67
Count values 69
Reset inventory defaults 69
Host groups 70
Use attributes to organize your data 71
System attributes 71
Attribute value limitations 72
Use host groups to organize your data 73
Create custom object lists 75
Working with the inventory reports view 77
Accessing inventory reports 78
Choose reports to display 78
Filter within available reports by category 80
Pin reports - saving reports with inventory objects 81
Rename pinned tabs 82
Reset default report names on pinned tabs 82
Inventory reports and the action menu 83
Work with the inventory list view 84
Filter the inventory list view 85
Manage objects in the inventory list view 85
Assign attributes in the inventory list view 86
About attributes 86
Attribute limitations 87
Set attribute values in the inventory list view 88
Import host attribute values 90
Export objects from the inventory list view 93
Delete objects using the inventory list view 94
Customize columns in the inventory list view 95

Chapter 4 Get acquainted with reports 96

Overview 97
Reports and templates 97
Cloud reports and templates 98
Report library 98
Report orientation 99
Shared reports and dashboards 100
About dashboards 100
Work in the Reports tab 101
New badge 101
About caching 101
Sections in the navigation panel 102
About badging 102
    Home section 103
    Cloud section 103
About the action bar 104
Reports tab keyboard shortcuts 104
Folder types 104
Report template icons 105
Navigate through reports 106

Chapter 5 Generate and maintain reports 109

Report Scope 110
Select Report Scope 113
    Time Period Considerations 113
Configure report scope with attributes 116
Group hosts by attributes 119
    Host attributes examples 120
    Find hosts without an attribute value set 120
Search for hosts in the report Scope Selector 120
    Find hosts example 120
Alert scope selector settings 121
Amazon Web Services (AWS) scope selector settings 122
Backup Manager scope selector settings 123
Backup Manager advanced scope selector settings 126
    Examples of Advanced Options Job Types 127
Capacity Manager scope selector settings 128
Data Collection Status Details scope selector settings 129
Host scope selector settings 130
Ransomware Scorecard scope selector 131
Solution reports scope selector settings 132
    Cascade to Sub-Domains 132

Storage Optimization Scope Selector 133
Risk Mitigation Scope Selector 135
System health check scope selector 137
Generate reports 137
Generate reports from the Cloud 140
Units of Measure in Reports 140
Detail Reports 140
Advanced Filtering for Tabular Reports 141
Customize report filter logic 142
Report Filter Operators 143
Filtering Indicators, Multiple Expressions, Dates, Percentages, Yes/No 144
Sort columns in reports 147
To sort columns in a table report 147
Refresh and set refresh intervals 147
Modify Reports 148
Convert tabular report to chart 149
Convert a chart into a tabular report 153
Add and remove custom header and footer to a report or dashboard 153
Set or edit the drill down preference 155
Access NetBackup web interface from the IT Analytics Portal 156
View report statistics 158
Interact with Reports 159
Work with Topology 161
Add a note to a Backup Job Details Report 161
Use new and updated User Report Templates 161
Use new and updated System Report Templates (Super User only) 163
Save Reports 165
Save Cloud Reports 166
Delete Reports 167

Chapter 6 Distribute, share, schedule, and alert 168

Export reports and dashboards 169
PDF report format 170
Known Limitations for Donut Charts 170
Schedule export of reports and dashboards 171
To schedule an export of report or dashboard 172
Using the Expression Builder 176
Email reports and dashboards 177
Schedule email of reports and dashboards 180

Viewing and Managing Your Scheduled Reports 183
Schedule Types 184
Deleting Scheduled Reports and Dashboards 186
Setting Up Alerts for Tabular Reports 186
Configure an alert 187
Deleting Report-Based Alerts 189
Share Reports, Dashboards, and Folders 189
Working with Shared Reports 193

Chapter 7 Organize reports 195

Work with dashboards 196
Understand the dashboard Icons 196
About Cached Reports on Dashboards 197
Create new dashboards and add reports 197
Add reports to an existing dashboard 200
Save a dashboard with a different name 201
Customize a Dashboard 203
Edit the scope for a report on a dashboard 204
Edit the dashboard scope 206
Expand a report to full size on a dashboard 207
Set the dashboard layout 208
Delete a dashboard 209
Remove reports from dashboards 209
Create a custom report folder 209
Create shortcuts to reports and templates 210
Manage My Home Pages 211
User Home Pages and User Group Home Pages 212

Chapter 8 Work with the dynamic template designer 217

Dynamic Template Designer Overview 220
Dynamic Template Designer Cheat Sheet 221
Modify an Out-of-the-Box Dynamic Template 222
Dynamic Template Designer Quick Start 223
Quick Start Step 1: Create a Basic Table Dynamic Template 223
Quick Start Step 2: Format Fields 225
Quick Start Step 3: Create a Calculated Field 227
Quick Start Step 4: Configure Table Column Totals 227
Quick Start Step 5: Define the Scope Selector 228
Quick Start Step 6: Generate a Report and View the Automatically Generated Help 229
Steps to Create Dynamic Templates 230
Working with Enterprise Objects and Fields 231

Select an Enterprise Object 232
Configure General Dynamic Template Designer Components 233
Converting to a Homogeneous, Product-Specific Template 236
    Copy from All 237
Add Fields and Methods to a Dynamic Template 237
Alias Names for Fields 239
Conditions When Hidden Fields are Included by Default 239
Dynamic Template Field Configuration 239
Valid Formatter Patterns 248
Examples of Dynamic Templates Containing Graphics Tool Tips 251
Dynamic Template Drilldown Configuration 253
Custom Drilldowns and Examples 255
Drilldown Parameters 259
Examples of Dynamic Templates Containing Drilldowns 260
Using Groups in Dynamic Templates 262
Add a Group to Create Separate Line Charts in a Report 263
Add a Group to Create a Double Header in a Tabular Report 263
Dynamic Template Function Configurations 264
    Aggregation Functions 267
    Comparison Functions 267
    Numeric Functions 268
    Character Functions 268
    Date Functions 270
    Character String Returning Numeric Values 271
    Null Function 271
    Decode Function 271
    Unique Function 272
Examples of Dynamic Templates Containing Functions 272
Create User-Defined Fields with the Field Builder 275
Create Fields with the Field Builder 275
    Convert Dates to Numbers in the Field Builder 275
    Syntax for Calculated Fields 276
    Examples of Calculated Fields in a Dynamic Template 276
Configure a Static Filter 277
Filter on Date Fields 277
Static Filter vs. Tabular Report Filter 278
Configure the Field Sort Order 278
Return Unique Results 279
Dynamic Template Scope Selector Components 279
Scope Selector Component - Custom Filter 280
    Define a Custom Filter at Run Time 281
    Custom Filter Operators for Numeric Fields 282
    Custom Filter Operators for String Fields 283

Data Domain Enterprise Object Scope Selector Components 284
Host Enterprise Object Scope Selector Components 285
Job Enterprise Object Scope Selector Components 287
Storage Array Enterprise Object Scope Selector Components 289
Customize and Export Dynamic Templates 292
Export/Import Dynamic Templates Containing Custom Attributes 292
Export a dynamic template that contains attributes 293
Import a dynamic template that contains attributes 293
Save a Dynamic Template After Edits 294
Format the Dynamic Template Output 294
Configure a Bar Chart Dynamic Template 295
Best Practices 295
Steps to Create a Bar Chart Dynamic Template 296
Bar Chart Dynamic Template: Use Case 296
Bar Chart Dynamic Template Step 1: Create a Dynamic Template 296
Bar Chart Dynamic Template Step 2: Configure Basic Field Functionality 297
Bar Chart Dynamic Template Step 3: Configure Functions for a Field 299
Bar Chart Dynamic Template Step 4: Configure Drilldowns for a Field 300
Bar Chart Dynamic Template Step 5: Configure a Line Field 301
Bar Chart Dynamic Template Step 6: Define the Scope Selector 301
Bar Chart Dynamic Template Step 7: Save and Run the Report Template 301
Examples of Bar Chart Dynamic Templates 302
Configure an Area/Stacked Area Chart Dynamic Template 303
Area/Stacked Area Chart Dynamic Template Step 1: Create a Dynamic Template 303
Area/Stacked Area Chart Dynamic Template Step 2: Configure Basic Field Functionality 304
Area/Stacked Area Chart Dynamic Template Step 3: Configure Functions for a Field 306
Area/Stacked Area Chart Dynamic Template Step 4: Define the Scope Selector 308
Area/Stacked Area Chart Dynamic Template Step 5: Save and Run the Report Template 308
Configure a Donut Chart Dynamic Template 309
Example of a Donut Chart Dynamic Template 309
Configure a Horizontal Bar Chart Dynamic Template 310
Example of a Horizontal Bar Chart Dynamic Template 310

Configure a Horizontal Stacked Bar Chart Dynamic Template 311
Example of a Horizontal Stacked Bar Chart Dynamic Template 312
Configure a Line Chart Dynamic Template 312
Line Charts for Performance Metrics 313
Identify Array Performance Data Fields for a Template 313
Rules and Guidelines for Performance Line Charts 314
Steps to Create a Line Chart Dynamic Template 314
Line Chart Field Requirements 315
Line Chart: Caption Field 315
Line Chart: Line Field 316
Line Chart: Group by Object Field 317
Line Chart: Hidden Field 318
Line Chart: Exclude Field 318
One Object Per Line Chart, One or More Metrics Per Chart 318
Configure One Object Per Line Chart 319
Field Configuration for One Object Per Line Chart 319
Scope Selector Configuration for One Object Per Line Chart 320
Multiple Objects Per Line Chart, One Metric Per Chart 321
Field Configuration for Multiple Objects Per Line Chart, One Metric Per Chart 322
Format Line Chart Fields 322
Line Chart Scope Selections 322
Examples of Line Chart Dynamic Templates 323
Configure a Pie Chart Dynamic Template 324
Examples of Pie Chart Dynamic Templates 326
Configure a Stacked Bar Chart in a Dynamic Template 326
Example of a Stacked Bar Chart Dynamic Template 327
Configure a Line in a Bar Chart 327
Configure a Table Dynamic Template 328
Examples of Tabular Dynamic Templates 329
Create a Sparkline Chart in a Tabular Dynamic Template 330
Example of a Client Backup Sparkline Chart in a Dynamic Template 330
Example of Using a Method to Create a Sparkline Chart in a Dynamic Template 331
Client Failure Sparkline Method 331
Configure Chart Axes 332
Using Methods in Dynamic Templates 333
Overview of Method Creation 334
System Methods 334
Stand-Alone Methods 340
Viewing Methods 341
Adding or Editing Methods 342

User-Defined Methods 342
System Methods 343
Accessing Methods 343
Tokens for Methods 347
Exception Tables for a Method's Join Clause 347
Validate and Save a Method 348
Method Validation 348
Delete a Method 349
Troubleshooting Dynamic Templates 349

Chapter 9 Work with the SQL template designer 350

SQL template designer overview 351
Database Published Views 352
Access the Database Published Views from the CLI 352
Access the Database Published Views from the SQL Template Designer 352
Steps to Create SQL Templates 352
Create a SQL Template 353
Configure SQL Template Scope Selector Components 356
Construct the SQL Query 367
Help for Database Views 368
Accessing the Published Database Views from the CLI 369
Special Characters in SQL Queries 369
Sample SQL Queries 370
Use Functions in Queries 375
Format the SQL Template Output 376
Number, Size, Date, and Time Formatting 377
Alignment, Aggregation, Bar Type, and Bar Type Color 381
Area, Line or Column Sparkline Format for a SQL Tabular Report Template 383
Configure a Bar Chart SQL Template 384
Example of a Bar Chart SQL Template: Host CPU Performance 385
Configure a Line Chart SQL Template 386
Configure a Donut Chart SQL Template 388
Example of a Donut Chart SQL Template: NetBackup Job Size by Job Type 388
Configure a Gauge Chart SQL Template 389
Configure a Horizontal Bar Chart SQL Template 392
Example of a Horizontal Bar Chart SQL Template: Host CPU Performance 392
Configure a Horizontal Stacked Bar Chart SQL Template 393

Example of a Horizontal Stacked Bar Chart SQL Template: Available/Allocated Capacity 394
Configure an Area Chart SQL Template 395
Example of Area Chart SQL Template: Allocated Available 395
Configure a Pie Chart SQL Template 398
Example of a Pie Chart SQL Template: Overall Job Status Summary 399
Configure a Pivot Table SQL Template 401
Example of a Pivot Table SQL Template: NetBackup Job Size 401
Configure a Stacked Bar Chart SQL Template 403
Example of a Stacked Bar Chart SQL Template: Host CPU Performance 404
Configure a Table in the SQL Template Designer 405
Example of a Table SQL Template: Host Group Membership List 405
Example of a Table SQL Template: Exposed Clients 406
Save and Share Report Templates 407
Advanced SQL Report Template Options 408
Creating Drilldowns 408
Using SQL Templates for Drilldowns 411
Create mouse-over hovers 411
Add an External Reference 413
Inventory Report Configuration 416
About the Reports in the Inventory 417
Edit a Custom Report Template 418
Export/Import SQL Templates 418
Export a SQL Template 419
Import a SQL Template 419

Chapter 10 Pipelined functions for report query building 421

About Pipelined Functions 422
General Functions 423
APTlistOfDates 423
Example Query 424
aptStringConcat 425
Example Query 425
collectString 425
getLicenseClientDetail 426
getServerAttributeValue 426
Example Query 426
getObjectAttributeValue 427
Example Query 427

getChildServerGroupContextById 428
Example Query 428
getServerGroupContextById 429
Example Query 429
secsToHoursMinSecs 429
Example Query 430
Backup Manager Functions 430
APTgetJobTypeName 431
APTgetTapeDriveStatusName 432
Example Query 432
getFullPathname 433
Example Query 433
listJobSummaryAfterRestart 433
Example Query 434
listJobSummaryAfterRestartNBW 435
Example Query 436
listJobSummaryAfterRestart for NetWorker Backup Jobs 436
Example Query 437
listOfBackupWindowDates 438
Explicitly Supply the Backup Windows 438
Example Query 439
listOfBackupWindowDates (by Backup Window ID) 439
Policy Auditing Functions 441
listClientChanges 441
listPathnameChanges 442
listPolicyChanges 443
listScheduleChanges 444
Capacity Functions 445
listChargebackCatByVOLSDetail 445
Example Query 446
listChargebackCatByNcVolDetail 447
Example Query 448
listChargebackCatByFSDetail (for HNAS) 449
Example Query 449
listChargebackCatByFSDetail (for EMC Isilon) 450
Example Query 451
listChargebackByLUNSummary 451
Example Query 452
listChargebackByLUNDetail 453
Chargeback in Cluster and Virtualization Environments 453
Example Query 454
listChargebackCatByLUNSummary 455
Example Query 455

listChargebackCatByLUNDetail 456
Example Query 456

Chapter 11 Alert configuration 458

Overview 458
Alert rules 459
Scheduled Rules vs. Run Time Rules 459
Alert rules sample set 460
Managing alert policies 463
Add/Edit an Alert Policy 463
License Usage Alert Rule 477
ServiceNow: Job finalized alert rule 478
Scope Example: Data Protection: ServiceNow Jobs 479
Copying an Alert Policy 480
Deleting an Alert Policy 604
Managing alert notifications 481
Adding an Alert Notification Delivery Method 481
Deleting an Alert Notification Delivery Method 483
Report-Based alerts 483
Edit notifications for Report-Based Alerts 483
Run a Report-Based Alert On-Demand 486
Delete Report-Based Alert 487
Syslog configuration 487
About using SNMP 492
About SNMP 492
About SNMP versions 493
SNMP versions supported in IT Analytics 493
About the Management Information Base (MIB) and IT Analytics support 493
SNMP traps 494
Detect alerts in the inventory 498
Suppress alert notifications 498
View suppressed alerts 499

Chapter 12 Manage hosts, backup servers, and host groups 501

About hosts, backup servers, and host goups 501
Plan your host group hierarchy 502
Manage host groups 502
Add Host Groups and Sub-Groups 503
About Host Group Operations 503
Add Host Groups 504

Automated Tool to Create Host Groups 506
Host Group Membership 507
Rename Host Groups 511
Delete Host Groups 512
Move Host Groups 513
Find a Host Group ID 515
Assign a User Home Host Group 516
Manage Backup Servers 517
NetBackup Primary Servers 517
Configure the Time Zone for a NetBackup Primary Server 517
Default Host Group Hierarchy for Veritas NetBackup 518
About the Primary Servers Group 519
Identifying Inactive Clients (For Veritas NetBackup) 519
Managing hosts and backup servers 520
Adding and Editing Hosts and Backup Servers 521
Import Hosts 523
Decommission/Recommission Hosts and Backup Servers 526
Deleting Hosts 528
Manage host aliases 529
Bulk import host aliases 529
View and manage aliases for individual host 530

Chapter 13 Manage attributes and objects 535

About attributes 536
Attribute example 537
Host groups vs attributes 537
Bulk load utilities and attributes 538
System attributes 538
Attribute inheritance 539
Manage attributes 540
Set attributes on hosts 540
Add attributes 541
Edit or rename attributes 544
Enable granular user access using authorization attributes 546
Delete attributes 549
Attribute naming rules 550
Examples of attributes and values 551
Understanding the attribute management view 552
Attribute management view (Top-Level domain) 552
Attribute management view (Sub-Domain) 553
Export/Import dynamic templates with custom attributes 554
About object maintenance 554

Product specific objects for reporting 555
Maintaining objects 556
Assign attributes using object maintenance 557
Search for objects by type in object maintenance 560
Customize library objects 561
Permanently remove devices, libraries, and drives 561

Chapter 14 Provide Portal access and user privileges 563

Providing user access to the portal 564
Creating portal user accounts 565
About user types 565
Selecting user groups for users 568
Working with user groups 569
Creating user groups 570
Assigning users to user groups 571
About user privileges 572
Assigning user privileges 578
Setting user group privileges 581
Enabling new product report templates 586
Importing reports 588
Enabling cloud privileges 590
Granting access to template designers 591
Enabling all privileges in a category 592
Managing users and user groups 593
Searching for users and user groups 593
Editing user accounts 593
Impersonating user accounts 595
Setting / Resetting passwords 597
Password notes 598
Account lockout 598
Managing user group home pages (Administrator) 599
My Home pages and user group home pages (Administrator) 599
Removing portal users 604
Deactivating user accounts 604

Chapter 15 Configure primary schedules and backup windows 605

About primary schedules 605
Configure primary schedules 606
About custom backup windows 607
Custom backup window requirements 607

Configure custom backup windows 608

Chapter 16 Add, edit, and move policies 610

Add/Edit a threshold policy 610
Threshold Example 611
Add/Edit a capacity billing and usage policy 612
Capacity Chargeback policy types 613
Example of Capacity Billing and Usage Policy Logic 617
Drive Types for Capacity Chargeback 618
Add/Edit a backup SLA policy 620
Add/Edit a SAN Fabric Chargeback policy 622
IT Analytics Billing and Chargeback policies 623
Add/Edit a Billing and Chargeback policy 623

Chapter 17 Solutions administration 625

Storage Optimization solution overview 625
Understand the value of Storage Optimization historical data 626
Configure Storage Optimization rules 627
Edit Storage Optimization Rule Example 636
Common storage optimization parameters 637
Enable Storage Optimization rules 638
Storage optimization rule prerequisites and logic 641
Risk Mitigation solution overview 646
Configure Risk Mitigation rules 647

Chapter 18 Manage and monitor data collection 653

Data collection overview 654
Data collection component configuration 654
About data collection tasks 655
Backup collection tasks 655
Backup event data collector 655
Data collection installation summary 656
Update the Local Hosts file for data collection 656
Data Collector security and data encryption 657
Manage Data Collectors and collection policies 657
About validation and status 658
Add/Edit Data Collectors 658
Generate key file to configure Data Collector v11.2.xx or lower 665
Data Collector encryption 666
Enable/Disable data collectors 667

Enable and disable data collection policy schedules 668
Review collectors and collection status 669
Known Limitations for Policy States in Collector Administration 669
Deleting a data collector 672
Upgrade Data Collectors 673
Upgrade aptare.jar 673
Update Upgrade Manager 673
Monitoring data collection status 674
Organize the collection status view 674
Quick Filters 674
View data collection status 675
Troubleshoot data collection status 676
Use reports to monitor data collection status 677

Chapter 19 Work with Capacity Manager host data collection 678

Understand the host data collection process 679
Host resources prerequisites and configurations 680
Host Access Privileges, Sudo Commands, Ports, and WMI Proxy Requirements 680
Access Requirements by OS 681
Host access requirements 681
For Linux Hosts Only 682
For Windows Hosts Only 682
Command path verification 683
Host discovery and collection configuration steps 684
Host discovery and collection setup overview 684
Host discovery and collection maintenance overview 686
Before discovering hosts 687
Configure/Search 687
Manage credentials 688
Example of credentials for Windows hosts 689
Manage WMI Proxy 691
Manage paths 694
Manage access control 696
Host management 700
Configure host discovery policies to populate the host discovery and collection view 701
Discovery policy considerations 702
Configure a Discovery policy 702
Execute and monitor host discovery 705

Discovery processes 706
Method 1 707
Method 2 707
Validate host connectivity 708
Validate hosts 708
Validation history 709
Show errors 710
Filter the host discovery and collection window - Hide/Unhide, remove 711
Search and export in host discovery and collection 712
Basic Search 712
Pre-Defined Search 712
Advanced Parameters 713
Export in host discovery and collection 714
Configure and edit host probes 714
Propagate probe settings: Copy probes, paste probes 717
Example of Probe Copy/Paste 717
Probe Settings 718

Chapter 20 Discovery policies for Veritas NetBackup 722

Task overview: Configure and monitor discovery policies 722
Discovery policies overview 724
About Discovery types 724
Client drive discovery 724
Media server disk discovery 725
Backup policy coverage 725
Activate a discovery license 725
Exclude devices from discovery policies 726
Activate discovery probes in the NetBackup Data Collector policy 726
Monitor discovery processes 726
View client protection status 727

Chapter 21 View and manage system notifications 728

Overview 728
Viewing system alerts 728
Defer or suppressing notifications 729

Chapter 22 Customize with advanced parameters 731

Overview of advanced parameters 731
Use cases for advanced parameters 732
Domain/Collector 733

Operating system 733
Host 733
Commands 734
Adding an advanced parameter 734
Access control advanced parameters 736
ACCESS_CONTROL_FALLBACK_STRINGS 736
ACCESS_CONTROL_PASSWORD_PROMPT 737
ACCESS_CONTROL_SKIP_COMMAND_OUTPUT 737
ACCESS_CONTROL_TTY_REQUIRED 737
ACCESS_CONTROL_TTY_REQUIRED_ERROR 738
ACCESS_CONTROL_TTY_REQUIRED_NUM_LINES_SKIP 738
General Data Collection advanced parameters 738
DC_START_CONNECT_RETRIES 739
DC_UPGRADE_HANDLE_UTILITY_PATH 739
DC_UPGRADE_KILL_FILE_LOCKING_PROCESSES 739
COMMAND_TIMEOUT 740
DCLOGS_CLEANUP_LIMIT_IN_MONTHS 740
DATARCVR_READ_TIMEOUT 740
LOGGING_LEVEL 741
METACOLLECTOR_MAX_THREADS 741
METACOLLECTOR_MIN_THREADS 741
PATH_CONTROL_PARAM 742
RAWDATA_CLEANUP_LIMIT_IN_MONTHS 742
SSH_ALTERNATE_PKG_USE 742
SSH_CHANNEL_WAIT_TIME 742
SSH_CIPHERS 743
SSH_DEBUG 743
SSH_PORT 743
STDOUT_ERROR 743
STDOUT_FILTER_LIST 744
THREAD_IDLE_TIMEOUT 744
WMI_PROXY_VERSION_TIMEOUT 745
HISTORIC_TIME_INTERVAL_HOURS 745
Cloud data collection advanced parameters 745
Amazon Web Services (AWS) 746
Microsoft Azure 746
Host discovery and collection advanced parameters 748
HOST_COLLECTION_ORACLE_CONN_STR 748
APP_DISK_PATH 748
HBAVERIFY_NUM_DELAY_SECS 749
HOST_CONNECTION_PROTOCOL 749
HOST_CONNECTION_TIMEOUT 749

HOST_SOCKET_TIMEOUT 749
IGNORE_FILE_SYSTEMS 750
IS_EMULEX_MAPPING_STRICT_CHECKING 750
NUM_HBAVERIFY_RUN 751
STRICT_HOST_KEY_CHECKING 751
USE_VIO_RESTRICTED_SHELL 751
Backup Manager advanced parameters 752
COMMVAULT_OLDEST_JOB_HOURS 755
MMINFO_MOVE_BACKWARD_MIN 755
NBU_AUDIT_LOOKBACK_DAYS 755
NBU_DUP_JOB_PARTIAL_STATUS_OVERRIDE 755
NO_AVAMAR_CLI 756
RMAN_BACKUP_LOOKBACK_DAYS 756
RMAN_BACKUP_LOOKBACK_OVERRIDE 756
SUPPRESS_KERBEROS_PROMPT 757
USE_ALT_NBU_INCL_EXCL 757
USE_NTML_V2 757
VEEAM_BACKUP_LOOKBACK_DAYS 757
VEEAM_BACKUP_LOOKBACK_DAYS_OVERRIDE 758
Capacity Manager advanced parameters 758
EMC_MULTIPATH_OFF 759
EMC_VMAX_INCLUDE_ALL_MASKINGVIEWS 759
HDS_HTM_JPCCTRL_PATH 759
REMOTE_RAID_AGENT 760
SYMSTAT_COUNT 760
SYMSTAT_INTERVAL_SECONDS 761
File Analytics advanced parameters 761
FA_CLEAN_UP_TEMPORARY_FILES 761
AFS_PATH_EXCLUDE_DIRECTORY_NAME 762
FA_HOST_CAPTURE_REMOTE_SHARES 762
FA_HOST_DIRECTIVE 762
Virtualization Manager advanced parameters 763
VMWARE_COLLECT_UNUSED_PROPERTIES 763
VMWARE_DATASTORE_TIMEOUT_SECONDS 763
Fabric Manager Advance Parameters 764

Chapter 23 Manage your Portal environment 765

Root folder and domains 765
Root folder 765
Domain 766
Multiple domains (for managed services partners only) 766
Add/Configure a domain 766

Configure alerting for a domain 768
Community performance statistics for a domain 769
Configure the domain password policy 770
Manage your profile and set a language preference 772
Change your login password 774
View Downloads 774

Chapter 24 Manage ransomware scorecard 775

Ransomware Scorecard overview 775
Ransomware administration 785
Answer a ransomware query 785
Override a data query 786
Add a ransomware question 787
Utility to add Ransomware Scorecard questions and answers in bulk 789
Add Ransomware Scorecard questions in bulk 789
Add answers in bulk to Ransomware Scorecard questions 792
Edit a ransomware question 795
Add ransomware data query 796
Edit ransomware data query 798

Chapter 25 Analyze files 799

Overview 799
File categories 799
Adding/Editing file categories 805
File list export 806
File list export output in CSV format 808

Chapter 26 Troubleshoot the Portal 810

Support tools 810
Requesting a support package 812
File selector wildcards 817
Retrieving log files 819
On a Windows platform 819
On a Linux platform 819
Archiving script 819
Debug 820
Debugging Level Scope 820
Login errors 823
Add notes to failed jobs 823
Host name processing - Filters and aliases 824

Change the Portal server's IP address 824
Changing the port for Portal server communication 825
Determine report rendering statistics 826
Clearing the inventory and report cache 826
Troubleshoot Oracle performance issues in IT Analytics 826
Troubleshoot test connection locale issue 828

Chapter 27 Attribute inheritance overrides 836

Overview 836
Override inherited attribute values 837
Use case 1 838
Use case 2 838

Chapter 28 Understanding report data caching 839

Understand report data caching 839
How reports and caching work together? 839
Store and purge data 840
Disable optimistic caching with portal.properties 841

Chapter 1

Introduction

This chapter includes the following topics:

- Overview
- Licensed modules
- IT Analytics components

Overview

IT Analytics is the open and extensible storage resource management platform designed to provide actionable intelligence to IT organizations who are integrating new storage and backup solutions and delivery methods into their environment to address rapid growth, declining budgets, and the push to reduce costs.

The system 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 system's platform, an integrated Oracle DB that maintains current and historical data on the entire storage environment. From this database, each licensed module can analyze and correlate relevant information. The result is consistent analysis and reporting, enabling you to manage all aspects of your storage environment.

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

Virtualization Manager enables 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
- current usage and forecasting of virtual resource demands

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

IT Analytics components

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

- Portal Server. The physical server on which the 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.

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 to provide framework components.
- Binaries to provide 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 to instruct the Oracle SQL scripts to create the 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. The preferred configuration is to have the database on the same server as the Portal. If you use a Managed Services Provider, the Reporting Database and Portal Server are usually off site and hosted by that third-party provider.

Data is managed in the Reporting Database with automatic purging scripts that run, with specific retention periods per product and even data type. 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).

Data Collectors

A Data Collector is a centralized and remotely managed data collection mechanism. This Java application is responsible for interfacing with enterprise objects, such as backup servers, storage arrays, and switches, gathering information related to storage backup and recovery, and capacity management.

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.

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 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.

Chapter 2

Understand the Portal

This chapter includes the following topics:

- Overview
- Portal access
- Portal overview
- Navigating with search
- Valid entries for search
- Search results view: Narrowing the scope
- Filtering data in management grids
- Advanced filtering
- Customizing advanced filter logic
- Saving advanced filters (Inventory)
- Advanced filtering examples
- Advanced filter operators
- Clear advanced filters
- Delete saved advanced filters
- Select columns on management pages
- Detect alerts in the inventory
- About the Inventory tab
- About the Reports tab

- Reports tab keyboard shortcuts
- Reports tab basics
- Templates, reports, and dashboards overview
- About the Alerts tab
- About the Admin tab

Overview

The **Inventory** serves as a browser for your infrastructure. All objects in the database can be seen and managed in this single management window. These objects can be individually selected to see details or to run reports against providing a broader view of your infrastructure.

You can securely view and centrally manage the storage in your data center through a library of report templates, interactive dashboards and tools to help you create your own custom reports.

Portal access

The portal allows more than one Administrator and End User to be logged in simultaneously. The Portal Administrator provides users with credentials.

1. To log into the portal, do one of the following from your browser:
   - If your company hosts the portal, use this URL in the browser: http://itanalytics.<domainname>
   - If a third-party company hosts your portal, use the URL provided by your Managed Services Provider.
2. On the Portal login page, enter your User Name and Password, then click **Log In**.

Note: You can control portal access using privileges as well as the `portal.properties` file. Using settings in the `portal.properties`, you can restrict a user ID from signing on multiple times with the same or different browsers. In this instance, the last user ID to login will have access to the portal. Other users with same login will be logged out.

Multi-language support and access considerations

During your first access, the portal login screen shows the text based on the browser language setting. The supported languages are Simplified Chinese, French, Korean, and Japanese. If the browser is set to any other language, the login screen appears in English. This happens only during the first login. After your login, you can change the language preference from the user menu. To set the language preference after accessing the portal, see *Manage your profile and set a language preference* section in the *IT Analytics User Guide*. After specifying your preference, you need to log out and login again for the preferences to take effect.

See "Manage your profile and set a language preference" on page 772.

See "Troubleshoot test connection locale issue" on page 828.

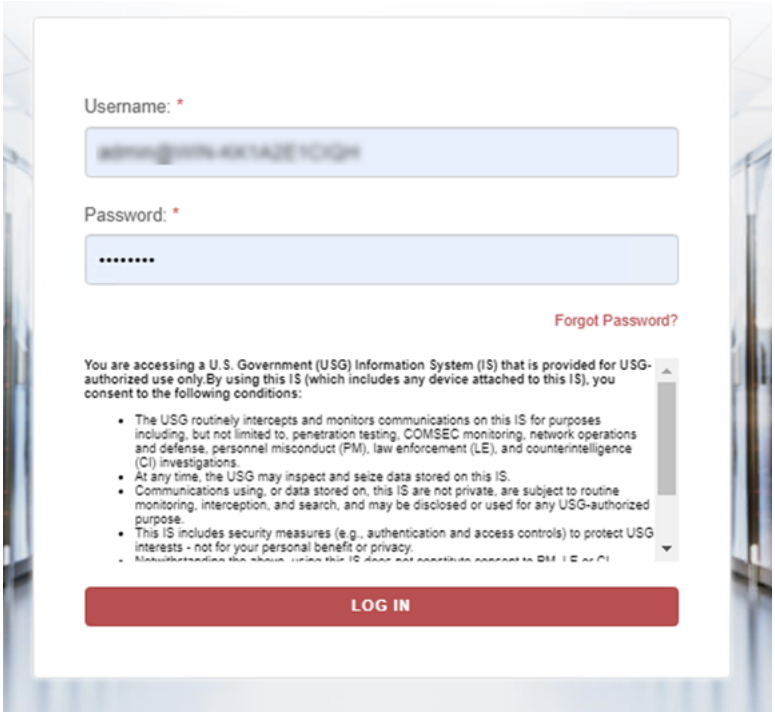
Department of Defense notice or custom disclaimer on the login page

The IT Analytics Portal login page may display a Department of Defense (DoD) notice or a custom disclaimer on the login page. The content of the notice may differ based on your organizational rules and regulations. To customize the notice, add your content to the `LoginConsent.html` file in HTML format.

File path:

- Windows: `C:\opt\aptare\portalconf\LoginConsent.html`
- Linux: `/opt/aptare/portalconf/LoginConsent.html`

For example, a sample portal login page and its HTML content are provided below.



Sample HTML text:

You are accessing a U.S. Government (USG) Information System (IS) that is provided for USG-authorized use only.By using this IS (which includes any device attached to this IS), you consent to the following conditions:

The USG routinely intercepts and monitors communications on this IS for purposes including, but not limited to, penetration testing, COMSEC monitoring, network operations and defense, personnel misconduct (PM), law enforcement (LE), and counterintelligence (CI) investigations.

 At any time, the USG may inspect and seize data stored on this IS.

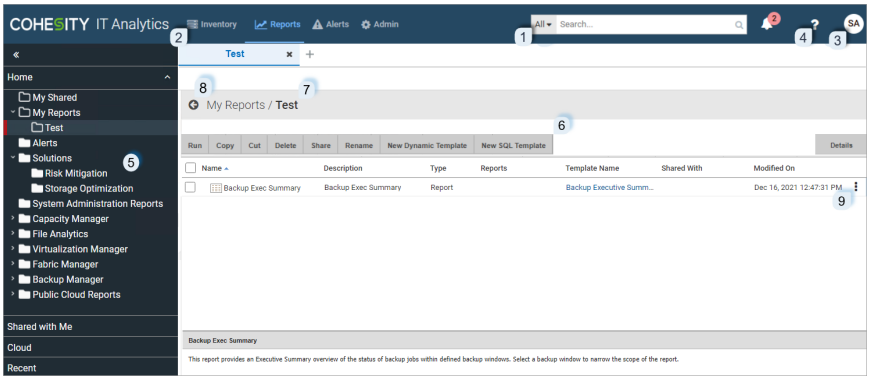
 Communications using, or data stored on, this IS are not private, are subject to routine monitoring, interception, and search, and may be disclosed or used for any USG-authorized purpose. This IS includes security measures (e.g., authentication and access controls) to protect USG interests - not for your personal benefit or privacy.

Notwithstanding the above, using this IS does not constitute consent to PM, LE or CI investigative searching or monitoring of the content of privileged

communications, or work product, related to personal representation or services by attorneys, psychotherapists, or clergy, and their assistants. Such communications and work product are private and confidential. See User Agreement for details.

Portal overview

Understanding Portal terminology is a prerequisite for getting started. The following figure highlights naming conventions for each area in the interface.



The Portal is organized into the following major sections, accessed from the top of your browser window.

1. **Search** - Use Search to discover, navigate to, and perform operations on a number of elements within the product. For example, navigate to the reports you require or find specific **Inventory** objects and Data Collectors/Policies. Search is case insensitive, supports partial entries, and will display a list of potential matches using a string match. If the Cloud is enabled, search results will also include user report templates available from the cloud. Because you can directly run a report from your **Search** results, this enables you to preview any cloud template before you save it locally. Results are displayed in a full search page that allows you to refine results further or start a new search.

   See “Valid entries for search” on page 37.

2. **Menu Bar**

   - **Inventory**: Provides a browser view of your collected datacenter infrastructure inventory. This view allows you examine your data starting from a logical entry point to narrow down, find, and run the reports you require. You can quickly find the inventory object type within your infrastructure and view the relevant reports. Object-specific reports, summary

pages are available, plus Host and Host Group management is enabled through the Inventory.

See “About the Inventory tab” on page 52.

- **Reports**: A management window for your reporting inventory that allows you to browse, manage, create and customize templates, reports and dashboards. This tab provides access to designers that enable custom report template creation: the Dynamic Template Designer and the SQL Template Designer, an advanced feature, requiring experience in SQL (Structured Query Language) development.

  See “About the Reports tab” on page 52.

- **Alerts**: IT Analytics collects a variety of data on various aspects of your datacenter and provides reports to visualize the results. Alerting enables you create and configure real time alerts for specific conditions within your datacenter allowing for faster resolutions to issues that are important to your organization.

  See “About the Alerts tab” on page 55.

- **Admin**: Provides consolidated access to functions that allow administrators to manage the Portal, create users, groups, set up data collection, plus set privileges for report and feature access. Also available, the Methods Designer, to create special processing methods to be included in report templates created with the Dynamic Template Designer and the File List Exporter (if you have licensed the File Analytics product). The exporter assembles File Analytics metadata into a comma-separated values (.csv) file, and allows you to export for further analysis.

  See “About the Admin tab” on page 55.

3. **User Account Menu**: This drop-down menu provides convenient access to account-specific features such as managing your profile information, your homepage, language selection, and access to your scheduled reports.

   For those assigned a Super User role, a **Clear Cache** selection is also displayed for the purpose of debugging.

   See “Clearing the inventory and report cache” on page 826.

4. **Help**: Displays context sensitive help for the page. This menu also displays the software version (**About**).

5. **Navigation Panel**: Browse through the reports categorized by product and functional area.

6. **Action Bar**: Displays functionality available for the selected element. As you select an item on a management grid, the available functions are displayed.

7. **Bread crumb**: Displays the navigation trail of the active tab. Items within the trail can be clicked to visit the history.
8. **Back arrow**: Restores the previous view of the active tab. The arrow is disabled if the active tab has no navigation history.
9. **Actions menu**: Invokes a quick actions menu with options similar to those present in the Action Bar. The menu disappears if the row is selected using the selection checkbox.

Navigating with search

Search enables you to discover, navigate to, and perform operations on a number of elements within the Portal. All objects within the Inventory are discoverable by Search. Reports, templates and dashboards can be found and generated on-demand using Search. Data collectors and their policies can be quickly located, providing immediate access for analysis or updating. Users can also be located using Search.

See "Valid entries for search" on page 37.

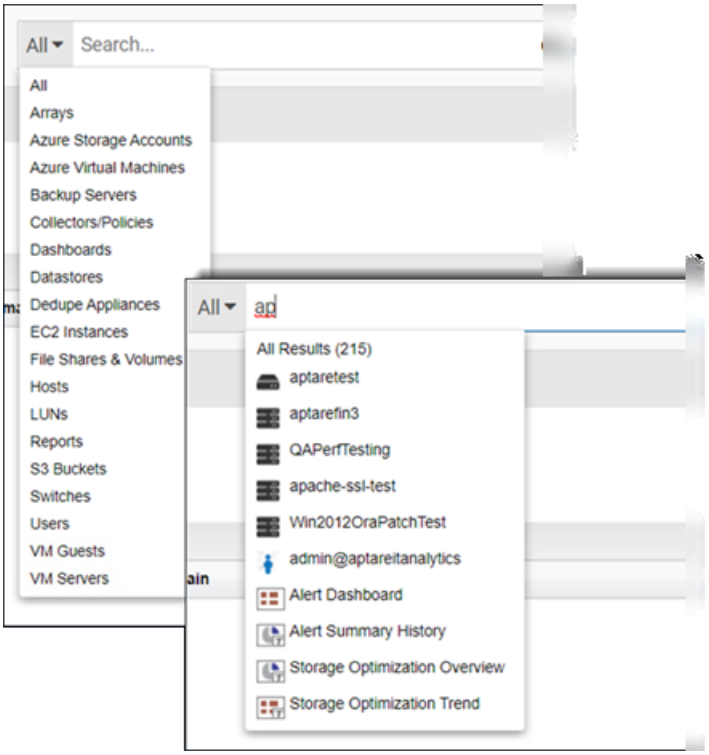
Additional filtering is available on those items displayed in management grids.

- See "Filtering data in management grids" on page 39.
- See "Advanced filtering" on page 40.

To search the across the Portal

- Enter a value in the **Search** field. You can select a category, such as Data Collectors or Hosts, from the drop down menu to further narrow the scope of your results. As you type, the first 10 potential matches are displayed. These are string matches. You can also click **All Results** in the match list to go directly to the Search Results.

See "To Refine Your Search Results" on page 107.



Searching for...	Search Results...
Reports and Dashboards	For reports and templates, search crawls the report/template name, short and long description. The search functionality allows you to search for the OpsCenter reports. Specify the OpsCenter report name in **Search** bar and press *Enter* key. The search result displays relevant IT Analytics reports. On the **Search Results** page, click the link in the **Source Location** column to navigate to the source report or template. You can edit or customize the report if applicable. Double-click a template to display the scope selector. This allows you to set the parameters before running it. Double-click a report or dashboard to run it. Click **Details** to toggle on the view panel to read a long description for a report or template. (This is only available for Backup Manager and Capacity Manager).

Searching for...	Search Results...
Inventory Objects	For Inventory objects, search crawls object name. For Hosts, search also crawls primary IP address. For VM Servers, search also crawls Cluster name. **Note:** In Search, Primary IP address requires the full address and does not support a partial entry. Search by, the page-level operation, supports a partial entry. On the **Search Results** page, if you locate the required inventory object, you can double click it and go directly to its specific Reports within the **Inventory**.
Data Collectors and Policies	For Data Collectors, search crawls the collector name. For collector policies, search crawls the policy name. **Note:** The policy name is system generated. The name is derived from a combination of vendor name and in most cases, a server IP address. On the **Search Results** page, if you locate the Data Collector or policy you're interested in, double click it to go directly to the **Collectors** page with the result selected.
Users	For Users, search crawls First Name, Last Name, Login name and email. A user can be active or inactive. On the **Search Results** page, if you locate the User you're interested in, double click it to go directly to the **Users** page with the result selected.

Watch this video to know the functionality of **Global Search** in IT Analytics

http://video.symantec.com/services/player/bcpid292374537001?bckey=AQ~~,AAAABuIiy9k~,l8Bhas-Vwr9zYL9V36WFi86fR_NoePscn&bctid=6301526417001

Valid entries for search

- Partial entries are supported

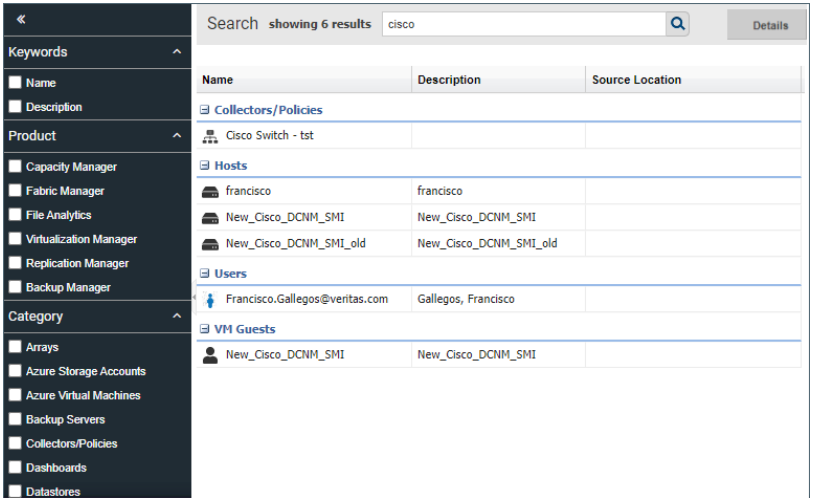
  Note: Primary IP address requires the full address and does not support a partial entry in the main **Search**.

- Case insensitive
- Alphanumeric a-z, 0-9
- Use double quotes for exact search (retain word order, reserve punctuation and all stop words such as _-., for, if, ...)

Search results view: Narrowing the scope

On the **Search Results** page, you can refine the results based on **Keywords, Product, Category and Source**. By default, all reports available to your user role, and those that are shared with you are displayed.

You must enter a search value to access the **Search Results** page. Click the Search icon beside the **Search** field. The search results are displayed on the **Search** page. Results are listed under categories.



On the **Search Results** page you can refine your results or start a new search. The **Search Results** page allows you to:

- Further refine results from a search
- Run a report
- Go to a location to manage the data collector/policy, user profile, inventory object, report, template, or dashboard
- View long descriptions for certain reports. (This feature is only available for Capacity Manager and Backup Manager).

To narrow your search results

1. Select a category and the results automatically update based on the selection.
2. For **Keywords**, enter a name or description in the field, and select **Name** and/or **Description**.
3. Double-click to run a report. Click an inventory object to navigate to its **Summary** page.

   Click the link in the **Source Location** column to navigate to the actual report. You can edit or customize the report if applicable.

Filtering data in management grids

Because the system collects and presents such a large volume of data, views often require additional filtering. On the management grids, filtering at the page-level provides the ability to further refine the list of data presented. Management grids include:

- Inventory list view
- Data Collection - Collection Status and Collector Administration
- Solutions: Storage Optimization and Risk Management
- Users and User Groups

See “Advanced filtering” on page 40.

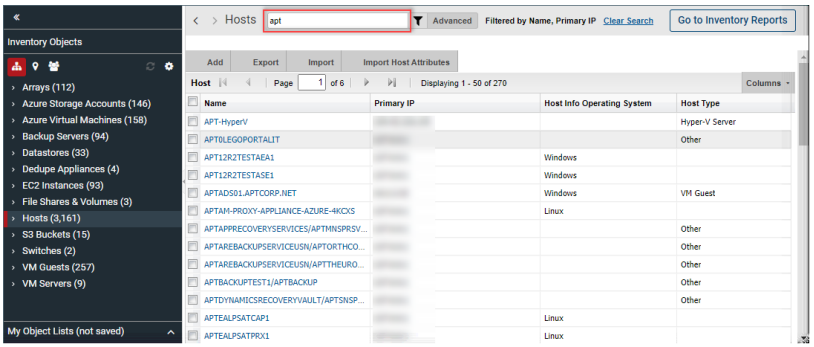
Valid Entries for Filtering

- Partial entries are supported
- Case insensitive
- Alphanumeric a-z, 0-9
- Use double quotes for exact match for a single entry (retain word order, reserve punctuation and all stop words such as _-., for, if, ...)
- Comma separated and space separated strings are allowed. Double quotes cannot be used in comma separated or space separated filters.

To filter at the page-level

1 Click a management grid.

2 Enter a value or comma/space separated string of values in the **Filter by** field.



3 Press **Enter**. The results are shown. A message is displayed to indicate the view has been filtered.

See “Clear advanced filters” on page 47.

Advanced filtering

- See “Customize report filter logic” on page 142.
- See “Saving advanced filters (Inventory)” on page 42.
- See “Advanced filtering examples” on page 44.
- See “Advanced filter operators” on page 46.
- See “Clear advanced filters” on page 47.
- See “Delete saved advanced filters” on page 48.

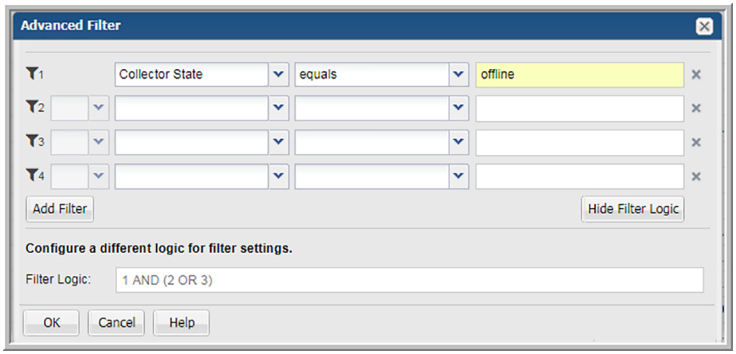
Advanced Filtering is another, more granular level of discovery available on the management grids. This powerful option uses database fields as the criteria along with the use of logical operators to build custom queries. Not all fields are searchable. Those fields are not offered as a selection when building an Advanced Filter.

An option to customize the filter logic is also provided.

In the Inventory, advanced filters are tied to a specific Inventory Object Type. You must either select one in the Hierarchy panel prior to building an Advanced Filter, or if your Inventory is not organized by Object Type, you must specify one when constructing the query.

To build an advanced filter

1. Navigate to a grid that supports the Advanced Filter:
   - Inventory list view
   - Data Collection - Collection Status and Collector Administration
   - Solutions - Storage Optimization and Risk Mitigation
   - Users and User Groups
2. Click **Advanced**.
   - If you are in the Inventory list view, choose an Inventory Object Type and then click Advanced.
3. Select a field name in the first drop down list.



4. Select an operator in the second drop down list.

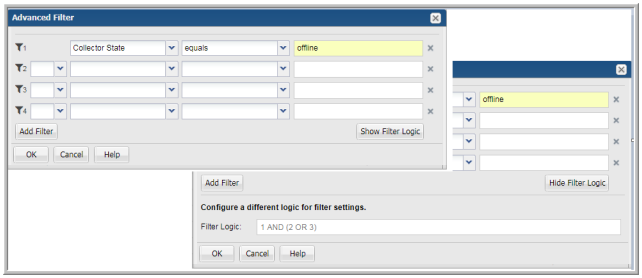
   See “Advanced filter operators” on page 46.
5. Enter a value in the third drop down list.
6. Choose a logical operator, **And/Or**, if you are adding another set of filters to your query.
7. Click **Add Filter** to add up to 12 additional filters.
8. Click **Show Filter Logic** to customize the filter logic if required. Logic defined in this field will override any setting established on the top of the dialog.

   See “Customize report filter logic” on page 142.

Customizing advanced filter logic

Customize filter expression order and the operators using the **Filter Logic** field. Logic defined in this field will override any setting established on the top of the dialog. Use the numbers on the left of the filter expressions to construct your **Filter Logic**.

1. Click **Show Filter Logic** to expand the window. This action disables the operators you set, and the Filter Logic field becomes mandatory.



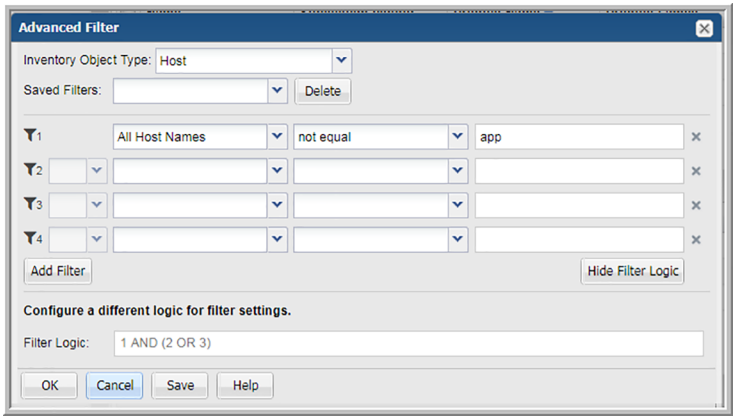
2. Edit the logic using the filter numbers and by adding parentheses or changing the operators. For example, you can change “1 AND 2 OR 3” to “1 AND (2 OR 3)”.

Saving advanced filters (Inventory)

In the Inventory, filters are tied to a specific Inventory Object Type. You can save Advanced Filters and access them by selecting the associated object in the Hierarchy panel. For example, if you create a Saved Filter for Hosts, you must select Hosts in the Inventory hierarchy to access the filter again.

1. Click the **Inventory** tab.
2. Select an object type in the Inventory Hierarchy panel. If your Inventory is not organized by Object Type, you can select one when constructing an Advanced Filter.

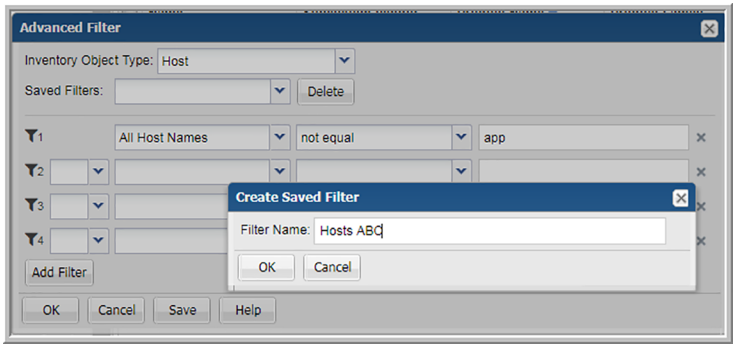
   See “Hierarchy toolbar to organize your data” on page 63.
3. Click **Advanced**. Choose an **Inventory Object Type** if the field is displayed.
4. Select a field name in the first drop down list.



5. Select an operator in the second drop down list.

   See “Advanced filter operators” on page 46.

6. Enter a value in the third field.
7. Choose a logical operator, **And/Or**, if you are adding additional criteria.
8. Click **Add Filter** to add up to 12 additional filters.
9. Click **Save**. The **Create Saved Filter** dialog is displayed.



10. Enter a name and click **OK**.

This **Saved Filter** becomes available for selection in the **Advanced Filtering** dialog for this **Inventory Object Type**.

Advanced filtering examples

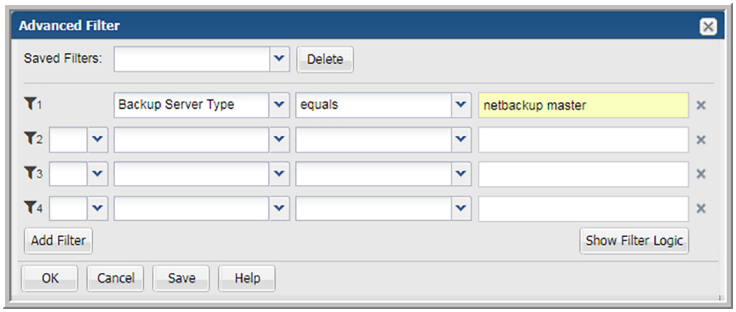
Example 1

Let us say you want to find all NetBackup Primary Servers to edit a property or assign an attribute.

1. Choose the Default Hierarchy (object, subsystem) in the **Inventory**.

   See “Hierarchy toolbar to organize your data” on page 63.

2. Select Backup Servers in the Hierarchy Panel.
3. Click **Advanced**.



4. Select **Backup Server Type**.
5. Set the operator to **equals** and enter **NetBackup Primary Server**.

If you want to see the Backup Server Type Name column in your grid display, you must explicitly add it.

See “Select columns on management pages” on page 48.

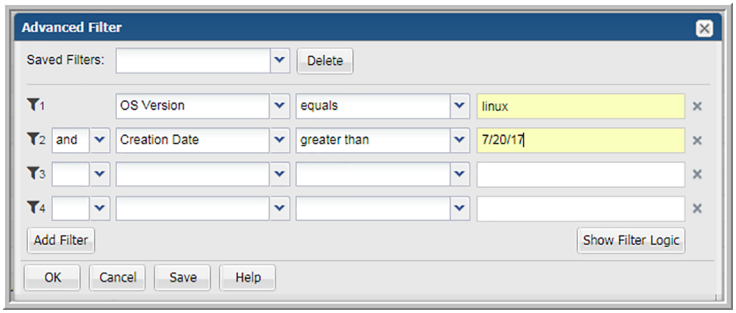
Example 2

Let’s say you want to find all hosts on a LINUX platform that were created after July 20, 2017.

1. Choose the Default Hierarchy (object, subsystem) in the **Inventory**.

   See “Hierarchy toolbar to organize your data” on page 63.

2. Select **Hosts** in the Hierarchy panel.
3. Click **Advanced** and the selections should be:



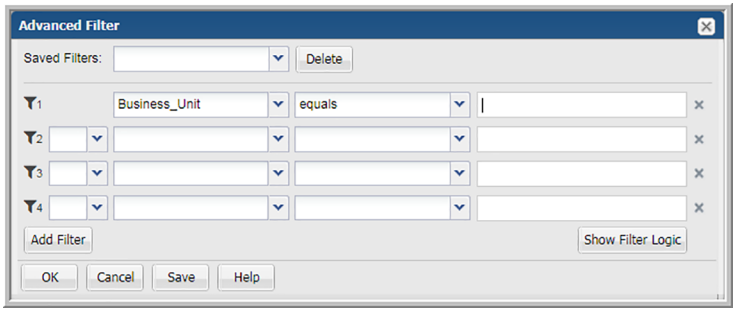
If you want to see the OS Version and Creation Date as columns in your grid display, you must explicitly add them.

See "Select columns on management pages" on page 48.

Example 3

Advanced Filtering is useful for discovering what's been left out of a criteria set. For example, find backup servers that have not been assigned to a Business Unit. Note, Business Unit would be an attribute you must assign to an Inventory Object.

1. Choose the Default Hierarchy (object, subsystem) in the **Inventory**.
2. Select **Backup Servers** in the Hierarchy panel.
3. Click **Advanced** and the selection should be:



You set the value as blank, but you must select an operator. In the example, the operator is equals.

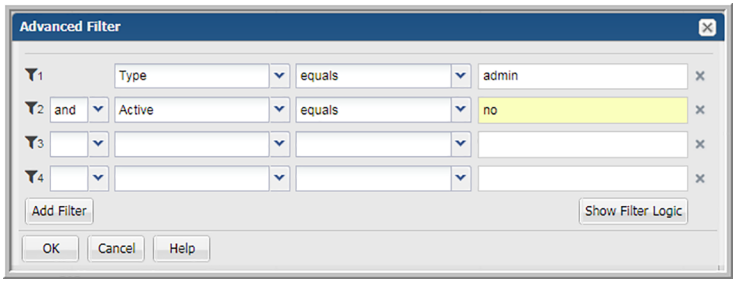
If you want to see the Business Unit as a column in your grid display, you must explicitly add it.

See "Select columns on management pages" on page 48.

Example 4

Let's find all the users who are assigned as an Admin but are no longer active.

1. Navigate to **Admin>Users>Users and Privileges**.
2. Click **Advanced** and the selection should be:



Advanced filter operators

Select from the following operators:

Table 2-1 Advance filter operators

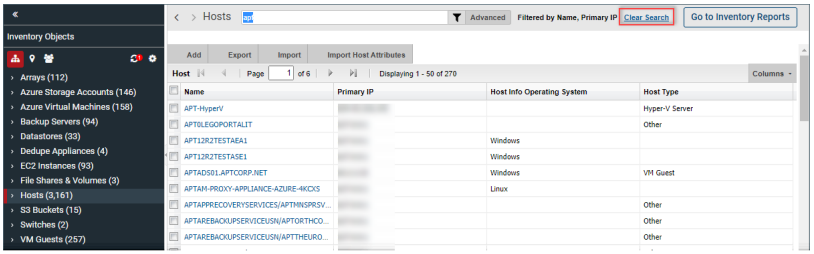
Operator	Description
equals	The result value is equal to the value entered. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
not equal	The result value is not equal to the value entered. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
greater than	The result value is greater than the value entered. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
less than	The result value is less than the value entered. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
greater than or equal to	The result value is greater than or equal to the value entered. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
less than or equal to	The result value is smaller than or equal to the value entered. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.

Table 2-1 Advance filter operators *(continued)*

Operator	Description
contains	The value of the entry is present anywhere in the value of the result. Example: If the entry is “rattle” and the value entered is “rat” or “at” will be displayed. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
does not contain	The value entered is not present anywhere in the value of the result. The entry can be a number, date, string or a blank. The match is case insensitive. Date entry format is based on your locale and the settings in your profile. To set the value as a blank, you must select an operator.
is a member of	The result value is equal to any of the comma-separated values entered. This string of values is the equivalent to a multi-select. There is no limit to the number of characters that can be entered for this value. The entry can be numbers and strings. To set the value as a blank, you must select an operator.
is not a member of	The result value is not equal to any of the comma-separated values entered. This string of values is the equivalent to a multi-select. There is no limit to the number of characters that can be entered for this value. The entry can be numbers and strings. To set the value as a blank, you must select an operator.
matches regular expression	The result value matches the regular expression entered. The entry can only be a string. See also: http://docs.oracle.com/javase/7/docs/api/java/util/regex/Pattern.html For example: To find all hosts that start with qa_, use ^qa_.*
does not match regular expression	The result value does not match the regular expression entered. The entry can only be a string. See also: http://docs.oracle.com/javase/7/docs/api/java/util/regex/Pattern.html

Clear advanced filters

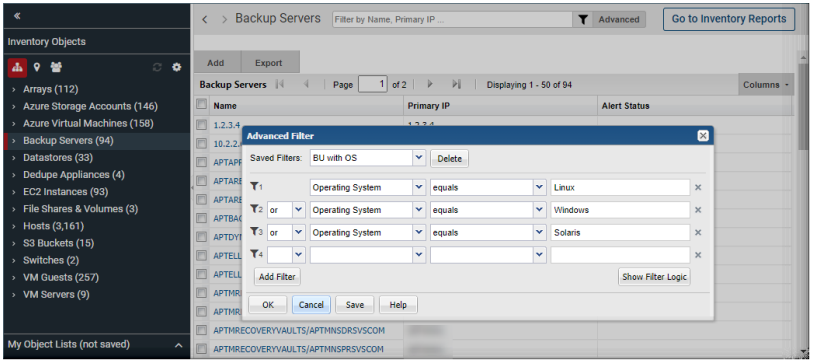
The **Filter by** page-level operation enables you to create a custom view on a management grid. You can quickly determine if your view has filters applied, and clear it if required. A message and a link to clear the filter is displayed beside the **Filter by** field.



Click the **Clear Search** link to remove all filter criteria and reset the view.

Delete saved advanced filters

1. Navigate to the Inventory.
2. Choose the Default Hierarchy.
3. Select the object type for which the **Advanced** Filter was created.
4. Click **Advanced**.
5. Select the **Saved Filter** from the drop down list.



6. Click **Delete**.

Select columns on management pages

The **Inventory** view allows you to create a custom views and choose the columns to display on the pages. The columns available for selection are based on a variety of factors, including your product licenses, collected data within the your domain, and pre-built custom attributes. Once selected, sort and move these columns to your preference.

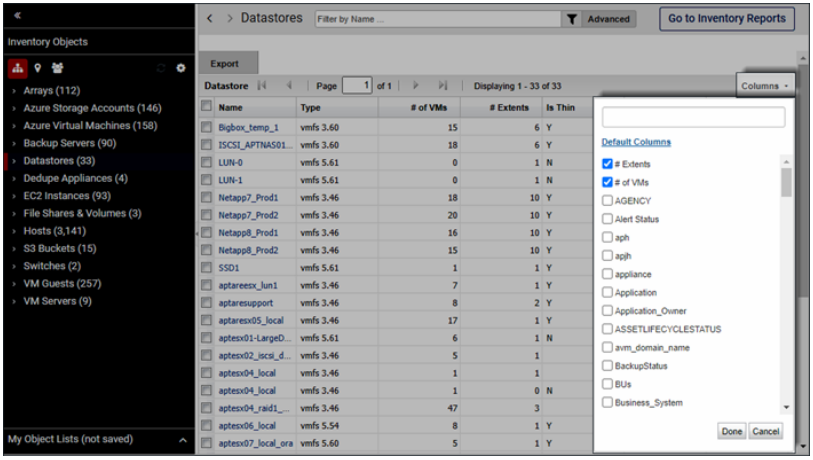
Note: The **Host Group** column, displayed in the Inventory for the Host Group management view, has sorting disabled to improve portal performance.

To add columns to the view

1 Navigate to **Inventory**.

2 Click **Columns**.

Note: If you select more columns than will fit on the current view, you can roll your cursor over the top of the **Columns** button to quickly see all column names that are selected for the active page.



3 Enter the column names you'd like to add to the view. The results shown in the drop-down list are listed alphabetically and filtered by keywords as you enter a value. Your entry is highlighted as it is found in the list. You can also scroll and browse through the list without searching.

4 Click the checkbox beside the column name to display it on the management page.

5 Click **Done** to display your choices.

To remove columns from the view

1 Navigate to the **Inventory**.

2 Click **Columns**.

3 Enter the column names to remove from the view. The results shown on the drop-down list are filtered as you enter the value and suggestions are displayed. You can also scroll and browse through the list without searching.

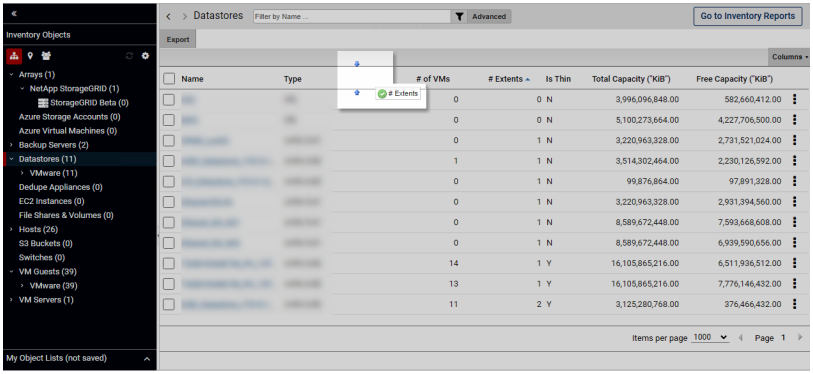
4 Deselect the checkbox beside the column name to remove it from the management page. The column is removed from the view when you click the page.

To reorder columns in the view

Once you've added/removed columns on the view, order them as you like. This custom order is persisted for your session.

1 Click the column name to move.

2 Drag and drop the column to the desired location.



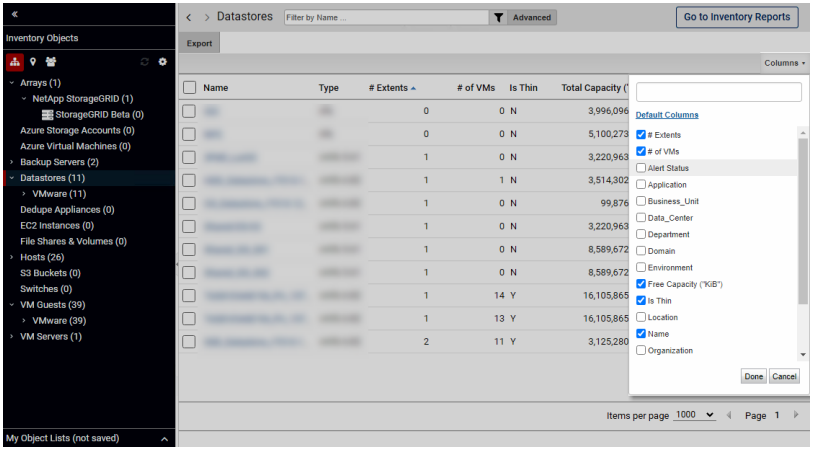
To reset to default columns

After customizing the columns shown, you can easily reset to the default columns.

1 Navigate to the **Inventory**.

2 Click **Columns**.

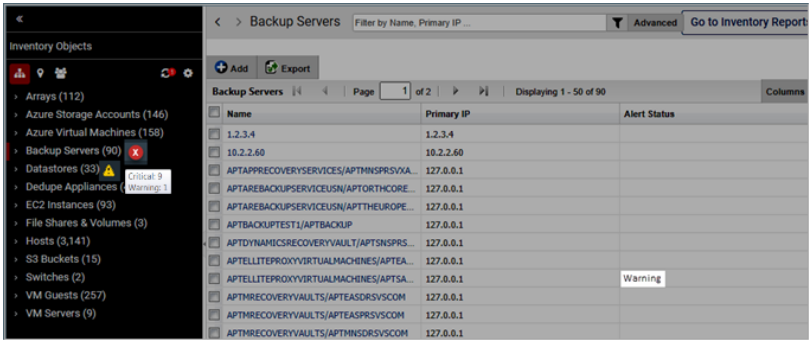
3 Click the link **Default Columns**.



Detect alerts in the inventory

Alerts can be delivered through a defined notification method. You can also quickly detect alerts in the Inventory with badging. Icons displayed in the Inventory Objects panel identify which objects have triggered a critical or warning alert using the thresholds you defined. Click the badge to display Alert Details.

By selecting the object, you can display reports that help you to identify where the problems are occurring and take action.



By activating the Alert Status column, you can also sort and filter on the status.

See “Select columns on management pages” on page 48.

About the Inventory tab

The **Inventory** serves as a browser for your infrastructure. All objects in the database can be seen and manipulated in a single navigation window. These objects can be individually selected to see details or to run reports against. Our reporting and dashboards provide a broader view of the state of the infrastructure, and the **Inventory** provides an object level view.

Access to the **Inventory** is controlled by user privileges and product licensing. Your user type dictates if the **Inventory** tab is displayed by default.

See “Inventory privileges” on page 61.

The **Inventory** provides an accurate picture of the assets discovered within your data center. This interface not only displays counts from collected data, but provides access to relevant information in the form of summaries, and templates (both custom and standard). This catalog of data center items allows you explore trends, capacity, system level information and more within the context of your inventory.

You can review data for short term goals such as performance and longer term goals such as trends in backup storage failures. The view into the **Inventory** can be customized by grouping the inventory objects by a variety of attributes such as inventory items type, location or business unit to further visualize the data in a way that is the most convenient to you.

Inventory objects are associated with default reports that offer relevant data or information for the selected object. These default reports can be customized to display a set of specialized reports that you choose and associate to the object. You can also quickly toggle between the landing pages and a list view that displays all the collected inventory objects for a category.

See “Exploring your inventory” on page 60.

About the Reports tab

The Reports tab is a management window into your reporting catalog, allowing you to browse through and manage templates, reports and dashboards. You can quickly identify when new/updated templates are introduced and choose to update or not. You can view the mappings between reports and templates, if sharing is enabled, detailed report descriptions and other information. On the management side, you can manage, create and customize templates, reports and dashboards. This tab provides direct access to designers that enable custom report template creation: the Dynamic Template Designer and the SQL Template Designer, an advanced

feature, requiring experience in SQL (Structured Query Language) development. The Reports tab also allows you to delete reports and organize them into a folder structures relevant to your business.

See “Templates, reports, and dashboards overview” on page 54.

Reports tab keyboard shortcuts

The Reports tab supports standard keyboard shortcuts to work more efficiently. Delete is available on other grid-based tabs in the Portal. Shortcut operations are controlled by Privileges.

The following keyboard shortcuts are available on the Reports tab:

Operation	Keyboard Shortcut
Copy	Ctrl+C
Cut	Ctrl+X
Delete	Del or Delete
Paste	Ctrl+V
Select All	Ctrl+A

Reports tab basics

There are a few basic navigation features that are important to know:

- Scrolling through tabs - As the number of opened tabs increases, a scrolling icon is displayed on either side of the page. This icon allows you to move quickly through your open tabs, even if they are not visible on the page.
- Changing the tabs order - To move a tab to a different place on the tab strip, drag it using your mouse. While you’re dragging a small indicator shows where the tab will land.
- Show tab list - As many tabs are opened, the browser will become filled with tabs and some will not be visible on the page. At this point, an arrow at the far right of the tabs enables a selection of the full list of open tabs.
- Create Dashboard - Launches a new blank dashboard.
- Select Report - Duplicates the currently open tab within the **Reports** tab.
- Closing one or all tabs - You can close all tabs at once, close a single active tab, or close all tabs but the active one by right-clicking directly on top of a tab.

Templates, reports, and dashboards overview

Access the rich set of collected data through reports, templates, and dashboards.

- Templates - Report templates provide the basis for all reports. The product comes packaged with a set of over 200 report templates, based on the IT Analytics products you've licensed. These templates are designed to provide a meaningful display of data, supporting business solutions for your enterprise. Users also can create custom templates using the SQL Template Designer or the Dynamic Template Designer.

  New or updates to system report templates are periodically published and automatically pushed to their appropriate folders. This provides instant access to new and improved report templates without waiting for a formal product release. These new and updated templates are badged with a NEW or UPDATED icon to provide a visual alert that templates are available. For new system templates, the Super User must then enable new reports in the privileges for individuals or user groups.

  See "Enabling new product report templates" on page 586.
  See "About badging" on page 102.

- Cloud Templates - New report templates are periodically published and made available through the Cloud section in the Reports tab. You can run reports directly from the Cloud, preview the data, and if you like report, save the template locally. Once saved, Cloud templates work exactly like out-of-the-box templates, except these are connected to the Cloud, so when improvements are introduced, your Cloud templates can also be updated if you choose to do so. Users with cloud privileges, will see a **NEW** badge on the **Cloud** section when new user report templates are introduced.

- Reports - Reports are generated from templates, with a scope selector that enables you to define the data set to be included in the output. A saved report includes the selected scope. Several reports can use the same template. Reports always have a template associated with them. You can quickly see this connection on the **Reports** tab. This allows you to map back to the original template if you'd like to make changes.

- Dashboards - Create your own dashboards by running reports and dragging them into a Dashboard window. Dashboards enable you to combine several reports into a single-pane view. Once constructed, you can customize the row height and determine if report headers are shown.

  See "Work with dashboards" on page 196.

About the Alerts tab

IT Analytics empowers you to intelligently and pro-actively ensure operational wellness. Through real-time symptom detection and notification, you can quickly spot problems across your datacenter, rapidly identify their causes, and minimize service degradation and disruption. The Alert feature enables you to craft a policy and watch specific conditions within your datacenter and create custom alerts to trigger based on set thresholds and time periods.

Once Alert policies are in place and thresholds defined, alert notifications can be sent and badging occurs to quickly identify the trouble spot in the Inventory.

About the Admin tab

The **Admin** tab provides integrated access to all administration functionality. Certain features are product-specific and are only available if that product is licensed in your environment. Also, access to administration functions is driven by the permissions specified for the user's login account.

The areas are structured as follows:

- See “Users” on page 55.
- See “Domains” on page 56.
- See “Chargeback” on page 56.
- See “Solutions” on page 56.
- See “Data Collection” on page 56.
- See “Reports” on page 57.
- See “Advanced” on page 57.

Users

- Users and Privileges - Allows you to create user accounts and assign privileges. When you create a user account, you create the user details, add access privileges, assign group membership and a password. Also, using this function, you select the system report templates that a user will be permitted to run and enable/disable cloud report access.
  See “Providing user access to the portal” on page 564.
- User Groups - Allows you to create groups and organize users into these groups. Once you set up a User Group, you can add the access privileges for the members of that group. A group is recommended when giving multiple users or

organizations access to reports. This is especially useful when newly supported subsystems and reports become available.
See “Creating user groups” on page 570.

Domains

- Domains - A Domain identifies the top level of your host group hierarchy. Most environments will have only one domain. Multiple domains enable organizations such as Managed Services Partners (MSPs) to compartmentalize access to data.
  See “Add/Configure a domain” on page 766.

Chargeback

- Backup- Allows you to create backup utilization billing and chargeback policies.
  See “ IT Analytics Billing and Chargeback policies” on page 623.
- Capacity - Allows you to add and edit capacity chargeback policies to allocate costs for storage array usage.
  See “Add/Edit a capacity billing and usage policy” on page 612.
- SAN Fabric - Allows you to create and manage SAN fabric chargeback policies to associate a cost with fabric and port usage. These policies are used by the SAN Fabric Usage report.
  See “Add/Edit a SAN Fabric Chargeback policy” on page 622.

Solutions

- Storage Optimization - Allows you to have visibility into storage utilization over time. This solution discovers on-prem and cloud resources that can be reclaimable or that can be optimized.
  See “Storage Optimization solution overview” on page 625.
- Risk Mitigation - Enables you to view curated analytics to enable proactive management for backup compliance and storage performance.
  See “Risk Mitigation solution overview” on page 646.

Data Collection

- Collection Status - Allows you to monitor data collection from a single window.
  See “Monitoring data collection status” on page 674.
- Collector Administration- Allows you to set up and manage data collection policies for a variety of enterprise objects, such as storage arrays and switches.
  See “Manage Data Collectors and collection policies” on page 657.

- Host Discovery and Collection - Allows you to configure and manage host inventory data collection as part of your Capacity Manager reporting.
  See “Understand the host data collection process” on page 679.
- Collector Updates - Allows you to view and manage the downloading of updates for collectors.
  See “To deploy updates to collectors” on page 673.

Reports

- Thresholds - Allows you to manage Capacity Manager threshold policies for file system, host group, host, LUN, Database, or NetApp Aggregate objects. Threshold Policies enable you to establish Low, Warning, and Critical levels from which to manage the state of your capacity utilization.
  See “Add/Edit a threshold policy” on page 610.
- Backup SLA - Allows you to create and edit SLA Group policies to manage service level agreements.
  See “Add/Edit a backup SLA policy” on page 620.
- File Categories - Allows you to define/edit file categories to report on groups of file metadata that has been collected by File Analytics. (Super-User only and File Analytics specific)
  See “File categories” on page 799.
- Primary Schedules - Allows you to maintain a primary report schedule, where reports are scheduled to run automatically on a regular basis.
  See “Configure primary schedules” on page 606.
- Backup Windows - Allows you to define custom backup windows for Backup Manager reporting.
  See “About custom backup windows” on page 607.

Advanced

- Parameters - Do not make adjustments to Advanced Parameters without assistance from Cohesity Support. These provide a mechanism for customizing data collection parameters to meet the needs of your organization.
  See “Overview of advanced parameters” on page 731.
- Attributes - Allows you to add and manage enterprise object characteristics to facilitate specific filters for reporting and serve as an organizer for the Inventory Navigator. For example, you could associate a data center attribute to servers/hosts and storage arrays to narrow the scope of a report or view your enterprise inventory by data center.
  See “ Manage attributes ” on page 540.

- Ransomware - Modify the default values of ransomware detection and NetBackup supportability. Ransomware administration allows you to set:
  - The percentage threshold up to which the storage pool space can be utilized for reporting on ransomware.
  - The lookback duration in days for ransomware RPO-RTO for job recovery.
  - The time duration in hours within which the service disruption must be restored.
  - The minimum supported version of Cohesity NetBackup primary server required to generate data for the ransomware scorecard.

  See "Ransomware administration" on page 785.

  Note: Access to ransomware administration is enabled through the user privileges or user group privileges. See "Assigning user privileges" on page 578.

- Object Maintenance - The database contains a variety of enterprise objects on which reports can be generated. These objects also can have attributes associated with them to further refine the scope of a report. Using the Object Maintenance tool, you can manage these objects.
  See "About object maintenance" on page 554.
- Support Tools - To help facilitate troubleshooting when working with Support, you can download data collector/portal log and configuration files for a specified time period.
  See "Support tools" on page 810.
- System Notifications - Provides access to system notifications with the ability to view and suppress them.
- System Configuration - Provides access for Super Users to configure a number of portal properties and system parameters that control everything from data retention periods to email settings to security.
- License Administration - Request a new license, view license details and upload it from this location.

Chapter 3

Explore your inventory

This chapter includes the following topics:

- Exploring your inventory
- Inventory privileges
- Getting started with the Inventory navigator
- Hierarchy toolbar to organize your data
- Show objects
- Reset inventory defaults
- Host groups
- Use attributes to organize your data
- Use host groups to organize your data
- Create custom object lists
- Working with the inventory reports view
- Accessing inventory reports
- Choose reports to display
- Filter within available reports by category
- Pin reports - saving reports with inventory objects
- Inventory reports and the action menu
- Work with the inventory list view
- Filter the inventory list view

- Manage objects in the inventory list view
- Assign attributes in the inventory list view
- Set attribute values in the inventory list view
- Import host attribute values
- Export objects from the inventory list view
- Delete objects using the inventory list view
- Customize columns in the inventory list view

Exploring your inventory

Serving as a window into your infrastructure, the **Inventory** offers an accurate picture of the assets discovered within your data center. Objects in the database can be seen and manipulated through this single navigation window providing access to information in the form of summaries and reports. This inventory of your data center allows you to view asset counts, explore trends, plus monitor capacity and system-level information. The Inventory also displays Alerts on objects as defined by you - with direct access to reports that help define root cause.

Inventory objects and groups are associated with default reports that offer relevant data or information for the selected object. Custom report templates created using the SQL Template Designer or Dynamic Template Designer are also displayed. If you've assigned a category, report templates are grouped and if not, they are shown as uncategorized report templates.

You can customize the display of specialized reports associated with the object. For groups of objects, you can review data for short-term goals such as performance and longer term goals, such as trends in backup failures. You can also quickly toggle between the Inventory reports and an Inventory list view, which displays all the collected inventory objects for a category.

Watch this video to explore more about the **Inventory** module.

http://video.symantec.com/services/player/bcpid292374537001?bckey=AQ~~,AAAABuIiy9k~,I8Bhas-Vwr9zYL9V36WFi86fR_NoePscn&bctid=6301527923001

Inventory privileges

Access to the **Inventory** is controlled by user privileges and product licensing. Your user type dictates if the **Inventory** tab is displayed by default. There are multiple levels of access:

- Restrict access to the entire **Inventory**
- Restrict access to individual inventory objects
- Restrict the ability to view reports (controlled through Report privileges)
- Restrict the ability to assign attributes
- Restrict the ability to permanently delete objects
- Restrict host management functions

See "About user privileges" on page 572.

Getting started with the Inventory navigator

Customize the view into your **Inventory** by grouping the inventory objects in a way that is the most relevant to your business. Once you've defined the view and the hierarchy structure, the Portal presents only the reports relevant to that structure. This allows you to see the reports and analyze the data faster and more efficiently.

View your inventory by the following organization structures:

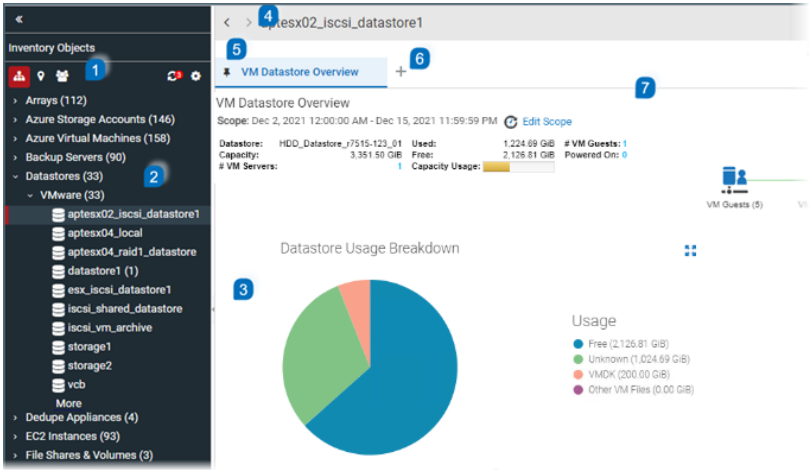
- Object type and subsystem
- Location and object type
- Domain
- Cloud
- Custom attributes
- Host groups

You can also create ad hoc custom object lists for homogenous objects and see the relevant reports.

See "Create custom object lists" on page 75.

Further, the information is presented in two view options, a list view and a reports view.

The following figure highlights key terms and functional areas in the Inventory.



1. Hierarchy Toolbar - Use the toolbar to quickly organize and refresh the display of your inventory. The icons stay highlighted to indicate the view you selected. Select from:

   Object type and subsystem (default organization)

   Location and object type

   Host groups.

   See “Use host groups to organize your data” on page 73.

   Refresh - When the icon displays a notification badge, rollover to identify what changes have been made in the Portal.

   Configure. Set the hierarchy manually using any of the defaults or create a custom organization.

   See “Hierarchy toolbar to organize your data” on page 63.

   Note: You can also create ad hoc lists for like objects, for example a specific set of hosts, and view relevant reports for just that selection. These object lists are displayed in the **Hierarchy Panel**.

   See “Create custom object lists” on page 75.

   ---

2. Hierarchy Panel - Navigate through the inventory by expanding and collapsing categories. Select a category or an individual object to display relevant reports. If an attribute is used to group objects, this is easily identified in the panel. Clickable Alerts are also displayed in the hierarchy.

See “Use attributes to organize your data” on page 71.

See “Detect alerts in the inventory” on page 51.

3. Details View - Displays the relevant information as it applies to the selection in the Hierarchy panel. You can toggle this area to display the information in a grid, or Inventory List view, or as a set of preselected reports, or Inventory Reports view.

   See “Work with the inventory list view” on page 84.

   See “Working with the inventory reports view” on page 77.

4. Navigation Buttons - Navigate within a single session, to jump back and forward through your **Inventory** history. This is useful when trying to complete a root cause analysis and you need to jump between multiple objects/pages within the Inventory.

   Note: Changing the hierarchy or logging out clears the **Inventory** history.

5. Pin a Report - Allows you to keep your favorite reports open and associated with an Inventory Object Type. Pinned Reports can also be renamed.

   See “Pin reports - saving reports with inventory objects” on page 81.

6. Select Reports - Click to view and choose each new report for the selected object in the **Hierarchy Panel**.

   See “Choose reports to display” on page 78.

7. Data Age Indicator - When a report is served from the browser cache, an indicator icon is displayed on reports and dashboards. You can roll over the indicator to show the age of the report from the cache. Click the icon to purge the old report from the cache and rerun the report from the database.

   When you run a report, the Portal takes the scope of the report, and checks if the cache contains the same report, for the same scope. If it does, the results are displayed from the cache. If the combination does not exist, the report is run from the database, saved in the cache, and then sent to the user interface. Cached reports are shared across users who belong to the same home host group.

   See “How reports and caching work together?” on page 839.

Hierarchy toolbar to organize your data

Use the **Hierarchy** toolbar to organize the view of your collected inventory. You can customize the view using the provided default categories or create custom

attributes to organize your data. Once set, these organizational categories will display only the reports relevant to your Inventory selection.

For example, if you choose to organize by Host Groups, once you select a group, your collected data center inventory is displayed and only the reports relevant to the selected Host Group are displayed.

The toolbar includes:

1. Default Hierarchy - Lists the inventory object type as the first-level category and vendor subsystem as the second. For example, Hosts is the top category and EMC Symmetrix would be the second-level category.
2. Location - Assigns a two-level selection for the location attribute, Location and Inventory Object Type. Working with this attribute, you can organize the view into your data center by location.

   Note: You must still create values and assign them to your inventory objects

   (See "Assign attributes in the inventory list view" on page 86.)

   Values do not display until this step has occurred.

   ---

3. Host Groups - Organize your data by Host Groups. Host groups can be nested within other host groups, thereby organizing hosts into sub-groups. In addition, a host can appear in multiple host groups so that it can be included in reports for the different groups. Refer to

   for complete information about Host Groups and their usage.
4. Refresh - When additions or deletions occur in the **Inventory**, the Hierarchy Panel does not automatically reflect the update. These can be changes made by you or someone in your host group. The **Refresh** icon displays a numeric badge to indicate changes have been detected. You can roll your cursor over the badge to see what's changed.
5. Configure - Set the hierarchy manually using any of the defaults or create a custom organization with your own multi-level organization. You can define up to three levels using the Hierarchy Panel.

   Predefined organization settings include, object type and subsystem vendor, location and object type, domains, cloud, and host groups. Configure also allows you to set the object display in the **Inventory**. Settings enable you to pick and choose which objects to show in the **Inventory**. You can also choose to hide decommissioned hosts.

Attribute value limitations

If an attribute has more than 400 values, and is assigned to an object, they are displayed as one value and labelled Aggregated for <n> values. This is applicable for:

- The **Host Overview** report
- A value displayed in a grid column for the attribute

- The **Inventory Hierarchy** will show as a single group with the label as the count value

Best practice is to limit attribute values to 50 total. Contact Cohesity Support for additional information.

Selection combinations

Note: Not all selection combinations are valid. Invalid combinations may result in duplicate folder names. For example, using VMware, if you remove the top-level **Inventory Object Type** from the hierarchy selections, duplication can occur as follows:

VMware

Guest1

Guest2

VMware

ESX1

ESX2

Inventory object limits

The **Inventory** is designed to manage large-scale environments and has been evaluated to support up to a minimum of 250,000 objects. Your browser vendor/version may actually support more objects without an impact to performance.

If your browser cannot accommodate the volume, you can reduce the total number of items displayed in the **Inventory**.

See “Show objects” on page 66.

for details about how to enable and disable object display.

Show objects

Use **Show Objects** to control the display of items in the **Hierarchy Panel**. If your browser limitations, (browser type, version and available memory) cannot accommodate the volume in your data center, you can reduce the total number of items displayed in the **Inventory** by de-selecting objects. You must select at least one object. Objects are controlled by privileges - meaning, for an object to be shown as a selection, you must have the privilege for that object. A user is assigned access

to individual object types. Administrators can also define which objects are selected by default.

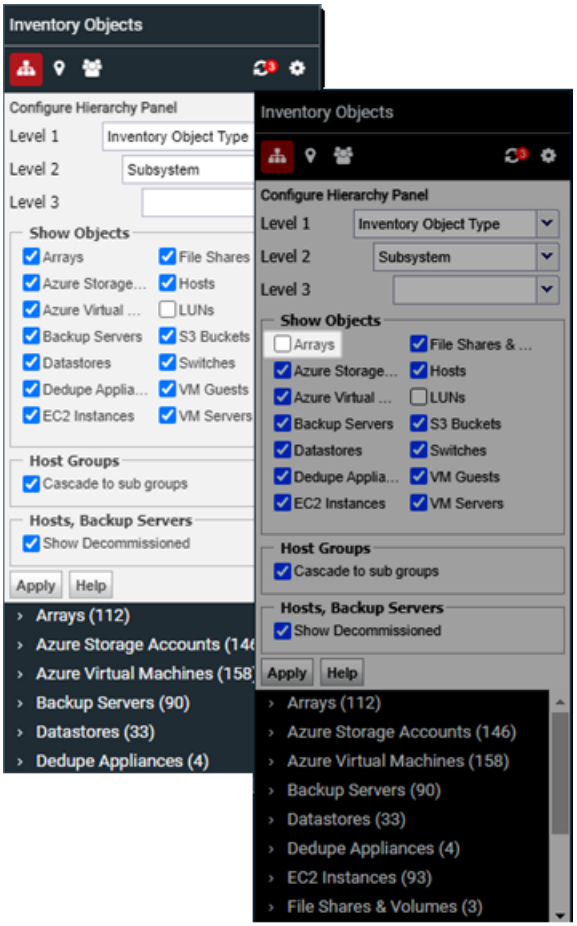
See “About user privileges” on page 572.

Note: LUNs are de-selected by default for most user roles. This is to address possible performance issues as related to the volume of LUN data. You can always select LUNs and click **Refresh** to display the objects in the Inventory.

Backup Servers in the Inventory

Backup Servers displayed in the Inventory are based on:

- The associated Host Type
- If at any point in time, jobs were collected against the server, even if the host type has been changed.



Cloud-based objects are displayed under **Show Objects** regardless of installed products. If you do not collect data from these objects, your Inventory hierarchy will display them with a count value of zero. Cloud-based objects include:

- Amazon EC2 Instances
- Amazon S3 Buckets
- Microsoft Azure Storage Accounts
- Microsoft Azure Virtual Machines

Note: You can also choose to hide decommissioned Hosts and Backup Servers. This is a selection in the **Hierarchy Panel** along with **Object Types**.

Count values

Groups with a zero (0) value on their display may be because:

- Data has not been collected. You must install data collectors and set up collection policies.
- Host object types may not be directly assigned to the host group. When you select **Cascade to sub-groups**, the top-level host group count reflects the count of its immediate members (all hosts and backup servers assigned to the group) plus the members of its child host groups. The counts are rolled up to the top level host group. When **Cascade to sub groups** is not selected, the count for each group only reflects hosts/backup servers directly assigned to it.

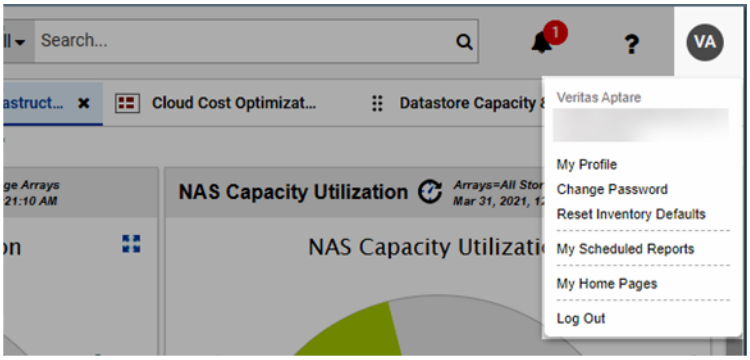
Reset inventory defaults

Default object selection may be set by your Administrator. You can always select any displayed objects. In some cases, environments with large volumes of data may be impacted by browser limitations and all object selections may not display in the Inventory. You can reset the Inventory to the default selections to work around issues with the browser.

To Reset Inventory Defaults

1 Click the User Account menu.

2 Select **Reset Inventory Defaults**.



3 Log in again to view the new selections.

Host groups

Host Group structures are one of the default organizational structures for the Inventory. Choose this setting to view your inventory by Host Groups you've setup. Only the reports relevant to the selected Host Group are displayed.

See "Use host groups to organize your data" on page 73.

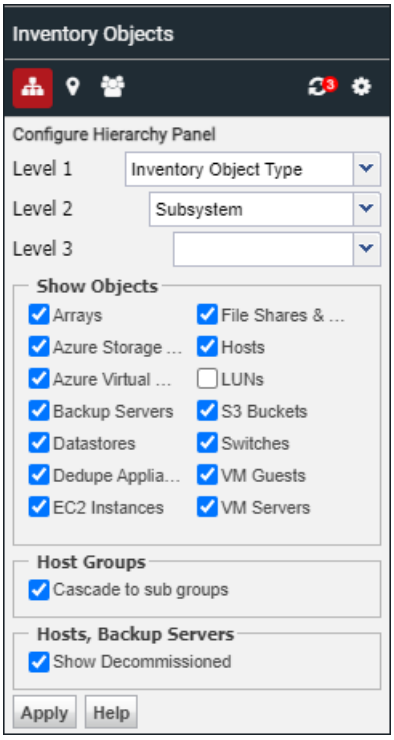
for complete information.

Select **Cascade to sub groups** to display host groups in an aggregated mode. This structure determines how host counts are reported.

When you select **Cascade to sub groups**, the top level host group count reflects all hosts and backup servers assigned to the group and all directly assigned to its sub groups. The counts are rolled up to the top level host group.

When **Cascade to sub groups** is not selected (non-aggregate mode), the count for each group only reflects hosts/backup servers directly assigned to it.

See "Count values" on page 69.

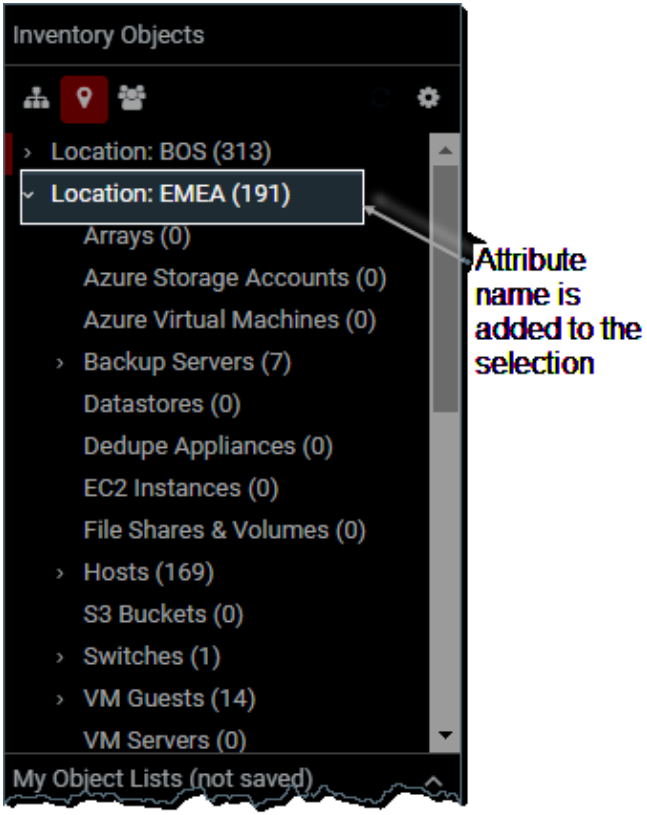


Use attributes to organize your data

You can use attributes to organize your inventory data. Once set up, these attributes are available for selection in the **Hierarchy Panel**. Attributes must have values associated with them before they can be used to organize your Inventory view. When viewing the **Hierarchy Panel**, the attribute name is prepended to the attribute value.

See "Hierarchy toolbar to organize your data" on page 63.

See " Manage attributes " on page 540.



System attributes

Predefined system attributes include:

- Application

- Business_Unit
- Data_Center
- Department
- Environment
- Location
- Organization
- Owner
- Region

Create custom attributes to further characterize an object with relevant properties. For example, you could create an attribute named Service_Required and set the attribute values as dates, and then assign it to your different Backup Servers. This would enable you to assign these properties to specific inventory types and organize the **Inventory View** by Service_Required.

See " Manage attributes " on page 540.

See "Host groups vs attributes" on page 537.

Attribute value limitations

If an attribute has more than 400 values and is assigned to an object, they are displayed as one value and labelled Aggregated for <n> values. This is applicable for:

- The **Host Overview** report
- A value displayed in a grid column for the attribute
- The **Inventory Hierarchy** will show as a single group with the label as the count value

Best practice is to limit attribute values to 50 total. Contact Cohesity Support for additional information.

To set up a custom hierarchy view

1. Click **Inventory**.
2. Click the **Configure** icon.

3 Select a top-level category from the **Level 1** drop down. This is a required selection. For multi-tenancy environments, such as MSPs, the first selection could be Domain. You can select up to three custom levels, but are only required to choose one.

Not all selection combinations are valid. Invalid combinations may result in duplicate folder names or inaccessible folder structures.

See "Selection combinations" on page 66.

4 Select the Inventory Objects you'd like to display. You can also choose to hide decommissioned hosts.

Note: Click the **Default Hierarchy** icon to return to the initial **Inventory Object Type > Subsystem** vendor hierarchy.

5 Click **Apply** when you complete all selections.

Use host groups to organize your data

Organize your data center view by Host Groups. Host groups represent relationships between hosts such as a business unit or a department. Host groups can be nested within host groups, resulting in sub-groups. In addition, a host can appear in multiple host groups. Host Group view is available in two modes - aggregate or non-aggregate. This is set using the **Configure** button.

See "Hierarchy toolbar to organize your data" on page 63.

This structure determines how host counts are reported.

See "Count values" on page 69.

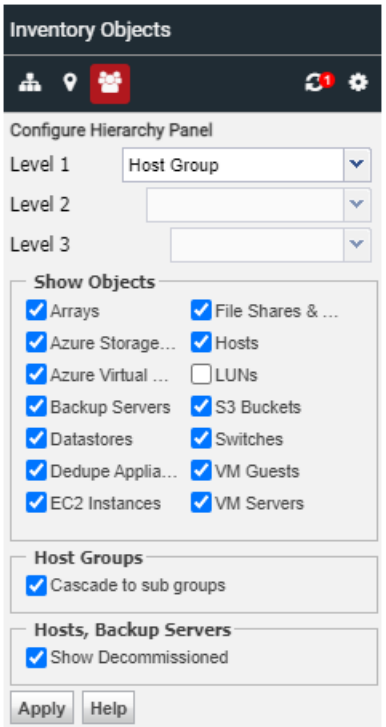
for more information.

The Inventory also provides management functionality for Hosts, Backup Servers and Host Groups.

See "Host groups vs attributes" on page 537.

To access Host Groups view

1. Click **Inventory**.
2. Click the **Host Groups** icon.



3. Click the **Configure** icon.
4. Select **Cascade to sub groups** to display host groups in an aggregated mode.

   This structure determines how host counts are reported. When you select **Cascade to sub groups**, the top-level host group count reflects all hosts assigned to the group and all hosts directly assigned to its sub groups. The counts are rolled up to the top-level host group. When **Cascade to sub groups** is not selected (non-aggregate mode), the count for each group only reflects hosts directly assigned to it.
5. Click **Apply** when you complete all selections.

Create custom object lists

You can create ad hoc custom groupings for homogenous object types and see the reports relevant to that custom set. This temporary list is only available until:

- you log out
- your session expires
- your browser is refreshed

If you are looking for a more permanent grouping, use Saved Filters in the Inventory.

See “Saving advanced filters (Inventory)” on page 42.

To set up a custom Object List

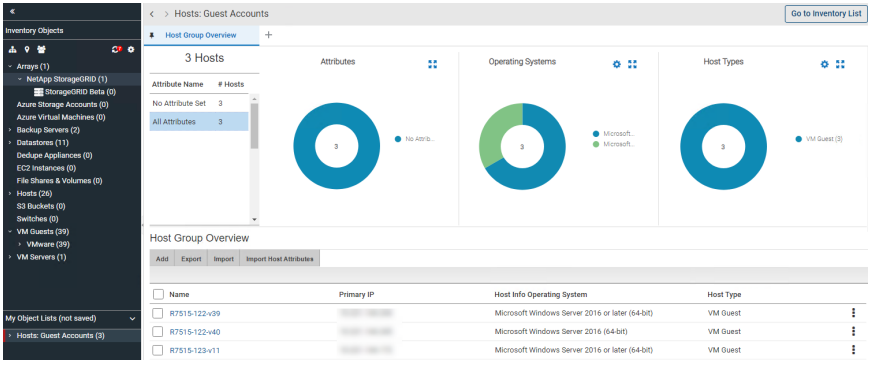
Note: Object types must be alike to create a custom group.

---

1. Navigate to the **Inventory**.
2. Configure the view to your preference.

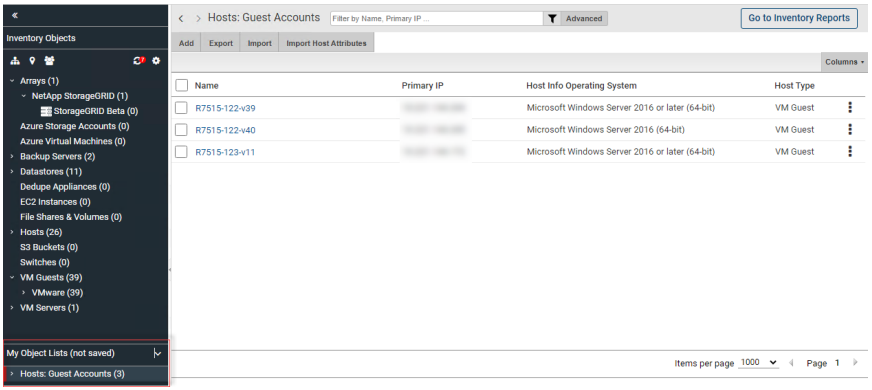
   See “Hierarchy toolbar to organize your data” on page 63.
3. If required, search for a specific set of objects within the Inventory.

   See “Navigating with search” on page 35.
4. Select the objects that you want to group.
5. Click **New List**. The **Add New Object** List dialog is displayed.



6 Enter a name and click **OK**.

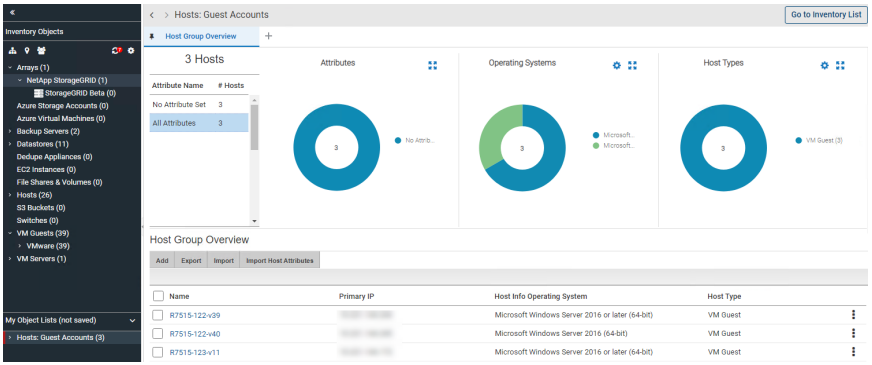
The new group is displayed under the **My Object Lists** section on the hierarchy panel.



If your object list is composed of hosts or backup servers, add individual ones to your list using the **Add** button. When you add them, they are added to your **My Object List** and to the host group you assign.

See “Manage objects in the inventory list view” on page 85.

7 Select the group and switch to Inventory reports to display the relevant reports.



The displayed report set is determined by the base object. For example, if you have selected Backup Servers and for that object type, the Inventory displays Current Media Summary and Backup Executive Summary those same reports are presented. You also customized what reports are displayed.

See “Choose reports to display” on page 78.

Working with the inventory reports view

Each inventory object category is associated with a default set of reports.

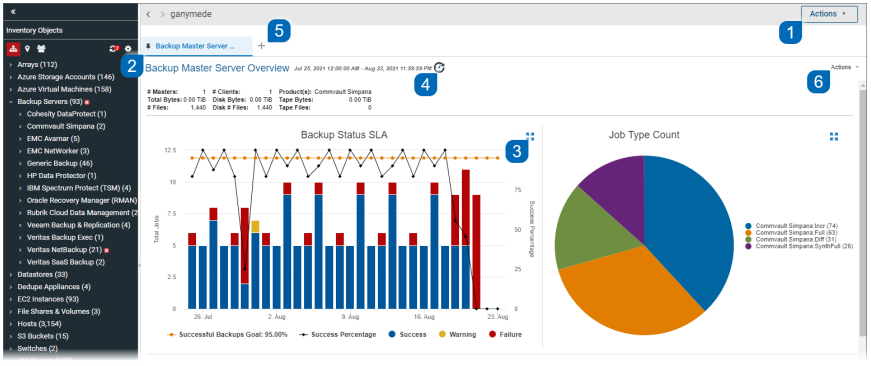
Note: Only a subset of the full report catalog is represented in the Inventory.

Groups with more than one object type do not have reports. For example, if you have sorted your Inventory by location, the object types could be different.

You can select what reports you want to readily access, by customizing and pinning an Inventory object's set of base pages. These reports offer relevant data for the selected category and are displayed as tabs across the top of the **Inventory** view. They are sorted into information categories such as performance, storage or forecasting. The report list is controlled by privileges and license restrictions. As you navigate within the **Inventory**, the system remembers your 10 most recently visited locations and displays them in the navigation panel.

A summary page, when available, is displayed by default for each object and category.

The **Inventory Reports** view enables you to:.



1. Toggle between the **Inventory List** and **Inventory Reports**.

   See "Accessing inventory reports" on page 78.

2. Pin a Report. This saves the report's association with this type of object and automatically displays each time you return and select an object of this type.

   See "Pin reports - saving reports with inventory objects" on page 81.

3. Expand the portlet outside of the dashboard for a full screen view of an individual chart.

4. Rollover the icon to view the age of the report data.

See “How reports and caching work together?” on page 839.

for details about report caching.

5. View the list of available reports for this object type and choose each new report to display in tabs.

   See “Choose reports to display” on page 78.

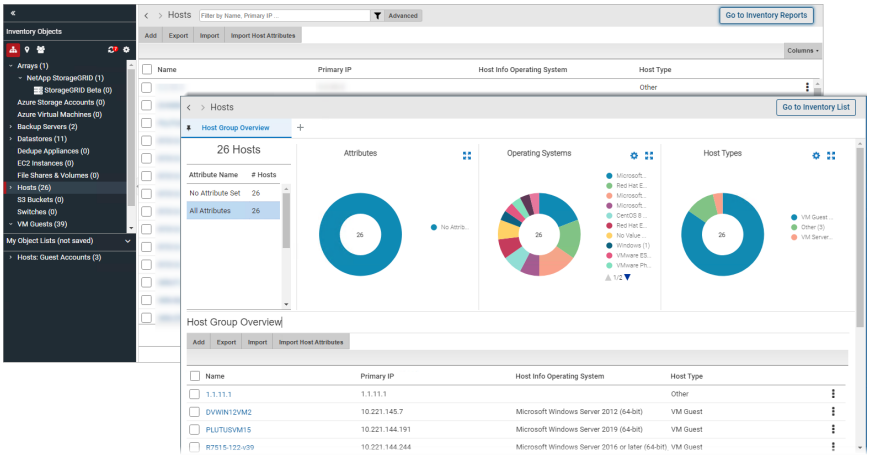
6. Click **Actions** to access the standard report functionality.

   See “Inventory reports and the action menu” on page 83.

- See “Interact with Reports” on page 159.

Accessing inventory reports

1. Click **Inventory**.
2. Select an object category or navigate to a single object in the **Hierarchy Panel**.
3. Click **Go to Inventory Reports** to toggle to the list of inventory objects within the category. An interactive summary page, when available, is displayed by default for each object or category.



Choose reports to display

You can choose the reports to display with an Inventory object. Once selected, your choices are displayed as tabs in the **Inventory Reports** view. Reports displayed under **Available Reports by Category** are controlled by privileges and license restrictions.

- See "Filter within available reports by category" on page 80.

Only a subset of the full report catalog is represented in the **Inventory**. A curated set of report templates are displayed and organized by a high-level category. When displayed in the **Inventory**, these templates are sorted into information categories such as performance, storage or forecasting.

Displayed reports are either tied to the object selection by subsystem or the report is generic in nature. The report list is controlled by privileges and license restrictions.

Custom reports built using the Dynamic Template Designer and the SQL Template Designer are also displayed.

Report templates built for arrays, hosts, backup servers, datastores, VM guests and VM servers using the Dynamic Template Designer or SQL Template Designer can be assigned an Inventory report category during the creation process.

See "Quick Start Step 1: Create a Basic Table Dynamic Template" on page 223.

See "Inventory Report Configuration" on page 416.

Object type and subsystem can also be defined for further classification.

If a category has not been assigned, these templates are displayed under the heading **Uncategorized**. You can always assign a category, by customizing an existing template.

See "Customize and Export Dynamic Templates" on page 292.

See "Edit a Custom Report Template" on page 418.

Your selections are displayed as functional categories in the **Inventory Reports** view.

To select the reports to display for an Inventory object

You can select any report to render, but you must explicitly Pin the report to save it once you log out of the session.

1. Click **Inventory**.
2. Select an object in the **Hierarchy Panel**.

3 Switch the view to the **Inventory Reports**, if required, by clicking the icon beside the object name.



4 Click the + tab to access a list of available reports to assign to the object.

5 Click the report name to select and display it in your **Inventory Reports**. If the name is not displayed as a link, it has already been selected.

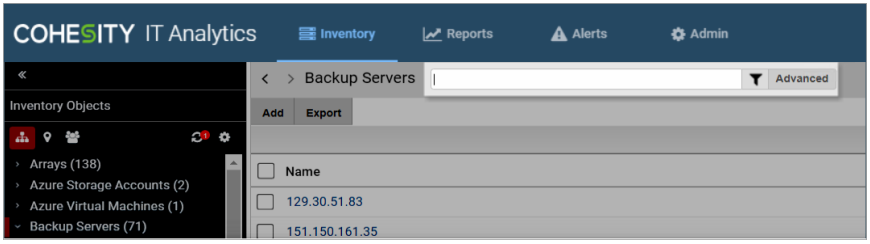
6 Pin the report to save the association with the Inventory Object Type to allow for quick access the next time you login.

See “Pin reports - saving reports with inventory objects” on page 81.

Filter within available reports by category

If you are looking for a specific report name or description within the list of **Available Reports by Category**, enter the value to refine the list. This search examines report names and descriptions. The list of reports is controlled by privileges and license restrictions.

1. Enter a value into the **Filter** field.



2. Click the **Filter** icon. The **Available Reports by Category** list is filtered to display only the reports that contain the specified value.

Pin reports - saving reports with inventory objects

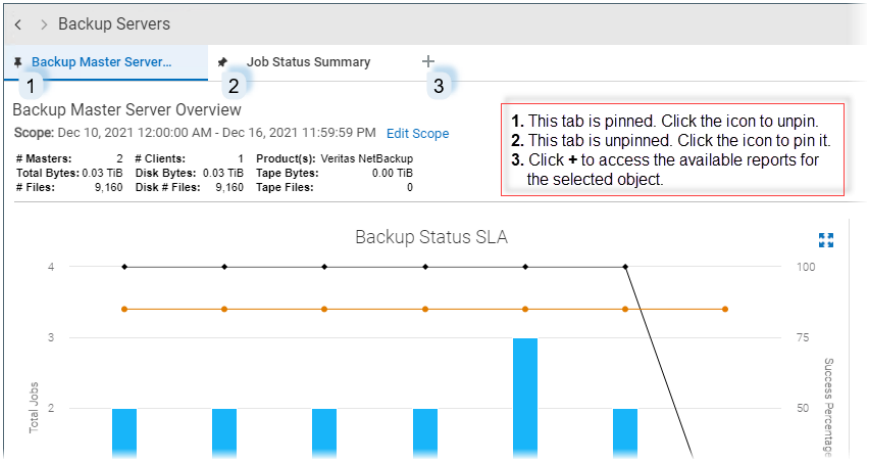
Once selected, Pin the report to save the association with the Inventory object. This action displays the selected reports as tabs in the **Inventory Reports** view each time you access the Inventory.

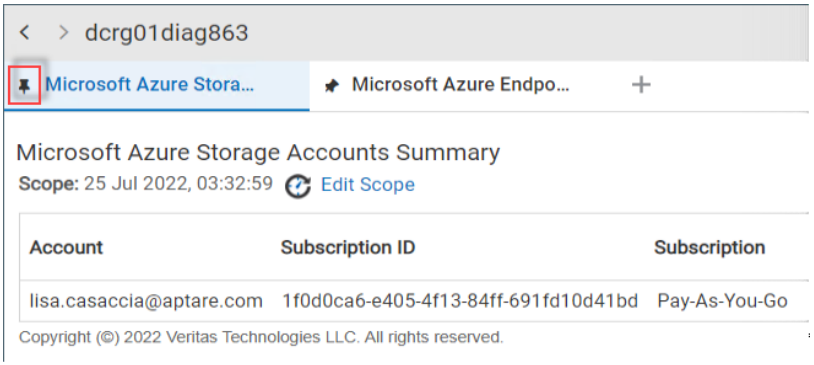
See "Choose reports to display" on page 78.

1. Click **Inventory**.
2. Select an object type in the **Hierarchy Panel**.
3. Click to switch the view if required.
4. Click the + tab to access a list of available reports to assign to the object.
5. Choose the reports to display.

   See "Choose reports to display" on page 78.

6. Click the Pin icon to save the report with the selected object. You can close a pinned report and it will continued to be pinned. When you login again, the pinned report will be displayed.





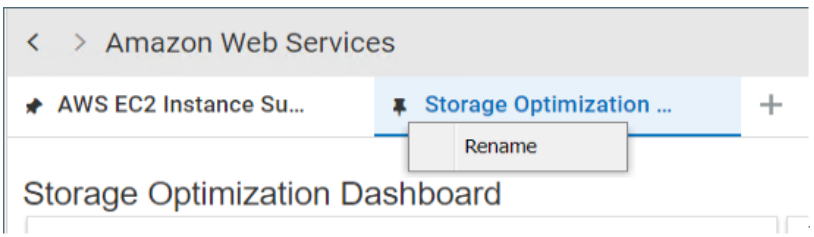
7. Click the Pin icon again to Unpin. One report must always be displayed, so if you try to unpin the last displayed report, an error message is displayed.

Rename pinned tabs

Once you pin a report, by default, the tab displayed in the Inventory is labelled with the report name. You can rename the tab to display a name that is more meaningful to your business. You can also quickly reset it to the default report name. This change only applies to the tab and does not impact the actual report name.

Note: Unpinned tabs cannot be renamed.

1. In the Inventory, right-click an active pinned tab.



2. Select **Rename** and enter the new name for the tab. The tab display name is updated.

Reset default report names on pinned tabs

1. In the Inventory, right-click an active pinned tab.
2. Select **Rename**.
3. Click **Default**. The tab display name is updated to the original report name.

Inventory reports and the action menu

Many functions available to standard reports and templates are available to the reports in the **Inventory Navigator**.

- Inventory Reports: **Edit Scope** - Your selection in the Hierarchy Panel sets the objects which are part of the scope. For some objects and object groups, you can edit additional parameters. **Edit Scope** is only shown when there are parameters available to change. A default time period is also set for those reports that require one. Use the Hierarchy Panel to revise the scope when the menu item is not displayed.
  See “Report Scope ” on page 110.
- Inventory Reports: **Save As** - When you save a report, you are saving a copy of the report into the **Reports** tab, not the report output. You can save a version in your Reports. Scope selections are saved with the report.
  See “Save Reports” on page 165.

  Note: Some reports provided by the Inventory are specialized and not available from the **Reports** tab. Due to the nature of these reports, they function like Detail reports and cannot be customized or saved, as they are specific to the report from which they were derived. **Save As** is not displayed for these reports.

- Inventory Reports: **Email** - After you generate a report and it renders, you can choose to instantly email the report. Emailed reports are not derived from the cache. These events are run in real time, so current data is always used. You can email a report to yourself, other individuals, or a distribution list. Scheduling reports to be regularly emailed is not available from the **Inventory** tab.
  See “Email reports and dashboards” on page 177.
- Inventory Reports: **Export** - You can export reports to make them available to external applications, such as Microsoft Excel or if you’d like to generate a hardcopy, you can export them to a file, such as a PDF. Exported reports are not derived from the cache. These events are run in real-time, so current data is always used.
  See “Export reports and dashboards” on page 169.
- Inventory Reports: **Filter** - In addition to the filtering that happens with your hierarchy selection, table-formatted reports can be further filtered on Rows and/or Columns, using advanced filtering. You can define the criteria for the data rows displayed in a report. Drop-down lists enable selections from the available columns. Next, you supply the operator--such as equals or does not contain--and a value for that column. Up to 16 selections can be joined to form the filter.

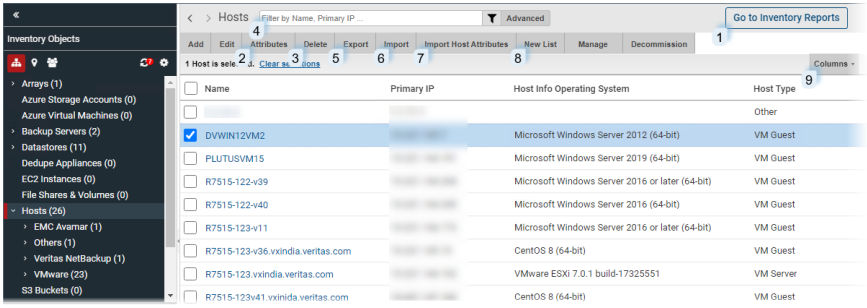
See “Advanced Filtering for Tabular Reports” on page 141.

Work with the inventory list view

Leveraging data collected from your data center, the **Inventory Navigator** provides an accurate picture of discovered objects. The database contains a variety of inventory object types from which data can be collected; for example hosts (and categories of hosts), switches, datastores and VM servers. For Amazon Web Service (AWS) users, EC2 Instances and S3 Buckets can be collected from; for Microsoft Azure, Storage Accounts and Virtual Servers. Individual categories are shown and populated only when associated objects have been discovered through data collection. Each category name has a count representing the number of discovered objects within that category.

The **Inventory List** view displays all the collected inventory objects for a category and provides the following:

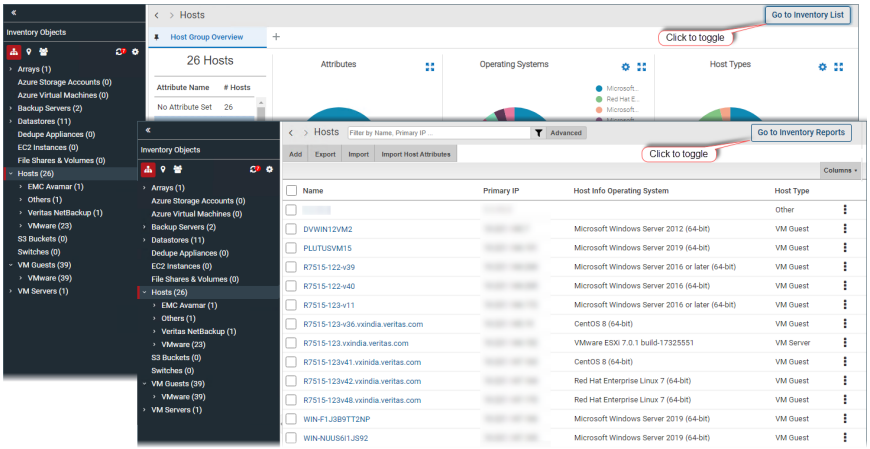
1. Toggle between the **Inventory List** view and the **Inventory Reports**.
2. See “Assign attributes in the inventory list view” on page 86.
3. See “Delete objects using the inventory list view” on page 94.
4. See “Filter the inventory list view” on page 85.
5. See “Export objects from the inventory list view” on page 93.
6. See “Import Hosts” on page 523.
7. See “Import host attribute values” on page 90.
8. See “Create custom object lists” on page 75.
9. See “Customize columns in the inventory list view” on page 95.



- See “Advanced filtering” on page 40.
- See “Set attribute values in the inventory list view” on page 88.

To access the Inventory List view

1. Click **Inventory**.
2. Select an item on the navigation panel.
3. Click the **Go to Inventory List** icon to toggle to the list of inventory objects within the category.



Filter the inventory list view

Advanced Filtering is available when you are in the Inventory List view.

See “Navigating with search” on page 35.

Manage objects in the inventory list view

The **Inventory List** view enables you to manage individual objects and groups of objects. For groups of objects, you must be in the **Inventory List** view to perform management operations. Individual objects display management functions when you drill into their report view. The following operations are available for some objects:

- See “Assign attributes in the inventory list view” on page 86.
- See “Export objects from the inventory list view” on page 93.
- See “Delete objects using the inventory list view” on page 94.

Note: These functions are privilege-based and not available for all object types.

Additional management operations are available for Hosts and Backup Servers:

- See “Adding and Editing Hosts and Backup Servers ” on page 521.
- See “Import host attribute values” on page 90.
- See “Decommission/Recommission Hosts and Backup Servers” on page 526.
- See “Import Hosts” on page 523.

Assign attributes in the inventory list view

Use the **Inventory List** view to assign attributes and select their values within the **Inventory Navigator**. Once attributes are created and their values defined, you can assign their values to objects represented in the **Inventory**. A set of predefined are included in the Portal software. Once values are defined for the system attributes, you can assign them. In the **Inventory List** view, you can select a single object or a group of objects and assign or modify attribute values associated with them.

See “System attributes” on page 71.

About attributes

Attributes provide a way of defining a set of objects that share a certain characteristic. Attributes represent logical relationships between objects and their relevant characteristics. Typically, you’ll set up attributes to aid in defining the scope of a report. For example, you might set up a “maintenance contract” attribute that you can associate with the hosts for which you have service coverage. Or, you might create an attribute to organize backup servers by geographical location so that the administrators responsible for the hosts at each corporate location can generate reports for their specific sites.

You can select hosts in bulk and assign or modify attributes associated with them. Use search or Advanced Filters to create a results set with the hosts you’d like to modify. For example, you can create a search query to find all Windows 2008 R2 systems as reported by NetBackup. When the search results are displayed, you can quickly apply a custom attribute such as patch applied.

- See “ Manage attributes ” on page 540.
- See “ Set attributes on hosts ” on page 540.
- See “Import host attribute values” on page 90.

Note: Attributes can also be displayed as a column in the **Inventory List** view.

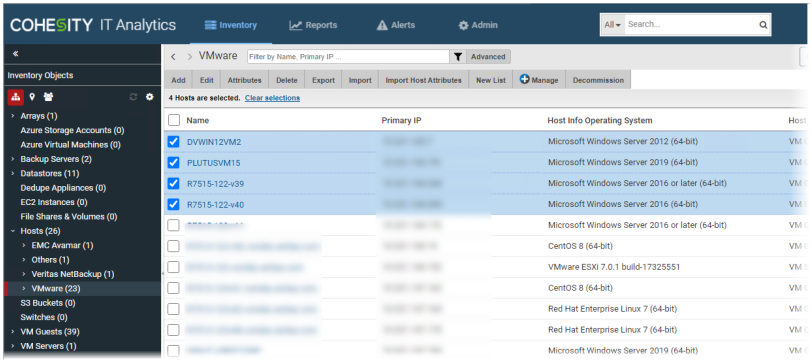
See "Filter the inventory list view" on page 85.

Attribute limitations

- Attributes cannot be assigned to LUNs, EC2 Instances, S3 Buckets, Azure Storage Accounts or Azure Virtual Machines.
- Attributes are associated with a Domain.
  See "Attribute inheritance" on page 539.
- If an attribute has more than 400 values, and is assigned to an object, they are displayed as one value and labeled Aggregated for <n> values. Best practice is to limit attribute values to 50 total. This is applicable for:
  - The Host Overview report
  - A value displayed in a grid column for the attribute
  - The Inventory Hierarchy will show as a single group with the label as the count value
  - Contact Cohesity Support for additional information.

To assign attributes and set values in bulk

1 Select **Inventory**.

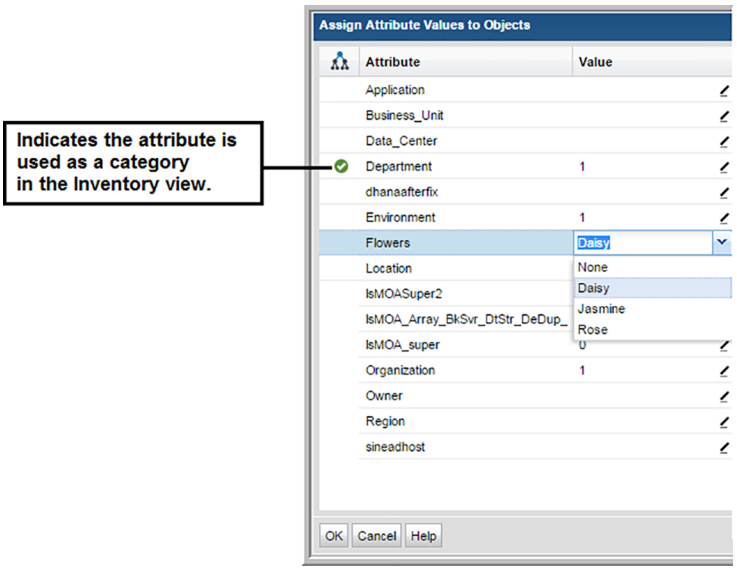


2 Navigate to the object category.

3 Toggle to the **Inventory List** view. You can select all the objects on the page by clicking the checkbox on the top of the management page. Note that only the objects displayed on a single page are selected.

4 Click **Attributes**. The **Assign Attribute Values to Objects** dialog shows if attributes are actively used in the Inventory hierarchy. It also allows you to set or modify an attribute (or multiple attributes) for all the inventory objects selected. If no value is set for an attribute, the field is blank. If different attributes' values are set for the selected objects, the value displays **Multiple Values**.

5 Double-click the **Value** field beside the attribute. A drop-down list is activated.



6 Choose a value to assign to the selected hosts. If the value is displayed as **Multiple Values**, the new value is applied to all selected hosts and will overwrite any previously set attribute. You can set multiple attributes at once.

Set attribute values in the inventory list view

You can select different attribute values or remove them entirely in the **Inventory List** view. You can do this on an individual object or on a set of objects. These instructions cover how to modify attributes in bulk.

To remove attribute values on objects in bulk

1 Select **Inventory**.

2 Toggle to the **Inventory List** view.

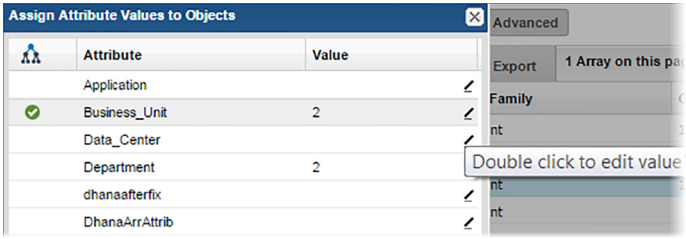
3 Enter your filter criteria to find your target set of objects.

4 Navigate to and select the objects to manage. You can select all the objects on a page by clicking the checkbox on the top of the management page. Note that only the objects displayed on a single page are selected.

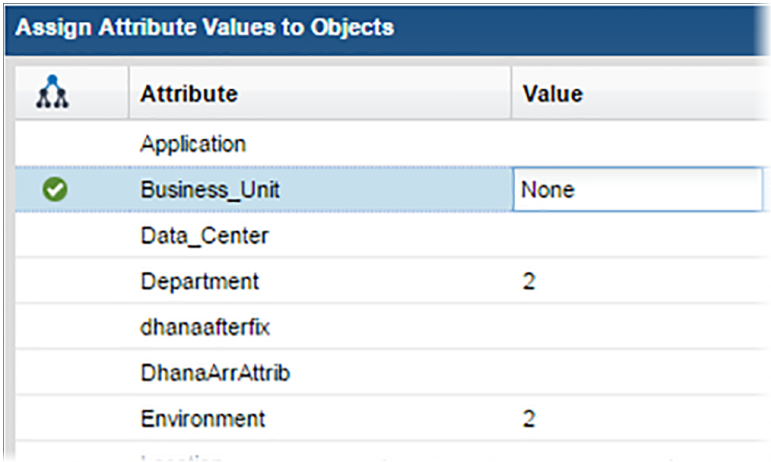
5 Click **Attributes**. The **Assign Attribute Values to Objects** dialog allows you to set or modify attribute values for all the objects selected.

6 In the **Assign Attribute Values to Objects** dialog, double-click an attribute to set an attribute value. If different attributes' values are set for the selected hosts, the value displays **Multiple Values**.

7 Click the **Value** field beside the attribute you want to remove. The field is activated.



8 Select **None** to remove the value. If you selected multiple objects, the attribute is removed from all your selections. You can remove multiple attributes and values at once.



Import host attribute values

When importing Host Attributes, the system will read the prepared CSV file and display a dialog to review the data results before the final import.

Import Host Attributes enables you to:

- assign attributes to existing hosts
- create new attributes if they do not exist
- unassign attributes from hosts

This operation may be scheduled to occur on a regular basis using the portal.properties file. This operation may also be completed using the bulk load utility. As a best practice, use the Portal for the import. See also the *System Administrator Guide* under Attribute Management.

See "To import host attribute values from a CSV file" on page 91.

This video demonstrates importing of host attributes using **Import Host Attributes** dialog box.

http://video.symantec.com/services/player/bcpid292374537001?bckey=AQ~~,AAAABuIiy9k~,I8Bhas-Vwr9zYL9V36WFi86fR_NoePscn&bctid=6301527985001

Prerequisites

- Hosts and attributes must exist.
  See "Adding and Editing Hosts and Backup Servers " on page 521.
- See "Add attributes " on page 541.
- The CSV file must be created. Refer to the following format specifications.

CSV format specifications

Enter the information into a spreadsheet from which you will create a comma-separated values file. The table in the spreadsheet should be in the following format:

- First column - list **Hosts** as the header title. Enter existing host names in the rows. This value must be the unique Host Name, not the value identified as Host Display Name.
- Subsequent columns - list the existing attribute name as the header title. Enter possible values for the attribute as a comma separate list in the corresponding rows.
  See "Examples of attributes and values" on page 551.

You can leave a table cell unpopulated to remove pre-existing values and the attribute will be unassigned from the corresponding Host.

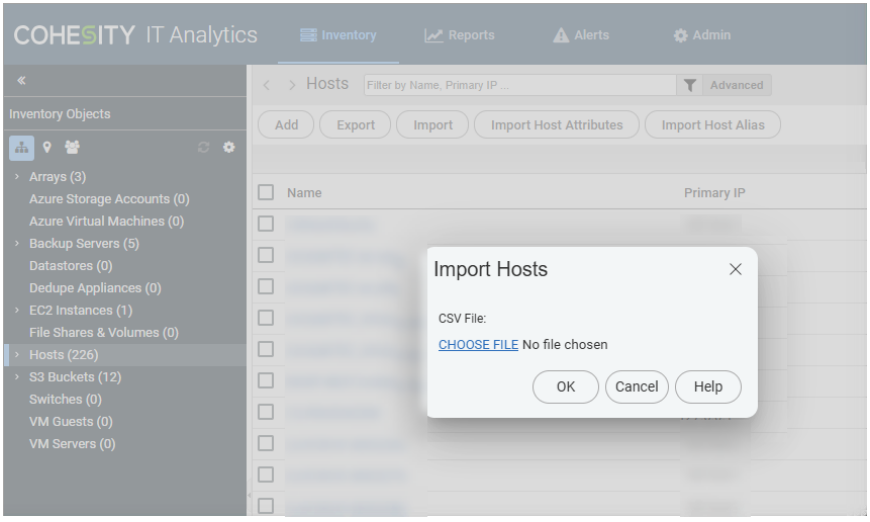
- Only list a host once. If a host name is listed in multiple rows, only the attributes from the last row with the same host name will be saved.

Table 3-1 CSV format specifications

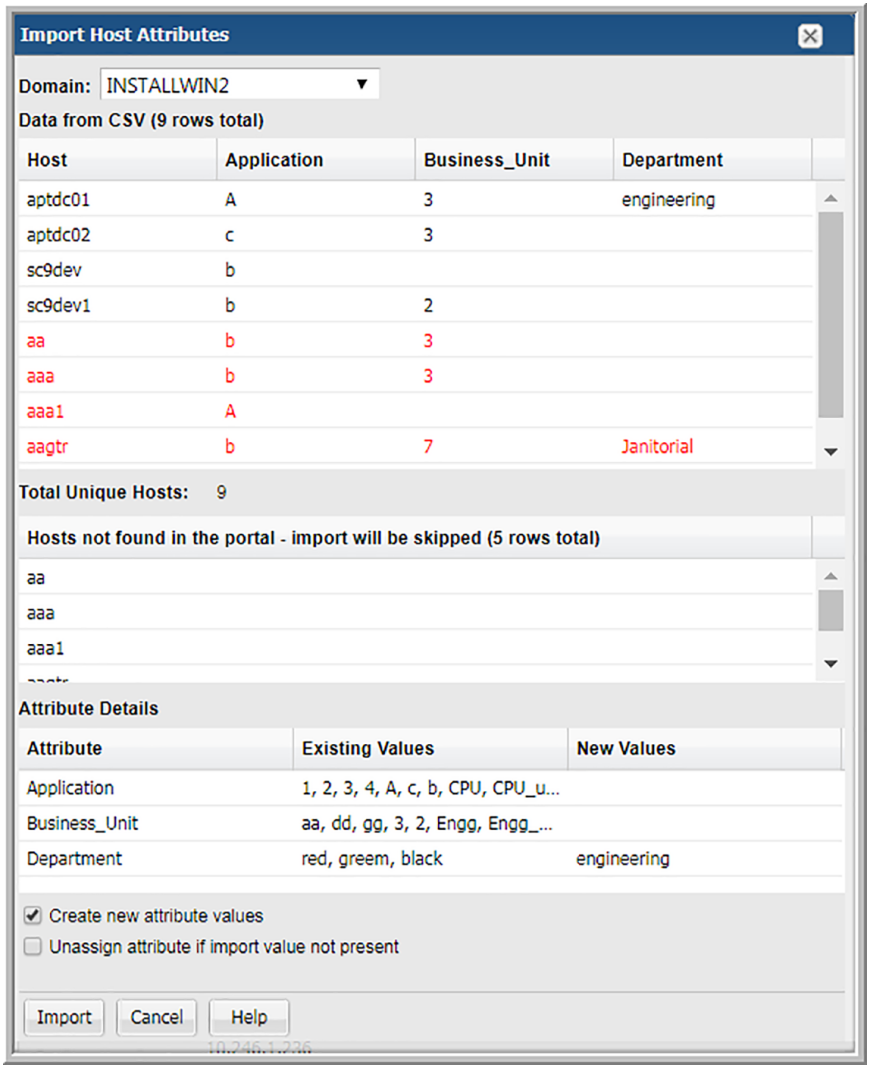
Hosts	Attribute Name 1	Attribute Name 2	Attribute Name X
Host name 1	Attribute value		Attribute value
Host name 2	Attribute value	Attribute value	
Host name X	Attribute value	Attribute value	Attribute value

To import host attribute values from a CSV file

1. Prepare the CSV according the previous format specifications.
2. Select **Inventory**.
3. Verify your hierarchy is set to display by **Inventory Object Type**.
4. Select the **Hosts** folder and set to list view.



5 Click **Import Host Attributes** to select the prepared CSV file. The results of the intermediate import are displayed.



6 Review the intermediate import results:

In the **Data from CSV** table, red entries indicate hosts that are not recognized in the Portal.

In the **Hosts not found in the portal - import will be skipped** table, the list of Host Names read from the CSV, but were not found in the IT Analytics database are displayed. No action is taken if Host Names do not exist in the Portal prior to importing Attributes and their values.

See "Adding and Editing Hosts and Backup Servers " on page 521.

Note: Hosts are listed by the unique Host Name. This is not the Host Display Name.

---

7 **Create new attribute values** is selected by default to use the import to create new values that do not currently exist in the Portal. If this option is not selected, new values are ignored when imported.

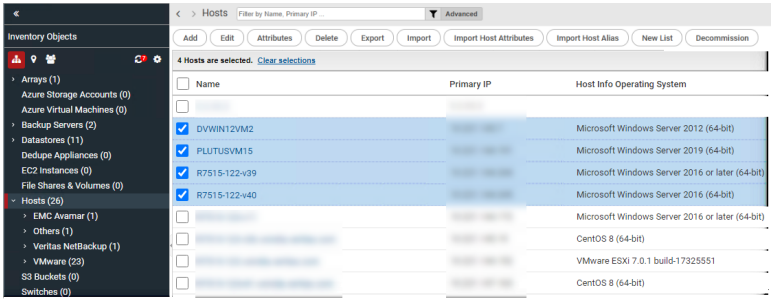
8 Click **Unassign attribute if import value not present** to remove the association between the Host and the Attribute. If the imported CSV file contains unpopulated cells for any attribute values, pre-existing values are removed and the Attribute will be unassigned from the corresponding Host. If this option is not selected, an unpopulated cell in the imported CSV file is ignored.

9 Click **Import** to complete the Import.

Export objects from the inventory list view

You can export selected objects from the Inventory List view and download them to an Excel spreadsheet.

1. Select a category in the Inventory hierarchy panel.
2. Click the **Go to Inventory List icon**.
3. Select an object type, for example **Hosts**.
4. Select the objects to export. Click the top checkbox to select all objects on the page.



5. Click **Export**. The objects and the displayed columns are exported to an Excel spreadsheet and automatically downloaded.

Note: The download status of the report will be displayed in the **Downloads** popup. To see all the download performed during the session, go to the user profile menu > **Downloads**.

Delete objects using the inventory list view

If you choose to delete an object (if you have privileges) you are irretrievably deleting the object from the reporting database. All related historical data will also be permanently deleted from the database and unavailable in all reports. A pop-up window warns you of this action to prevent inadvertent deletions. This is an audited action and requires a password before deletion can proceed.

When you delete a host, it removes everything related to the host except the VM Server. To remove the related VM Server, you just explicitly delete it using the Inventory window. This prevents servers from being orphaned in the database.

In most cases, you should remove a host from a group, thereby un-linking it from its relationship with other machines in that group.

To delete objects using the Inventory

Note: VM Guests, individual LUNs, S3 Buckets, EC2 Instances, Azure Storage Accounts, Azure Virtual Machines and datastores cannot be deleted.

---

1. Select **Inventory**.
2. Select the object type category and expand to view.
3. Toggle to the **Inventory List** view.
4. Select an individual object or you can multi-select.

5 Click **Delete**. Remember all related historical data will also be permanently deleted from the database and unavailable in all reports. Two confirmation dialogs are displayed.

6 Enter your password. This is the same password with which you logged into IT Analytics. This information is captured in the log files.

See "Decommission/Recommission Hosts and Backup Servers" on page 526.

Note: Deleting objects removes them your license count.

Customize columns in the inventory list view

To create a custom view of your data center inventory, choose the columns to display on the page. The columns available for selection are based on your licenses, collected data within the your domain, and pre-built custom attributes. Once selected, sort and move these columns to your preference. Attributes can also be displayed as a column in the **Inventory List** view.

See "Select columns on management pages" on page 48.

Note: The **Host Group** column, displayed in the Inventory for the Host Group management view, has sorting disabled to improve portal performance.

Chapter 4

Get acquainted with reports

This chapter includes the following topics:

- Overview
- Reports and templates
- Cloud reports and templates
- Report library
- Report orientation
- Shared reports and dashboards
- About dashboards
- Work in the Reports tab
- New badge
- About caching
- Sections in the navigation panel
- About badging
- About the action bar
- Reports tab keyboard shortcuts
- Folder types
- Report template icons

- Navigate through reports

Overview

The product includes more than 200 standard report templates and a small number of dashboards, along with the ability to build custom report templates and dashboards. Both built-in and custom reports can provide in-depth views based on the data you specify. You can tailor and save these custom report templates for easy access whenever you require them. You can also create dashboards to provide comprehensive data management tools for your complex, heterogeneous storage environment.

Reports and templates

Over 200 different out-of-the-box report templates plus periodic releases of cloud report templates are offered. Reports are created using pre-built templates that a user can customize to provide relevant, current data. A user can also create their own templates using either the Dynamic Template Designer or the SQL Template Designer. The Dynamic Template Designer is a tool that does not require Structured Query Language (SQL) knowledge to create custom report templates. Users can easily assemble a simple report template by dragging and dropping fields into the template. The SQL Template Designer is considered to be an advanced feature, requiring experience in SQL (Structured Query Language) query development. Using the SQL Template Designer and your knowledge of SQL, you can build advanced reporting solutions to support efficient storage resource management.

Reports, dashboards and templates function similarly, but the functions available to each object type varies. For example, if a report or dashboard is saved, you can share it, delete it and create shortcuts to it within your folder structure. Some of these options are not available for an unsaved ad hoc report. It is important to note, when you save a report, you are saving the definition and parameters, not the report output.

Reports always have a template associated with them. You can quickly see this connection on the **Reports** tab. This allows you to map back to the original template if you'd like to make changes. To help you organize templates and reports, you can create shortcuts to multiple templates and reports within your **My Reports** folder.

IT Analytics periodically publishes new user report templates and makes them available through the **Home** section in the **Reports** tab. This provides instant access to new and improved reports without waiting for a formal software release. System report templates are automatically added to the appropriate system folder and made available to the Super User. If the Super User chooses to update the report template, they can be made available to users through their privileges.

Watch this video explaining the process of executing the out-of-the-box report, saving that report in a template, and finally adding that report template to newly created dashboard.

http://video.symantec.com/services/player/bcpid292374537001?bckey=AQ~~,AAAABuIiy9k~,I8Bhas-Vwr9zYL9V36WFi86fR_NoePscn&bctid=6301528520001

See "About badging" on page 102.

Cloud reports and templates

IT Analytics periodically publishes new user report templates and makes them available through the **Cloud** section in the **Reports** tab. If you have the privilege enabled, you can run these reports directly from the Cloud, preview the data and if you like it, save the report locally. (The template is automatically saved locally when you save the report.) You can quickly identify these new or updated user report templates by their **New** or **Updated** badge.

See "About badging" on page 102.

Once saved, Cloud templates work exactly like out-of-the-box templates, except they are linked to the Cloud, so when improvements are introduced, the Cloud section is automatically refreshed and your local Cloud-linked templates can also be updated if you choose to do so. However, if you saved a Cloud template with a different name or made a copy, that saved template is detached from the Cloud template and will not automatically inherit the updates.

The **Cloud** section and its templates are available if you have the privilege enabled.

See "About user privileges" on page 572.

Proxy connections are supported.

Note: If you do not have an internet connection, you can still benefit from Cloud-based templates. Contact Professional Services for details.

Report library

The Report Library provides visibly into your IT environment through IT Analytics product reports such as:

- Chargeback metering platform helps manage costs and demand for resources
- Auditing and compliance reporting will get you the data viewing for every job
- Utilization and optimization reporting for detailed performance

- Trending and forecasting gives you the visibility on your data usage and storage consumption

The library allows you to filter by product or to search for the reports you need. There are at least 300 reports to choose from with new reports added regularly. You can use these directly or easily make changes to accommodate their unique environment. You can find the library on the website:

https://reportlibrary.veritas.com/

You can also subscribe using an RSS feed to stay on top of the latest additions.

Note: Some reports may not be suitable for large environments or MSP portals.

After downloading, you can import the report into your installed Portal.

Watch this video on leveraging the OpsCenter reports and process to import the report in IT Analytics

http://video.symantec.com/services/player/bcpid292374537001?bckey=AQ~~,AAAABuIiy9k~,I8Bhas-Vwr9zYL9V36WFi86fR_NoePscn&bctid=6301528424001

Report orientation

Reports are designed to enable views of high-level information, with drill down access to details in sub-reports.

- See “Report template icons” on page 105.
- See “Interact with Reports” on page 159.
- See “Understand report data caching” on page 839.
- **Tables** display data in rows and columns, with highlighted links that enable drill down access to additional details. When you click a link, a new report is generated in a new pane with a tab displayed at the top of the Content pane. These detail reports can have their own unit of measure - separate from the parent report settings. For example an Array and Capacity Utilization Summary report shows aggregated data in TiB.When you drill down, it may be more useful to switch to MiB.

  Tabular data can be sorted in ascending or descending order by clicking the column headers. Some tables use color to highlight potential problems. For example, the Host Utilization Detail report highlights values when they fall below the availability threshold.

- **Charts** use colored bar graphs and pie charts to enable at-a-glance analysis. Multi-colored bars can be deciphered with a few simple moves of your mouse. Move your cursor over each colored section of a bar or pie to view a text description of what that portion of the bar/pie represents. For example, in the Job Status Summary report, a single bar will reflect the number of event failures, the number of warnings, and the number of successful backups. In addition, some charts display a trend line graph with data points represented as percentages. If you move your cursor over a particular data point arrow, the performance percentage value will be displayed.

Shared reports and dashboards

You can share reports and dashboards that you create with other users and user groups. You can also share folders. This allows you to distribute data to other users in your business while maintaining control of the report. Sharing is controlled by permissions and can be retracted as required. You cannot share a report that you have not saved.

See "Share Reports, Dashboards, and Folders" on page 189.

About dashboards

Dashboards provide an at-a-glance overview by displaying many reports on a single page. You can drag and drop individual reports, to create a consolidated view into your storage environment. You can create a dashboard to contain reports you access regularly, or view multiple scenarios simultaneously. Dashboards can also help you troubleshoot a particular problem by gathering data from multiple reports into one view.

When you design a dashboard, you gather a variety of reports that will help you to correlate events, isolate problems, and forecast capacity. You simply create a new dashboard and add reports to it. There is no limit to the number of reports that you can add to a personal dashboard. Each time you add a report to a dashboard, the reports appear on the same page.

Note: Reports on dashboards are treated as stand-alone instances and are not linked to the original report. For example, when a name or scope change is made in the original report, those changes are not reflected in the dashboard version.

You can save Dashboards that you create and access them like any other report.

See "Work with dashboards" on page 196.

See "Manage My Home Pages" on page 211.

Work in the Reports tab

The **Reports** tab allows you to browse, view, run and manage Out-of-the-Box report templates, as well as your own custom reports. A management window for your reporting inventory that allows you to browse, manage, create and customize templates, reports and dashboards. This tab provides access to designers that enable custom report template creation: the Dynamic Template Designer and the SQL Template Designer, an advanced feature, requiring experience in SQL (Structured Query Language) development.

This area provides a single, consolidated view with relevant details and the ability to navigate to originating files. As you manage templates and reports, it's important to understand how they are related to other reports and how they are shared with other users. You can quickly assess this information from the **Reports** view panel.

Different navigation methods are available to slice and examine your collected data. You can explore the data by using parts of your IT infrastructure as the entry point or by using the customizable templates to provide a clear picture into your storage environment. The **Inventory Navigator** serves as browser mode for your infrastructure.

See "Exploring your inventory" on page 60.

New badge

New or updated report templates are periodically published and made available automatically. Folders may be badged with a **NEW** flag to indicate that new or updated report templates are available in your system.

See "About badging" on page 102.

About caching

When you run a report, the scope of the report is examined, and then the product checks if the cache contains the same report, for the same scope. If it does, the results are displayed from the cache. If the combination does not exist, the report is run from the database, saved in the cache, and then sent to the user interface. Cached reports are shared across users who belong to the same home host group. When a report is served from the cache, an indicator icon is displayed on reports and dashboards. You can roll over the indicator to show the age of the report from the cache. Click the icon to purge the old report from the cache and rerun the report from the database.

See "How reports and caching work together?" on page 839.

Sections in the navigation panel

The left navigation panel is divided into the following sections:

- **Home** contains the following folders:
  - **Recent** - Displays a list of the 20 most recently run, scheduled in a background run, or modified reports within the last 30 days as limited by your login. Note - if a report has been run and then your access to the report changes, for example through revoked sharing, that report is not displayed in the recent list.
  - **My Reports** - Contains folders that you have created, and the system report template catalog. System report templates refer to general product related report templates. User report templates refer to reports that users have either created using the Dynamic Templates Designer, the SQL Templates Designer or saved from system templates. These folders can also be identified by color. Blue folders contain items that can be edited by users. Yellow folders contain system report templates.
    See "Folder types" on page 104.
  - **Shared by You** - Contains folders or reports that you have shared with other users and user groups.
  - **Shared with You** - Contains a list of reports, dashboards, and templates that other users have shared with you. As the report/folder reader, you are limited in the actions you can perform on the report or folder. Shared reports can be emailed, exported or alerted on. You cannot edit or delete shared reports or folders.
- **Cloud** - Contains cloud-based report templates. New or updated report templates are periodically published and made available through the cloud. If this privilege has been enabled for your Portal account, you can view the Cloud section, access, run and save these as they become available in the **Reports** tab. Proxy connections are supported for Cloud reports. The **Cloud** section and its folders may be badged with a **NEW** flag to indicate that new or updated user report templates are available.
  See "About badging" on page 102.
  See "File list export" on page 806.

About badging

IT Analytics periodically publishes new or updated report templates and makes them automatically available within the product. There are two types of report templates, system and user. System report templates refer to general product related report templates. User report templates refer to reports that users have

either created using the Dynamic Templates Designer, the SQL Templates Designer or saved from system templates. You can quickly identify system versus user folders by color. User folders are blue. System folders are yellow.

The **Home** section and its folders may be badged with **NEW** icons to indicate that new or updated report templates are available. Badges on user report folders and templates are visible to every user. Badges on system folders and report templates are only displayed to Super Users.

The **Cloud** section and its folders may be badged with a **NEW** icon to indicate that new user report templates are available. This badging is visible to every user with Cloud privileges. The badge is dismissed when a copy is made or saved as a user report template.

Home section

In the **Home** section, badges are displayed on product folders indicating when a template change has been introduced. Badges on user folders are shown only for Cloud templates, that is, those linked to templates in the cloud by users. Badges serve as a visual alert that new or updated report templates are available. For new system templates, the Super User must then enable new reports in the privileges for individuals or user groups.

See “Enabling new product report templates” on page 586.

User report templates copied from the **Cloud** section into the **Home** section (**My Reports**), will also display **UPDATED** badges if revised versions of a user report template are available from the Cloud.

UPDATED badges are dismissed when you choose to update. **NEW** badges automatically expire in 14 days. You can customize the display duration for the **NEW** badge in the portal.properties file.

Cloud section

The **Cloud** section in the **Reports** tab will display a **NEW** badge to any user with Cloud privileges if a new user report template is available. The badge is dismissed when you copy the report template into your blue user folders or it is saved as a user report template.

- See “About user privileges” on page 572.
- See “Use new and updated System Report Templates (Super User only)” on page 163.

About the action bar

On the **Reports** tab, the element you select, such as a folder, a report, a template, or a shortcut, determines what operations are made available. For example, a report can be deleted, renamed and shared. A system report template cannot. To make this easier, the Action bar changes when an element is selected, showing you what operations are available based on your selection. Operations are also available through a right-click menu.

Reports tab keyboard shortcuts

The **Reports** tab supports standard keyboard shortcuts to work more efficiently. Delete is available on other grid-based tabs in the Portal. Shortcut operations are controlled by Privileges.

The following keyboard shortcuts are available on the **Reports** tab:

Operation	Keyboard Shortcut
Copy	Ctrl+C
Cut	Ctrl+X
Delete	Del or Delete
Paste	Ctrl+V
Select All	Ctrl+A

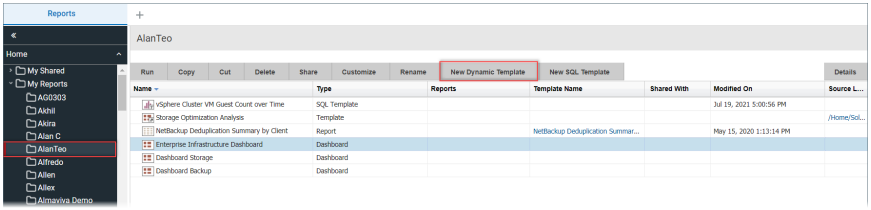
Folder types

In the **Home** section of the **Reports**, there are two folder types: system and user. The folders operate similarly, but the functions available with each type differ.

System folders are static and you cannot modify them. System folders contain product specific report templates. User folders can be created, modified, deleted, and shared. User folders contain reports that users have either created using the Dynamic Templates Designer, the SQL Templates Designer or saved from system template. These folders can also be identified by color.

Both System and User folders may display **NEW** badges when new or updated report templates are available. Badge display is controlled by privileges and roles.

As you select a folder, the available functions are displayed on the action bar.



Report template icons

Icons are used to represent report template types. You can quickly determine if a report will be produced as a dashboard, table, pie chart, donut chart, line chart or bar chart by the icons and roll-over help. These icons may display stamps to convey additional information. For example, shortcuts can be distinguished from the original file by the arrow that appears on the icon. This allows you to quickly scan lists of reports and make business decisions based on visual indicators.

Table 4-1 Table 1 Report Template Icon Definitions

Icon	Definition
	Tabular Chart - Displays data in rows and columns, with highlighted links that enable drill down access to additional details.
	Dashboard - Provides an at-a-glance overview by displaying many reports on a single page. This consolidated view can contain bar charts, pie charts and tabular reports.
	Pie Chart - Provides a visual representation of data in a circle divided into sectors which is proportional to the quantity it represents. Mouse-over sectors for details and drill down access.
	Donut Chart - Provides a visual representation of data in a circle divided into sectors which is proportional to the quantity it represents. Mouse-over sectors for details and drill down access.
	Bar Chart - Displays colored bar graphs to enable at-a-glance analysis. Multi-colored bars can be interpreted by using data displayed as you move your mouse over. This icon applies to all available bar charts such as stacked, horizontal and so on.
	Gauge Chart - This is a single value axis chart. Gauge charts have no x-axis. Normally, this displays a maximum value with a series of data point values across the gauge. Color can be set to indicate different sectors.
	Line Chart - Represented by a series of data points connected with a straight line. Line charts are most often used to visualize data that changes over time.
	Area Chart - Depicts a time-series relationship. Unlike line charts, they can also visually represent volume. Information is graphed on two axes, using data points connected by line segments.

Table 4-1 Table 1 Report Template Icon Definitions *(continued)*

Icon	Definition
	Shortcut Badge - Displays an arrow as a badge on the report rendering type icon. A shortcut is a link to a report or a template in a different folder.
	Template Badge - Displays a T as a badge on the report icon. Reports and report templates function differently. This badge identifies a template.
	Cloud Badge - Displays a cloud as a badge on the template icon. This badge identifies that the report originated from and is linked to the Cloud section.

Navigate through reports

Use Search to traverse names for all inventory, reports, dashboards and templates. You can narrow your search by selecting a specific category.

Use different navigation methods to slice and examine your collected data. You can explore the data by using parts of your IT infrastructure as the entry point or by using the customizable templates to provide a clear picture into your storage environment. The Inventory Navigator serves as browser mode for your infrastructure.

See “Exploring your inventory” on page 60.

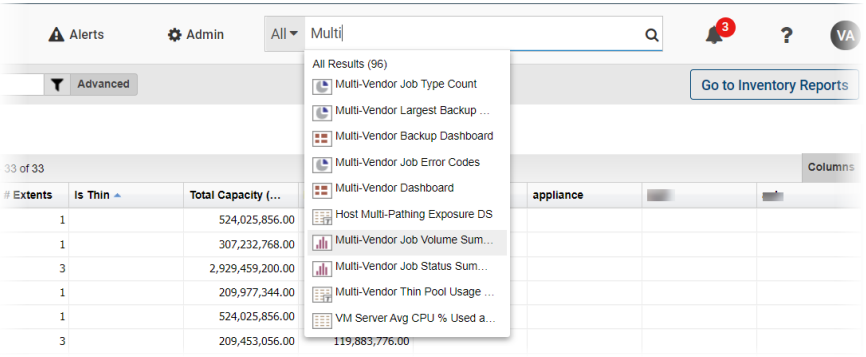
The platform includes more than 200 standard report templates, periodic releases of cloud reports, along with the ability to build custom reports. Using search is the most efficient way to navigate through the reports. If you know the name of a report, a template or dashboard use Search to find it. If you can’t remember a report name, you can click the **Recent** section in the navigation panel. This displays a list of the 20 most recently run, scheduled in a background run, or modified reports within the last 30 days as limited by your login.

See “About the Admin tab” on page 55.

To Find a Template, Report or Dashboard

1 Enter the name in the **Search** field.

As you type, up to 10 potential matches are displayed. These are string matches. If you locate the report or dashboard you require, you can select and run it directly from the match list. You can also click All Items in the match list to go directly to the search results.



2 Click the **Search** icon.

The search results are displayed on the **Search** page. Results are listed under categories, dashboards, reports and templates, and can be expanded to reveal additional reports. You can filter these results if required.

Click the link in the **Source Location** column to navigate to the source report or template. You can edit or customize the report if applicable.

3 Click **Details** to toggle on the view panel to read a long description for a report or template. (This is only available for Backup Manager and Capacity Manager).

4 Double-click a template to display the scope selector. This allows you to set the parameters before running it. Double-click a report or dashboard to run it.

To Refine Your Search Results

On the search results page, you can refine search results based on **Keyword, Product, Category and Source**. By default, all reports available to your user role, and those that are shared with you are displayed - your entire reporting inventory. Filtering these results allows you to quickly locate what you require. If you would like more information about a report before running it, select it, click Details, and a long description is displayed.

1. Select a filter and the results automatically update based on the selection.
2. For **Keywords**, enter a name or description in the field, and select **Name** and/or **Description**.
3. Double-click to run the report.
4. Click the link in the **Source Location** column to navigate to the actual report. You can edit or customize the report if applicable.

Chapter 5

Generate and maintain reports

This chapter includes the following topics:

- Report Scope
- Select Report Scope
- Configure report scope with attributes
- Group hosts by attributes
- Search for hosts in the report Scope Selector
- Alert scope selector settings
- Amazon Web Services (AWS) scope selector settings
- Backup Manager scope selector settings
- Backup Manager advanced scope selector settings
- Capacity Manager scope selector settings
- Data Collection Status Details scope selector settings
- Host scope selector settings
- Ransomware Scorecard scope selector
- Solution reports scope selector settings
- System health check scope selector
- Generate reports

- Generate reports from the Cloud
- Units of Measure in Reports
- Advanced Filtering for Tabular Reports
- Customize report filter logic
- Sort columns in reports
- Refresh and set refresh intervals
- Modify Reports
- Convert tabular report to chart
- Add and remove custom header and footer to a report or dashboard
- Set or edit the drill down preference
- Access NetBackup web interface from the IT Analytics Portal
- View report statistics
- Interact with Reports
- Work with Topology
- Add a note to a Backup Job Details Report
- Use new and updated User Report Templates
- Use new and updated System Report Templates (Super User only)
- Save Reports
- Save Cloud Reports
- Delete Reports

Report Scope

Report scope objects are grouped into a hierarchical format, enabling disparate groupings. For example, hosts can be grouped by vendor, geography, department, and so on. This object grouping also provides additional data partitioning (security) by restricting a user's view to hosts under the home host group. In addition, various built-in processes organize hosts based on information available from backup products; for example, group by master server or group backup servers by policy.

When you generate a report, a database query is initiated, based on the report scope that you specify. Since the nature of your IT environment is dynamic--for example, it's not uncommon to add host groups and hosts to your network--reports reflect the updates.

Before generating a report, you can select enterprise objects such as the storage arrays, hosts or host groups, fabrics, zones, or switches to include in the report scope.

- **Host Groups**: Host groups primarily serve as a means of enforcing security controls, limiting a user's access to only the hosts in a user's IT Analytics domain. In the context of report scope, host groups are relevant for backup, host capacity, File Analytics, and virtualization management reporting. In addition, several related sub-categories can be selected to narrow the report scope, depending on the type of data the report supports: Devices, Shares, Volumes, and Deduplication/VTL Appliances. When selecting the host groups in the Scope Selector, the selections are displayed in the bottom display panel. **Attributes** provide an alternate grouping mechanism for a report scope.
  See "Host groups vs attributes" on page 537.
- **Hosts**: Select specific hosts for a static report. A static report does not take into account changes in your network topology, so if you add clients to a network, you either have to explicitly include them in your report scope or add them to a host group. To ensure that host are always included in reports, assign them to a host group and select host groups for your report scope. For backup reports, a variety of components can be selected to be available in the Report Template's scope selector; for example, Consecutive Errors, Ignore Retries, and Backup Window. For certain host reports, additional options may be available to filter the scope of the report, such as OS Platform (such as Windows or Linux), Product Collected (find hosts that have been collected from a particular subsystem, such as Veritas NetBackup), and Product Not Collected (find hosts that may be unprotected because they have not been collected from a particular subsystem, such as EMC Avamar). Reports list the full path of the host group in report headers. When selecting host groups in the Scope Selector, the selections are displayed in the bottom display panel.
- **Backup Server**: You can select backup servers such as NetBackup Primary server, in the report scope, for Backup Manager reports.
- **Host Type**: Indicates how a host has been commissioned in an enterprise, such as VM Guest, VM Server, VIO Guest, VIO Server, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-V Guest or Others.
- **Datastores and Hypervisors**: You can select datastores and hypervisors in the scope dialog box for Virtualization Manager reports.

- **Clusters**: You can select clusters in the scope dialog box for Virtualization Manager reports. Clusters refer to the way ESX servers are grouped.
- **Arrays**: You can select any of the following in the scope dialog box for Capacity reports:
  - Storage array
  - Array family
  - Array vendor
  - Array product
- **SAN Fabrics (including Zones and Switches)**: Select specific SAN Fabric objects to narrow the scope of a report.
- **Attributes**: Attributes associated with objects such as hosts, storage arrays, libraries, drives, switches, host Oracle database, and host MS Exchange, enable logical groupings for reports; for example, OS version could be an attribute that you want to associate with hosts. Attributes enable you to define a data set based on a specific characteristic.

See " About attributes " on page 536.

See "Host groups vs attributes" on page 537.

Use the following guidelines to ensure that your reports include the data you expect to see:

- If you add new host groups after you generate a report, the next time you generate this same report, it will include different results. When you scope by host group or attributes, a dynamic report is produced.
- If you add hosts to your network, but do not assign them to a group, then generate a report without including the new hosts in your scope, you can be certain that the data for that host will not be in the report. This type of report scope produces a static report. However, if you add the host to a host group and this group is part of the report scope, the report automatically includes report data for that host.

The following are the scope selection configurations for various reports:

- See "Alert scope selector settings" on page 121.
- See "Amazon Web Services (AWS) scope selector settings" on page 122.
- See "Backup Manager scope selector settings" on page 123.
- See "Backup Manager advanced scope selector settings" on page 126.
- See "Capacity Manager scope selector settings" on page 128.
- See "Data Collection Status Details scope selector settings" on page 129.

- See “Host scope selector settings” on page 130.
- See “Ransomware Scorecard scope selector” on page 131.
- See “Solution reports scope selector settings” on page 132.
- See “System health check scope selector” on page 137.

Select Report Scope

You can create static reports or dynamic reports. To create static reports, define your report scope based on a list of individual enterprise objects, such as hosts or arrays. To produce dynamic reports, define your report scope based on domain, host groups, and attributes.

To learn about how report scope selects objects. Refer to the following.

See “Report Scope ” on page 110.

See “Advanced Filtering for Tabular Reports” on page 141.

Report scope operates with the following logic:

- Report scope is restricted to objects of the same type, such as arrays and array vendors, but not host groups.
- Detail reports have limited report scope capabilities. For reports that are the result of drilling into details from a parent report, those detail reports can have their own unit of measure, separate from the parent report settings. Once you’ve drilled into a child report, you can edit the report scope and use the Advanced option to change the measurement unit to see more accurate or easily understandable information. For example, an Array Capacity & Utilization report shows aggregated data in TiB. When you drill down to a detail report, it may be more useful to switch to MiB.

Time Period Considerations

- Certain reports, such as the Array Performance by RAID Group, query database tables for daily log data. When choosing a scope of less than 24 hours and no data is returned, make sure that the report scope time period includes the 12 a.m. midnight boundary.
- For Backup Manager reports: The time period used to retrieve a report’s data takes into account the time zone of the collected systems, if relevant. For example, if 15 minutes is selected for a report that has a backup server in Hong Kong, but the report is being generated in San Francisco, the time period reflects the 15-minute interval and the end time will be the end of the Hong Kong server’s

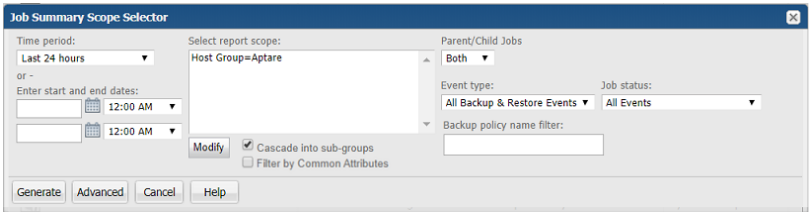
day. If there is no data collected from a different time zone, the Portal's local time is used.

- For non-backup reports and SQL Template Designer reports: The precise start and end times for the time period will be displayed, without consideration of time zones.
- For data collection reports: you can select the Run Time either for scheduled or on-demand collection to narrow the scope. Select from Last Run up to the Last 7 Days.

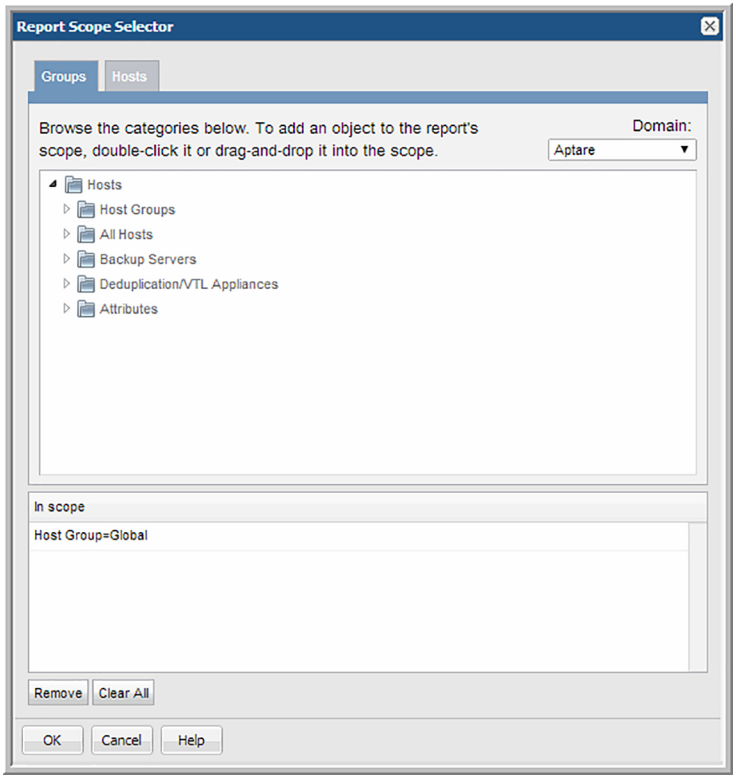
To select report scope

1 Search for a report to generate. For example: Job Summary

The Job Summary Scope Selector allows you to specify parameters, report criteria and generate a report as shown in the following example:



2 Click **Modify**. The **Report Scope Selector** window launches.



The tabs in the Report Scope Selector vary based on the report type. This window enables you to:

- Expand the list of available objects.
- Double-click to add an object to the scope - Report scope is restricted to objects of the same type, such as arrays and array vendors, but not host groups.
- Double-click to remove an object from the scope - This removes an object that has been place in the report scope.
- Drag and drop objects into the scope - Click the object and move it until you see a red dotted rectangle. Drop it into the In scope pane.
- Search for objects to add to the scope-

See “Search for hosts in the report Scope Selector” on page 120.

3 For reports with a host in the report scope, you can:

- Create a dynamic report. In the **Groups** tab, select the group of hosts that you want to include in your report:
- Search for specific hosts.
  See “Search for hosts in the report Scope Selector” on page 120.
- If you want to begin by including all hosts, select the entire domain--the top-level host group.
- If you want to begin by including specific hosts within a host group, select the check box for that host group.
- Create a static report. Select a list of individual hosts that you want to include in your report:
- In the **Groups** tab, drill down to each client’s host group and select the client’s check box; or
- Go to the **Groups** tab and create a list of hosts based on attributes.

4 Select attributes from the **Attributes** folder to include objects that have the attribute value assigned to them. Attributes provide a way of defining a set of objects that share a certain characteristic. When you include an attribute in a report’s scope, you include all objects that have this attribute value.

  See “Host groups vs attributes” on page 537.

5 Click **OK**.

Configure report scope with attributes

Attributes enable you to define a set of data to populate reports. In addition, attributes provide flexibility for categorizing data. For example, you may want to organize hosts by location and business unit. For more information on how to create and manage attributes refer to the following.

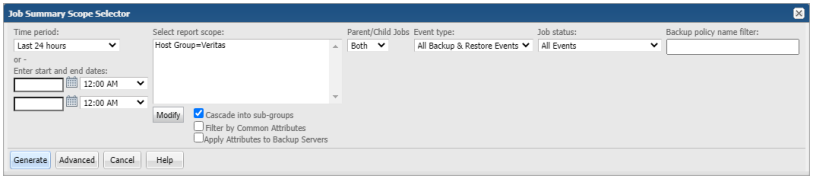
See “ Set attributes on hosts ” on page 540.

See “Host groups vs attributes” on page 537.

To configure scope using attributes

1 Search for a report or report template to generate the report. For example: Job Summary.

The Job Summary Scope Selector allows you to specify parameters, report criteria and generate a report as shown in the following example:

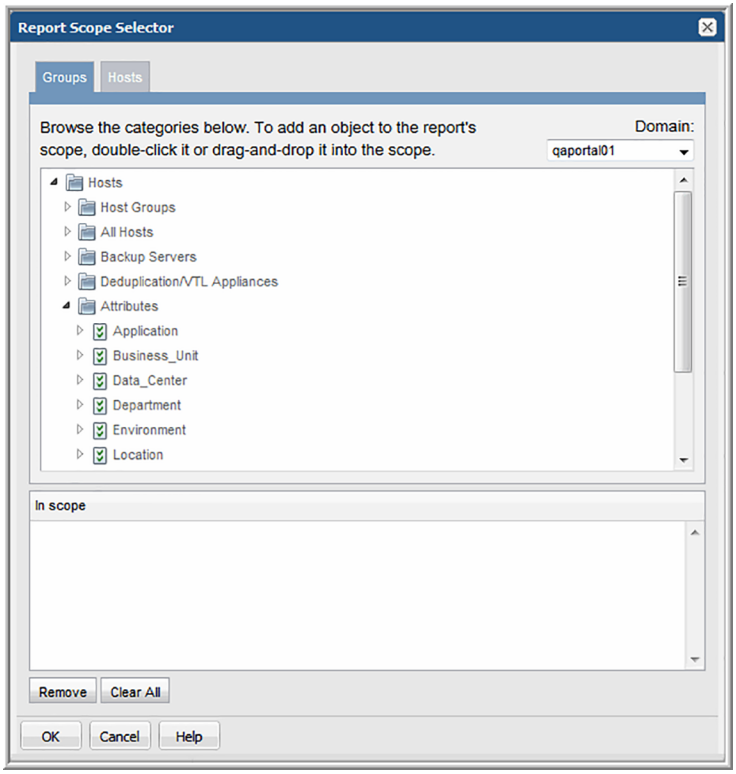


2 Select **Filter by Common Attributes** to define the scope of the report more specifically. For example, if you select attribute values, Location: Campbell, Department: Engineering, Business_Unit: Cost Center 1 and select **Filter by Common Attributes**, the report will display only the results that contain all 3 attribute values. If you do not select **Filter by Common Attributes**, the report will display all results with attributes values Campbell, Engineering, or Cost Center 1.

3 Select **Apply Attributes to Backup Servers** to apply the attributes only to the backup servers, instead of hosts.

4 Click **Modify**. The **Report Scope Selector** window launches.

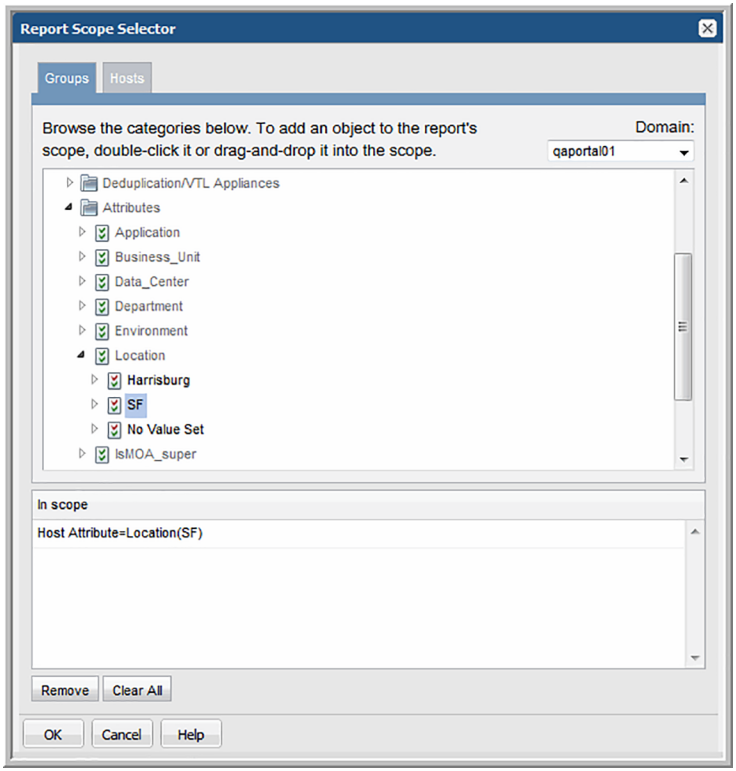
5 On the **Groups** tab, expand the Hosts folder to view the **Attributes** folder.



6 Expand the Attributes folder to view the list of attributes that have been created for your environment.

7 Drag the attribute values into the “In scope” pane.

For example, find all hosts where the attribute, Location, is set to SF.



Hosts have additional special considerations.

See “Group hosts by attributes” on page 119.

Group hosts by attributes

Because many environments have thousands of hosts to manage, it is often necessary to group hosts in a variety of ways to efficiently manage them. Use host attributes to add additional properties to a host to generate reports for discrete sets of hosts. Conversely, you may want to discover what hosts do not have an attribute set. This is also an option. By selecting the attribute No Value Set enables you to produce a results set without a set attribute value.

See “Host groups vs attributes” on page 537.

Host attributes examples

- Show all hosts owned by Marketing.
- Show the hosts in Europe that belong to Sales.
- Show hosts by location, business unit, and application.

Find hosts without an attribute value set

- Show all hosts that have not been tagged with a Location value. This enables you to find hosts that may not be grouped correctly for host management.
- Show all hosts that have not been tagged with a Business Unit.

Search for hosts in the report Scope Selector

When you are selecting a report scope, you may need to find and verify details about a particular host.

In the scope selector, the **Modify** button enables specific scope selection. Often the names of specific objects, such as arrays or switches, are not easy to find. The following search tabs are available (depending on the type of report):

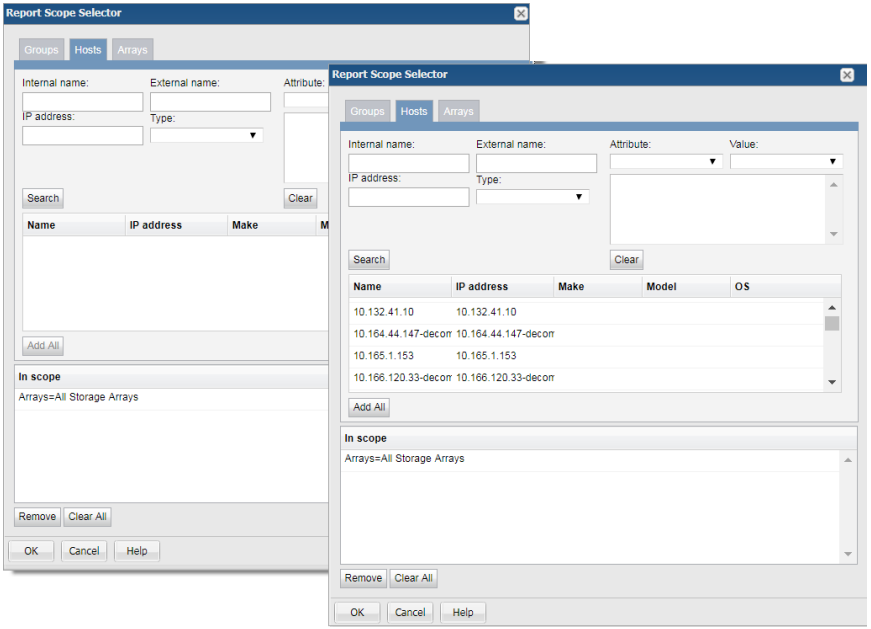
- Hosts
- Arrays
- Switches

Once the Search results are returned, you can add individual hosts to the scope by double-clicking to move them. Click Add All to move all the search results into scope. Report scope is restricted to objects of the same type.

Find hosts example

For reports that can be filtered by arrays and host groups, there is a search capability, as shown in the following example.

A similar search facility enables criteria selection for isolating a set of switches for a report’s scope. Enter or select the search criteria. Note that wildcards are supported. For example, Bi* yields a list of all objects/elements that begin with Bi.



Alert scope selector settings

The following selections are also available with the Alert reports, where relevant.

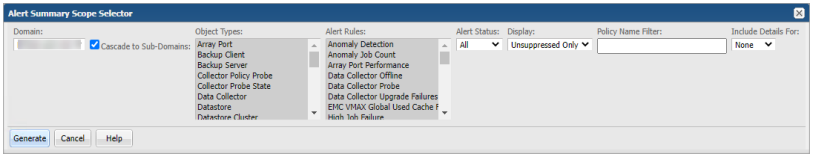


Table 5-1 Alert scope selector settings

Field	Description
Object Types	Use the drop-down list to specify the object types for the alerts. Use Ctrl+click to select multiple object types.
Alert Rules	Choose to filter the results by the Alert Rule. Alert Rules are determined by your product installation. For example, if you do not have Backup Manager installed, the Alert Rules for backup are not displayed.

Table 5-1 Alert scope selector settings *(continued)*

Field	Description
Display	Choose the filter for the Alert Detail, show All alerts, Unsuppressed Only or Suppressed Only.
Alert Status	Choose to filter the results by Alert Status. Select from Warning or Critical.
Policy Name Filter	Enter a value to filter the results by Alert Policy Name. This field accepts wildcards (*).
Include Details	Defines the extent of details to be displayed for each category of the chart: ■ **None**: No details are displayed. ■ **First 5**: The first 5 rows in each category are displayed. ■ **First 10**: The first 10 rows in each category are displayed. ■ **First 20**: The first 20 rows in each category are displayed. ■ **First 50**: The first 50 rows in each category are displayed. ■ **All**: All rows in each category are displayed.

Amazon Web Services (AWS) scope selector settings

The following selections are available in AWS reports, where relevant.

- Accounts
- EC2 Instances
- S3 Buckets

Group By
: Use the drop-down list to specify how you want the data in the bar charts or tables to be grouped. The available options are dependent on the type of report. Some reports group by time, while others group data by accounts, EC2 instances, or S3 buckets.

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.

Backup Manager scope selector settings

The availability of selections, options, and default values varies from report to report because not all selections and options apply to all reports.

Table 5-2 Scope selector field descriptions

Field name	Description
Time Period	The reporting time periods for this report. Your choices range from the last 5 minutes to the last 10 years. The options for this parameter depend on the report. For Backup Manager reports, the time period used to retrieve a report's data takes into account the time zone of the collected systems, if relevant. For example, if 15 minutes is selected for a report that has a backup server in Hong Kong, but the report is being generated in San Francisco, the time period reflects the 15-minute interval and the end time will be the end of the Hong Kong server's day. If there is no data collected from a different time zone, the Portal's local time is used.
Start Date / End Date	If you want a static report, specify a start and end date in the date format: **MM/DD/YYYY**. If you want a dynamic report--that is, a report that changes as the underlying data changes--do not specify a start and end date, as the default always represents the current day.
Hosts	Use this filter to specify what data will be included in a report. ■ Host Group. Includes data from default and user-defined host groups. ■ Attribute. Includes data from default, system and user-defined attributes or characteristics. ■ Hosts. Includes data for only specific hosts rather than data for host groups. See "Report Scope " on page 110.
Group By	Use the drop-down list to specify how you want the data in the bar charts or tables to be grouped. The available options are dependent on the type of report. Some reports group by time, while others group data by either host, client, or policy.
Parent/Child Jobs	Select from Parent, Child or Both. If a parent-child relationship exists, you can select to append Backup Manager job details with the details of parent or child jobs.

Table 5-2 Scope selector field descriptions *(continued)*

Field name	Description
Job/Event Type	Job types: ■ All Backup & Restore Events ■ All Backup Events ■ Full Backups ■ Incremental Backups ■ Restores ■ Unknown Events
Show Capacity in	Select the unit value to display the capacity in. Choose **Dynamic** to allow for the dynamic calculation of the units in the report based on size. If you choose **Dynamic** as an option and run the report, the data value will be displayed as follows: ■ If the value in the report is greater than say 1024MB then show in GB ■ If the value is greater than 1024 GB show TB
Media Status	The Media Status parameter is included in the Tape Media Summary Scope Selector. You can specify which media should be included in the report by selecting one of the following options: All Media, Expired Media, Available Media, In Use Media or Full Media. The default setting for this parameter is **All Media**.
Job Status	Applies to all backup event-driven reports. Select from the following event status options: ■ All Events ■ Successful Events ■ Warning Events ■ Successful or Warning Events ■ Failed Events

Table 5-2 Scope selector field descriptions *(continued)*

Field name	Description
Backup Window	Select a custom backup window to be applied to the report. Typically, backups begin at the end of the business day, but they do not finish before the end of the day--thereby skewing the success statistics for the day. To more accurately reflect backup SLA metrics, you can re-define a day with a custom backup window. These custom backup windows are defined by the IT Analytics Application Administrator. To create specific backup windows, See "About custom backup windows" on page 607.
Backup Policy Name Filter	This filter appears in the Job Summary and Job Volume Summary scope selectors. Enter a backup policy name to filter the report output by backup policy. This enables analysis of job status for a specific policy.
Policy Type	Select the policy types for which you want to generate the report. If no policy is selected, report is generated for all the policies displayed in this field. This filter is available only for NetBackup product reports. The field is visible by default for the following reports: ■ NetBackup AIR Replication Import Jobs ■ Running and Queued Job Summary ■ Consecutive Errors ■ Error Log Summary by Policy ■ Error Log Summary ■ Error Log Summary by Server ■ Job Volume Summary ■ Job Summary ■ NetBackup Policies For other reports the field can be added to the scope selector using steps described under *Work with the dynamic template designer*. See "Dynamic Template Designer Overview" on page 220.
# of Consecutive Errors	Use this filter to exclude consecutive errors. This option comes in handy when troubleshooting by limiting the report to hosts with excessive activity. Select a number from the drop-down list. Only sets of consecutive errors >= the specified number will be included in the report.

Table 5-2 Scope selector field descriptions *(continued)*

Field name	Description
Dynamic Report Start & End Times	Provides the user with the ability to specify dynamic report start and end day and time. Simply specify the number of days before or after the current day as well as the time for both the start and end day and times to be used to generate the report.
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.

Backup Manager advanced scope selector settings

The availability of selections, options, and default values varies from report to report because not all selections and options apply to all reports.

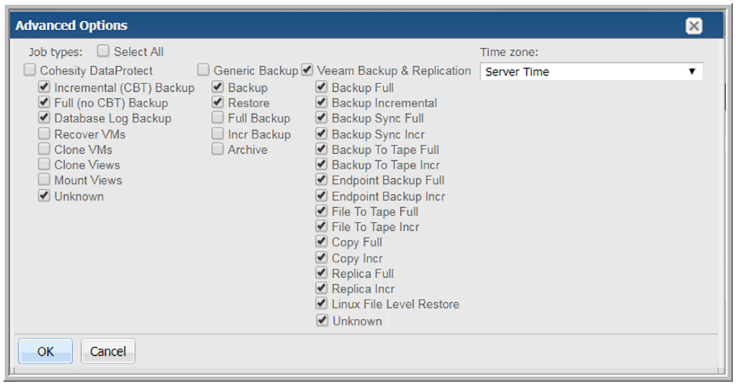
Table 5-3 Advanced scope selector settings

Field name	Description
Start Time / Finish Time	Specify the backup/restore event start or finish time to be used when culling the events within the time frame for the report. For example, if you only want to display events that occurred between 12AM and 6:30AM, then specify those times for the start and finish time parameters so that the application filters the events by those time constraints. Available in various backup/restore events reports.
Time Zone	Enables you to select a specific time zone to normalize the report by any time zone. The default setting for this parameter is the time zone setting of the Management Server.
Error Code Excludes	Check the **Set Error Excludes** box and the application will dynamically generate an Error Code ID/Descriptions Dialogue where you can check all the error codes that need to be excluded from the report. Then, click **Accept** and all the codes will auto-populate the Exclude Id text box. Any event which existed with an error code specified in the exclude list will not be included in the resulting report. This functionality currently is available only for the Consecutive Errors By Client Report.
Ignore Retries	Check this box to indicate that retries for backup jobs should not be incorporated in the statistics shown.
Job Types	Refers to the type of backup job. Each backup product has its own set of job types. The list represents the products for which data collection has been licensed for your Portal. Check the job types to be included in the report output. See “Examples of Advanced Options Job Types” on page 127.

Examples of Advanced Options Job Types

Job types enable more granular filtering within backup vendor subsystems.

Note: For Cohesity NetBackup, a previous job type, Incr Backup, has been replaced with more specific types: Cumulative Incr Backup and Differential Incr Backup. These job types more accurately reflect the backup types.



Capacity Manager scope selector settings

The availability of the selections, options within the selections, and default values for the options vary from report to report because not all selections and options apply to all reports.

Storage Arrays	Lists all storage arrays that the Data Collectors identified. Expanding the top-level folder enables access to other folders for: ■ All Storage Arrays ■ Storage Array Vendors ■ Attributes
Hosts	Use this filter to specify what data will be included in a report. ■ Host Group. Includes data from default and user-defined host groups. ■ Attribute. Includes data from system and user-defined attributes.
Host Type	Indicates how a host has been commissioned in an enterprise, such as VM Guest, VM Server, VIO Guest, VIO Server, or Other. This selection enables reporting of various types of virtual hosts.
Start Date / End Date	If you want a static report, specify a start and end date in the date format: **MM/DD/YYYY**. If you want a dynamic report--that is, a report that changes as the underlying data changes--do not specify a start and end date, as the default always represents the current day.

Threshold	For Capacity At Risk reports only, limit the report output to a threshold category: ■ Low ■ Warning ■ Critical
Group By	Use the drop-down list to specify how you want the data in the bar charts or tables to be grouped. The available options are dependent on the type of report. Some reports group by time, while others group data by either host, client, or policy.
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.

Data Collection Status Details scope selector settings

Data Collection Status Details report has the following scope options available:

Time Period	Select the data collection time period to display in the report results.
Status	Select the data collection status to display. Choose from: All Status, Successful, Warning, Successful or Warning, Warning or Failed or No Status, Failed or No Status.
Run Type	Select the type of collection run: All, Scheduled, or On-Demand.

Schedule	Select from All, Enabled, or Disabled. This indicates if the policy schedule is enabled or manually disabled.
Group By	Select from None, Collector or Policy.

Host scope selector settings

Certain reports that have been developed using the Dynamic Template Designer, will have the following scope options available.

Hosts	Specify hosts to be included in the report scope.
Backup Server Type	The type of host in a backup environment, such as Client or Media Server.
Domain	The domain to which the enterprise object belongs.
External Name	Name of the host that will appear in reports.
Group By	Provides a drop-down list to specify how the data in the charts or tables will be grouped. The available options are dependent on the type of report. Some reports group by time, while others group by either host, server, client, or policy.
Host	Name of the host, as defined in the collected product.
Host Type	Indicates how a host has been commissioned in an enterprise, such as VM Guest, VM Server, VIO Guest, VIO Server, Oracle Container, Oracle Zone, Hyper-V Server, Hyper-Guest, or Other. This selection enables reporting of various types of virtual hosts.
Make	Make of the host, such as Dell.
Maximum legends for Pie Chart	Specify the maximum number of slices for pie chart rendering.
Model	Model of the host, such as Windows-x86.
OS Platform	Host's OS type, such as Linux or Windows.
OS Version	Version of the host's operating system.
Product Collected	Find hosts that have been collected from particular subsystem, such as Veritas NetBackup.
Product Group Collected	Find hosts that have been collected by this product group, such as capacity collection.

Product Group Not Collected	Find hosts that have not been collected by this product group, such as capacity collection.
Product Not Collected	Find hosts that have not been from particular subsystem, such as Veritas NetBackup.
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 IT Analytics products, such as Backup Manager or Virtualization Manager.
Time Period	Provides a drop-down list to specify a time span, such as last 90 days or previous month.
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.

Ransomware Scorecard scope selector

Ransomware Scorecard scope selector has the following options

Table 5-4 Ransomware Scorecard scope selector field description

Field	Description
Domain	Domain for which you want to create the Ransomware Scorecard
Item Status	Filters the scorecard view based on item status as follows: ■ **All**: Displays both answered and unanswered questions on the scorecard. ■ **Completed**: Displays only the answered questions on the scorecard. ■ **Uncompleted**: Displays only the unanswered questions on the scorecard.

Table 5-4 Ransomware Scorecard scope selector field description *(continued)*

Field	Description
Item Visibility	Filters the scorecard view based on its visibility status as follows: ■ **All**: Displays both enabled and disabled questions on the scorecard. ■ **Enabled**: Displays only the enabled questions on the scorecard. ■ **Disabled**: Displays only the disabled questions on the scorecard. The authority to enable or disable a question lies with the super
Item Type	Filters the scorecard view based on the question item type as follows: ■ **All**: Displays all question types on the scorecard. ■ **Question**: Displays only questions on the scorecard. ■ **Data**: Displays only report-based questions on the scorecard.

Solution reports scope selector settings

The Solutions reports have scope selections relevant for the solution:

- See “Storage Optimization Scope Selector” on page 133.
- See “Risk Mitigation Scope Selector” on page 135.

These selections enable “what if” reporting, where you can modify parameters to see results before saving an instance of the report. When selecting the report scope, a list of the relevant solution rules can be modified by customizing and enabling only the rules for which you want to include data in the report.

It’s important to note that Rules enabled and configured on the **Admin** tab are the defaults for Adhoc report scope selection. You can use the adhoc reporting to report on one rule or many rules. Changing the Rule configuration while running Solution reports does not change anything on the **Admin** tab.

See “Add/Edit a threshold policy” on page 610.

See “Configure Risk Mitigation rules” on page 647.

Cascade to Sub-Domains

The scope selector includes a checkbox to Cascade to Sub-Domains. Depending on the context and the type of data associated with the report, this selection can have different meanings, as listed below.

- For host data, the sub-grouping refers to Host Groups.

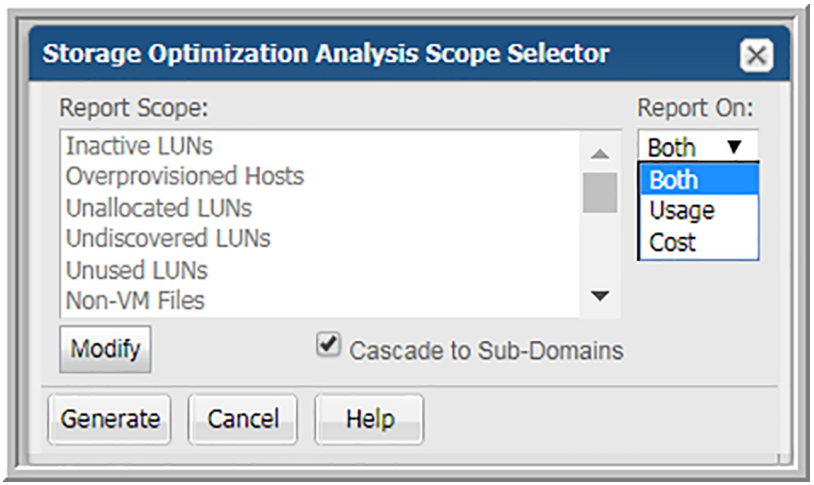
- For capacity data, the sub-grouping refers to IT Analytics Domains.
- The scope selector defaults to cascade to all child sub-groups when generating the report.

Storage Optimization Scope Selector

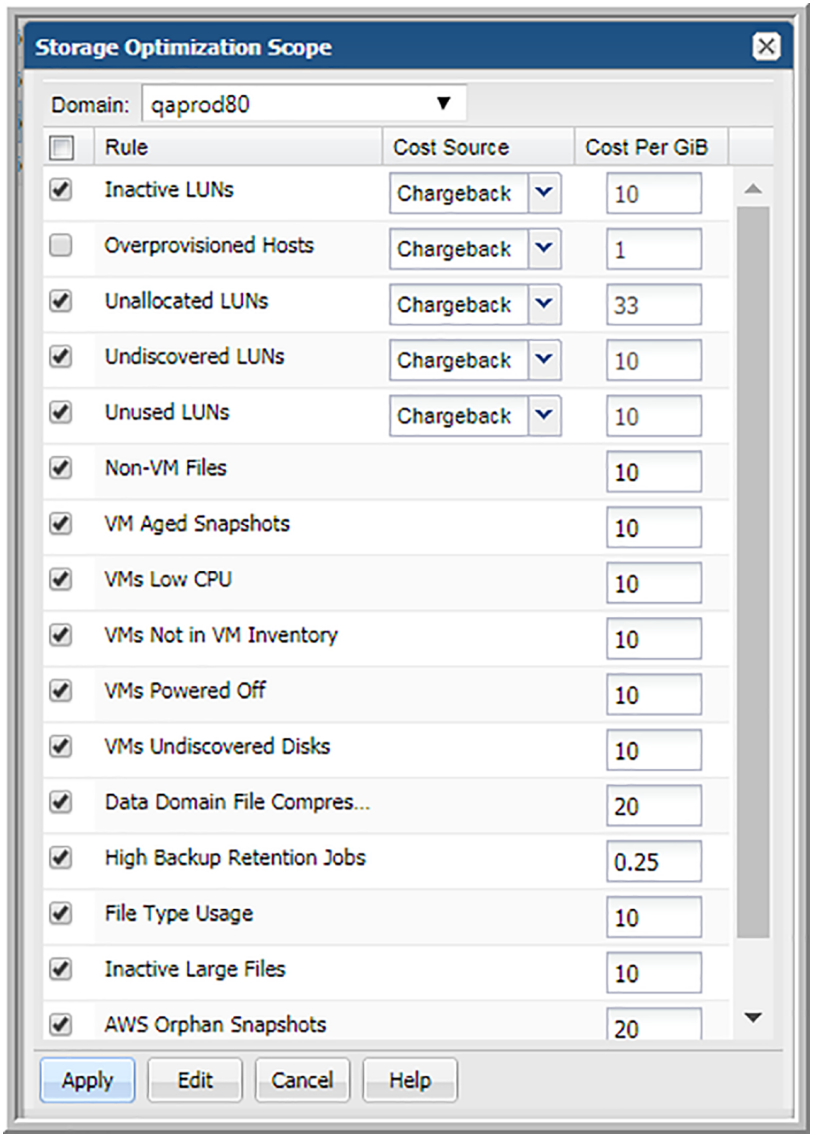
Storage Optimization reports allow you to modify the relevant solution rules by customizing and enabling only the rules for which you want to include data in the report. You are also able to select if the generated report focuses on Cost, Usage or Both.

See "Cascade to Sub-Domains" on page 132.

1. Double-click the report to launch the scope selector.



2. Choose the **Report On** focus.
3. Click **Modify** to view the relevant solution rules for customization for this report instance.



4. Click the top checkbox to enable all **Rules** or choose them individually.
5. Revise the **Cost Source** if required. Choose from **Chargeback** or **Rule**. If you select **Rule**, enter the value to use for **Cost per GiB**.
6. Double-click a **Rule** to modify the parameters; or select a rule and click **Edit**.

Changing the **Rule** configuration while running Solution reports does not change any setting defined under Storage Optimization administration. These selections are for this instance of the report allowing for what-if scenarios.

See "Add/Edit a threshold policy" on page 610.

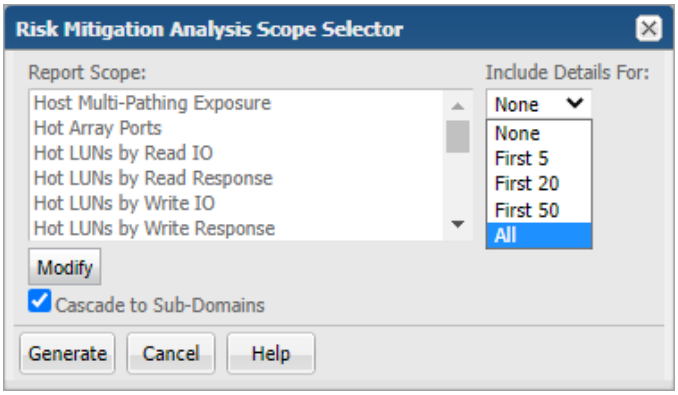
7. Click **Apply**.

Risk Mitigation Scope Selector

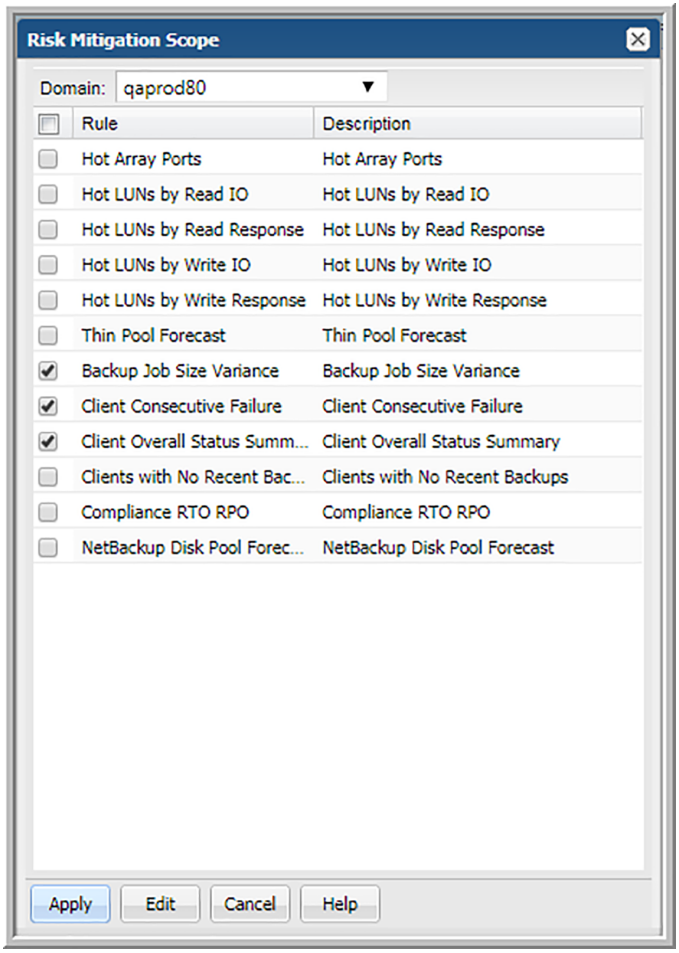
Risk Mitigation reports allow you to modify the relevant solution rules by customizing and enabling only the rules for which you want to include data in the report.

See "Cascade to Sub-Domains" on page 132.

1. Double-click the report to launch the scope selector.



2. Click **Modify** to view the relevant solution rules for customization for this report instance.



3. Click the top checkbox to enable all **Rules** or choose them individually.
4. Double-click a **Rule** to modify the parameters; or select a rule and click **Edit**.

   Changing the **Rule** configuration while running Solution reports does not change any setting defined under Risk Mitigation administration. These selections are for this instance of the report allowing for what-if scenarios.

   See “Configure Risk Mitigation rules” on page 647.
5. Click **Apply**.
6. Specify the extent of details you want to view for each category of the chart under **Node Details** based on the description below:

   - **None**: No details are displayed.

- **First 5**: The first 5 rows in each category are displayed.
- **First 10**: The first 10 rows in each category are displayed.
- **First 20**: The first 20 rows in each category are displayed.
- **First 50**: The first 50 rows in each category are displayed.
- **All**: All rows in each category are displayed.

7. To cascade the settings to all child domains when generating the report, select **Cascade to Sub-Domains**.
8. Click **Generate**

System health check scope selector

The following scope options available:

Time Period	The reporting time periods for this report. Your choices range from the last 5 minutes to the last 10 years. The options for this parameter depend on the report.
Start Date / End Date	If you want a static report, specify a start and end date in the date format: **MM/DD/YYYY**. If you want a dynamic report--that is, a report that changes as the underlying data changes--do not specify a start and end date, as the default always represents the current day.
Group By	Use the drop-down list to specify how you want the data in the bar charts or tables to be grouped. The available options are dependent on the type of report. Some reports group by time, while others group data by either host, client, or policy.
Mask User Details	Choose Yes or No to mask user details. Default is Yes, so the report is in Summary mode. Detailed mode displays all user data. Summary mode masks user names and email addresses.
Show Top	Select a value to indicate the number of entries to display under any of the System Health Check reports that define a Top value. For example, if a report shows Top 50 Largest Tables, use this selection to change that value to 100. Select from default, 50, 100 or 200.

Generate reports

When you run a report, IT Analytics takes the scope of the report, and checks if the cache contains the same report, for the same scope. If it does, the results are displayed from the cache. If the combination does not exist, the report is run from the database, saved in the cache, and then sent to the user interface. Cached reports are shared across users who belong to the same home host group.

When a report is served from the cache, an indicator icon is displayed on reports and dashboards. You can roll over the indicator to show the age of the report from the cache.

See “Understand report data caching” on page 839.

To find and generate a report

1 Use **Search** to find and run a report or a dashboard by name. Search is case insensitive, supports partial entries, and will display a list of potential matches using a string match. If the Cloud is enabled, search results will also include templates available from the cloud. Because you can directly run a report from your search results, this enables you to preview any cloud template before you save it locally.

As you enter the report or dashboard name in the **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.

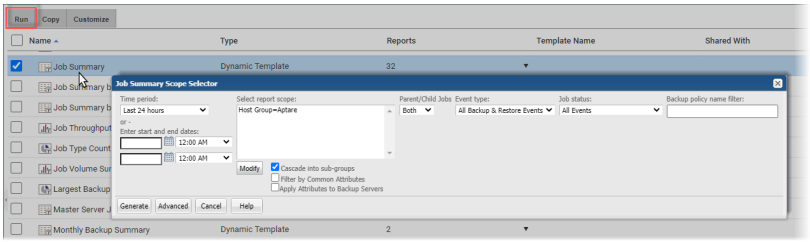
IT Analytics provides different navigation options to slice and examine your collected data. You can explore the data by using the customizable report templates or by using parts of your IT infrastructure as entry points. The Inventory Navigator serves as a browser for your infrastructure by object type.

See “Exploring your inventory” on page 60.

Use the **Reports** tab to examine the IT Analytics catalog of templates, dashboards and reports - organized by products along with user-created, and system folders.

See “Templates, reports, and dashboards overview” on page 54.

2 Click the report title and the **Scope Selector** is displayed. You can also select the report checkbox and click **Run** to launch the **Scope Selector**.



The scope selector is unique to the report you select and displays only the parameters that are relevant to the specific report. The image above is for Job Summary report.

See “Report Scope ” on page 110.

3 Specify the report parameters, and click **Generate**.

Note: When the Portal determines that a large amount of data has been returned for display, the resulting report window provides paging links: **Previous** and **Next**.

Generate reports from the Cloud

In general, cloud user report templates and their reports function the same as system templates and reports. There are a few cases where they behave differently. All templates, system, cloud and custom, are linked to their reports. This means that changes made to the template are inherited by the linked report. This rule is true for cloud linked reports, with one addition, when a cloud template is updated in cloud, you are notified by a **NEW** badge on your user folders.

For example, if you run a cloud template to create a report and save it, Cloud One, the cloud template and your report, Cloud One, are saved locally in your portal. Now let's say the cloud template has been modified by customer support and a new version has been promoted to the cloud. The next time you save Cloud One, you are prompted to accept the template modifications or refuse them.

Units of Measure in Reports

For certain reports, an Advanced option in the scope selector enables the selection of the unit of measure--MBytes, GBytes, or TBytes.

If you notice a discrepancy in the values reported by the Portal reports versus the values you see in the backup product itself, this is likely due to the conversion of the units of measure.

When 1024 is used for calculations, it refers to binary multiples: kibibyte (KiB), mebibyte (MiB), gibibyte (GiB), tebibyte (TiB), and pebibyte (PiB). When 1000 is used for calculations, it refers to decimal multiples: kilobyte (KB), megabyte (MB), gigabyte (GB), terabyte (TB), and petabyte (PB). By default, the units for values in reports are displayed as binary multiples. This default setting can be changed by modifying a user's profile.

Detail Reports

For reports that are the results of drilling into details from a parent report, those detail reports can have their own unit of measure - separate from the parent report settings. Once you've drilled into a child report, you can change the measurement

unit to see more accurate or easily understandable information. For example an Array and Capacity Utilization Summary report shows aggregated data in TiB.When you drill down, it may be more useful to switch to MiB.

Advanced Filtering for Tabular Reports

In addition to the filtering in the scope selector, table-formatted reports can be further filtered on Rows and/or Columns, using the following procedures.

Note: To optimize performance, be sure to use the filtering that's available in the scope selector before using any advanced filtering.

To filter reports by rows

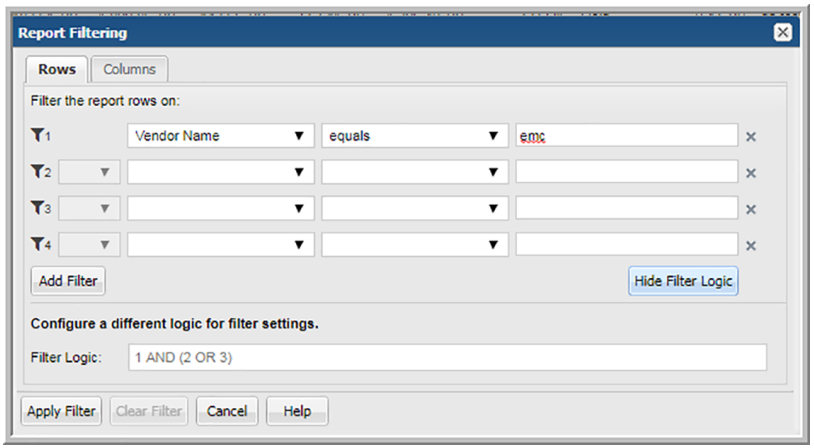
You can define the criteria for the data rows displayed in a report. Drop-down lists enable selections from the available columns. Next, you supply the operator--such as equals or does not contain--and a value for that column. Up to 16 selections can be joined to form the filter.

Filtering Limitations: Filtering on exact decimal values may produce inconsistent results, as reports display rounded values of higher-precision data. For example, a filter of MBytes equals 150.95 will return rows with rounded values. To configure a filter that more precisely narrows the content, include additional row-filtering criteria with > and < logic to limit the data that is displayed.

Note: Reports can have both row and column filters applied.

---

1. Generate a tabular report, such as the Job Summary or the VM Summary report.
2. Click Actions and select **Filter**.



3. Select a field name in the first drop down list.
4. Choose an operator.

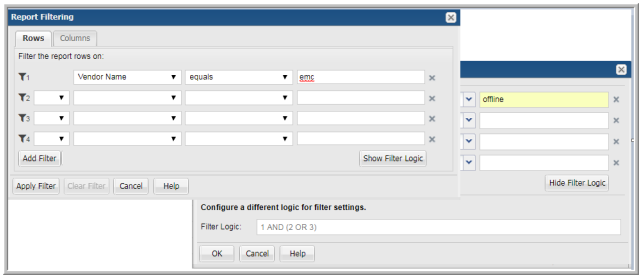
   See “Report Filter Operators” on page 143.
5. Enter a value in the third drop down list.
6. Select **And/Or** if you are adding another set of filters to your query.
7. Click **Add Filter** to add up to 12 additional filters.
8. Click **Show Filter Logic** to customize the filter logic if required. Logic defined in this field will override any setting established on the top of the dialog.

   See “Customize report filter logic” on page 142.

Customize report filter logic

Customize filter expression order and the operators using the **Filter Logic** field. Logic defined in this field will override any setting established on the top of the dialog. Use the numbers on the left of the filter expressions to construct your **Filter Logic**.

1. Click **Show Filter Logic** to expand the window. This action disables the operators you set, and the Filter Logic field becomes mandatory.



2. Edit the logic using the filter numbers and by adding parentheses or changing the operators. For example, you can change “1 AND 2 OR 3” to “1 AND (2 OR 3)”.

Report Filter Operators

Table 5-5 Report filter operators

Operator	Description
equals[1]	Filters rows where the value of the associated column is equal to the value entered. The column can be of the type Number, Date or String.
not equal[1]	Filters rows where the value of the associated column is not equal to the value entered. The column can be of the type Number, Date or String.
greater than	Filters rows where the value of the associated column is greater to the value entered. The column can be of the type Number, String, Date, Duration, and File Size. The match is case insensitive.
less than	Filters rows where the value of the associated column is smaller than the value entered. The column can be of the type Number, String, Date, Duration, and File Size. The match is case insensitive.
greater than or equal to	Filters rows where the value of the associated column is greater than or equal to the value entered. The column can be of the type Number, String, Date, Duration, and File Size. The match is case insensitive.
less than or equal to	Filters rows where the value of the associated column is smaller than or equal to the value entered. The column can be of the type Number, String, Date, Duration, and File Size. The match is case insensitive.
contains[1]	Filters rows where the value entered is present anywhere in the value of the associated column. Example: If the column value is “rattle” and the value entered is “rat” or “at” the row will be displayed. The column can only be of the type String.

Table 5-5 Report filter operators *(continued)*

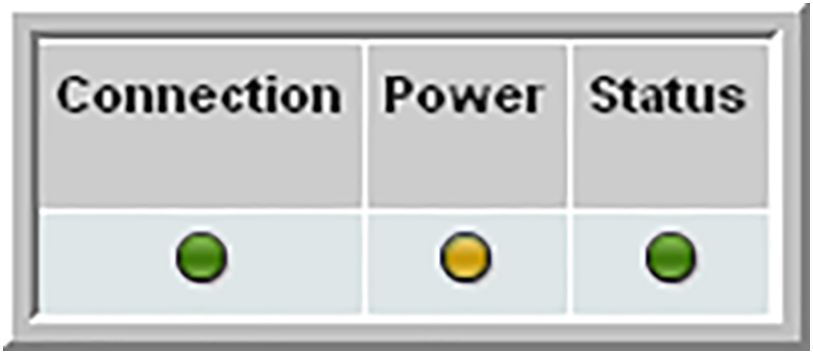
Operator	Description
does not contain[1]	Filters rows where the value entered is not present anywhere in the value of the associated column. The column only can be of the type String.
is a member of	Filters rows where the value of the associated column is equal to any of the comma-separated values entered. There is no limit to the number of characters that can be entered for this value. The column can be of the type Number, String, Duration, and File Size. Example: If you enter a value of chair, table, desk and the column value is table, the row will be displayed.
is not a member of	Filters rows where the value of the associated column is not equal to any of the comma-separated values entered. There is no limit to the number of characters that can be entered for this value. The column can be of the type Number, String, Duration, and File Size.
matches regular expression	Filters rows where the value of associated column matches the regular expression[2] entered. The column can only be of the type String. For example, to find a column starting with **S**, use **^S** or **^[Ss]**
does not match regular expression	Filters rows where the value of associated column does not match the regular expression[2] entered. The column can only be of the type String.

[1] -is case insensitive.

[2] -reference: http://docs.oracle.com/javase/7/docs/api/java/util/regex/Pattern.html

Filtering Indicators, Multiple Expressions, Dates, Percentages, Yes/No

- When **filtering on Indicator Lights**, as shown in the following example, filter on the color: Red, Yellow, Green, or Gray.



- To **combine several expressions** into one filter, select an and or or operator from the drop-down list at the end of the row.
- When **filtering on dates**, refer to the following table to determine which columns can be filtered on date. Be aware that filtering on Dates needs to be carefully designed in order to avoid degradation of performance. Date criteria can be specified in the following formats:

mm/dd/yyyy	Specified date
Now	Current date and time
Today	Current day
+nd -nd	Enables a variable time frame (plus or minus a number of days) For example: +2d This criteria can be combined with **Today**. For example: Today-5d
+nh -nh	This hour option enables a selection of a date and time, plus or minus a number of hours from the current time. For example, in the Job Summary Report, select all jobs where the Finish Time was greater than -12h. This would isolate jobs for only the last 12 hours.

- When **filtering on percentages**, such as in the NetApp Volume Summary Report, enter a decimal value equivalent to the percentage.

ex.	100% - 1.0
	75% - .75
	50% - .50
	25% -.25
	10% - 0.1

- When filtering on **Yes/No** values, enter **Y** or **N** for the values.
- Click **OK** to regenerate the report with the filter applied.

The report will now contain a **filter applied** link, enabling access to the defined filter. When you save this report, the filter is saved with it.

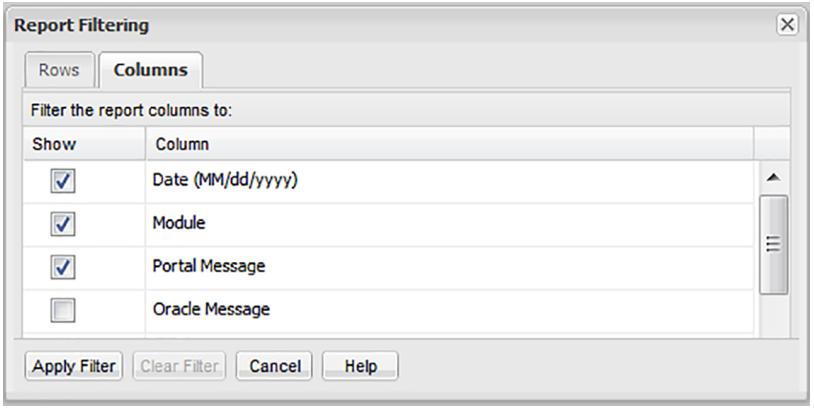
To filter reports by columns

Note: Reports can have both row and column filters applied.

The following tabular reports do not support column filtering:

- Billing and Chargeback Summary
- Chargeback by Host
- Array Performance by RAID Group
- Port Throughput by Array
- Datastore Performance Summary

1. Generate a tabular report, such as the Job Summary or the VM Summary report.
2. Click **Actions** and select **Filter**.



3. Click the **Columns** tab. By default, all columns are checked.
4. Uncheck the columns to remove from the report.
5. Click **OK** to regenerate the report with the filter applied.

   The report will now contain a **filter applied** link, enabling access to the defined filter. When you save this report, the filter is saved with it.

Sort columns in reports

For reports that display data in tables, you can easily change the order that data is displayed in columns.

To sort columns in a table report

- From the report, click any of the table headings to toggle the data's sort order: ascending or descending.
- To return to the report's original order, click **Actions** and select **Edit Scope**. Regenerate the report.

Note: Column sorting is disabled for tabular reports that have sub-total rows.

Refresh and set refresh intervals

Once you have generated and saved a report, you may want to periodically refresh the data that is displayed. You can also do this in real-time as required. The following characteristics apply to refresh intervals:

- A report or dashboard first must be saved before a refresh interval can be configured.
- Available refresh intervals include: 3, 5, 10, or 15 minutes.
- A report that is refreshing will have its last refresh time displayed in the tab of the window.
- Refresh intervals are not persistent--that is, once a report is closed, its refresh setting is cleared.
- To cancel a report that is configured to refresh, close the tab.

To set a report or dashboard to refresh

1. Generate the report and save it. A report must be saved to have the refresh capability associated with it.

   See "Generate reports" on page 137.

2. Click **Actions**.
3. Select **Set Refresh Interval**.
4. Set the refresh rate to 3, 5, 10 or 15 minutes and click **OK**.

To refresh a report or dashboard in real-time

1 Generate the report.

See “Generate reports” on page 137.

2 Click **Actions**.

3 Select **Refresh**.

Modify Reports

You can modify the parameters of a saved report. Use the information displayed in the **Reports** tab to help assess the impact of your changes. For example, you can determine if you’ve shared your report with other users or if templates have reports associated with them. Once you’ve changed the scope, regenerate the report. The report title is displayed in red italic to indicate a change has occurred. You can save your changes or just close the report to retain the original scope.

To modify a saved report

1 Use **Search** to find a report or a dashboard by name. Search is case insensitive, supports partial entries, and will display a list of potential matches.

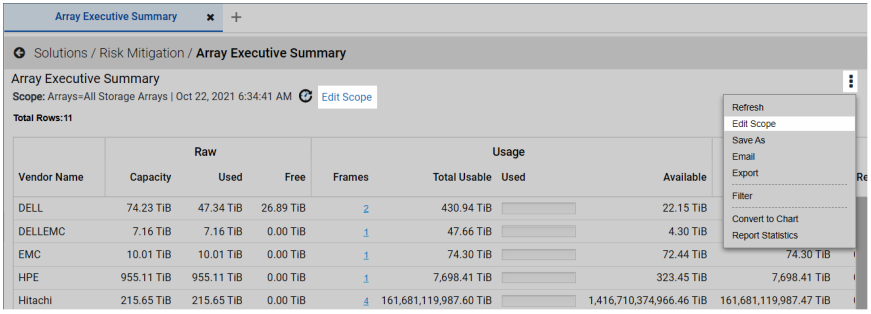
As you enter the report or dashboard name in the **Search** field, up to 10 potential matches are displayed. If the report or dashboard is shown, you can select and run it directly from the match list. You can also select **All Items** to display the full search results page.

You can use the **Reports** tab to browse through the IT Analytics reports. The reports are displayed by product.

See “Templates, reports, and dashboards overview” on page 54.

2 Select the report and click **Run**. Double-clicking also runs the report.

3 Click **Actions** menu and select **Edit Scope**.



4 Change the report parameters, then click **Regenerate**.

5 Save the changes to update the report with the new scope. You can also:

- Close the report to discard the scope changes and retain the original scope
- Select **Save As** from the **Actions** menu to save the new scope under a new name.

See "Save Reports" on page 165.

for details on saving report.

Convert tabular report to chart

The **Convert to Chart** option in the **Actions** menu enables you to view the data from a tabular report in a chart form. This feature helps you to visualize tabular data in different chart forms and view the data patterns. You can save the charts separately under **My Reports** folder and also revert back to the original tabular report through the same **Actions** menu. Once you save the chart, you can share or email it like other reports.

To convert a tabular report to a chart:

1 Login to the IT Analytics Portal and generate a tabular report. You can navigate to your report from the **Reports** tab or use the **Search** to generate the report.

2 After generating the tabular report, go to **Actions** menu > **Convert to Chart**.

3 On the **Report Converter** pop-up, choose the chart type and other data sources based on the descriptions in the table below:

Field name	Description
Display as	Select a chart to display from the list: ■ Line Chart ■ Bar Chart ■ Bar and Line Combo Chart ■ Area Chart ■ Pie Chart ■ Donut Chart ■ Solid Gauge Chart Chart description

Field name	Description
Chart type	Choose between: **1** **Simple**: Displays a simple plot of data points. **2** **Pivot**: Displays a plot of data points specific to each category. This option helps you to summarize data based on a specific category. Pivot charts can help you analyze data trends specific to the pivoted category and present meaningful insights.
Category	Appears when you choose to view a Pivot chart. Specify the column name based on which you wish to categorize the data displayed on your chart..
Horizontal axis	Specify the column name that must appear on the horizontal axis.
Vertical axis	Appears when you choose to view a Pivot chart. Specify the column name that must appear on the vertical axis.
Line	Specify the column name of the data points that must appear against the horizontal axis.
Bar	Appears when you select a Bar Chart. Specify the column name of the data points that you want to view as bars on the chart. You can use the **+** and **-** to add or remove data points on the chart.
Stack bars	Displays stacked bars on the chart.
Area	Appears when you select an Area Chart. Specify the column name of the data points that must appear as the shaded area on the chart. You can use the **+** and **-** to add or remove data points on the chart.
Stack areas	Select to stack the plotted areas on the chart.
Portion	Specify the column name that you want to view as pie or donut portions.
Value	Appears only for a Solid Gauge Chart. Specify the column name of the data points that will appear as gauge
Max value	Appears for a Solid Gauge Chart. Specify the column that will denote the maximum value of the gauge.

Field name	Description
Preview	Click to preview your chart before the report is actually converted to the chart.

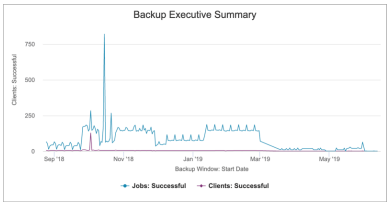
4 Click **OK** to view your chart.

The tabular report is converted to a chart. You can save this chart under the **My Reports** folder.

Chart description

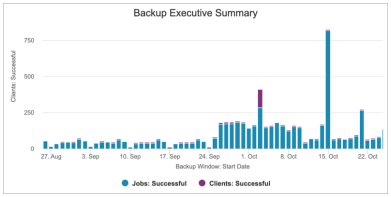
Each available chart type is described in the table below.

- Line Chart



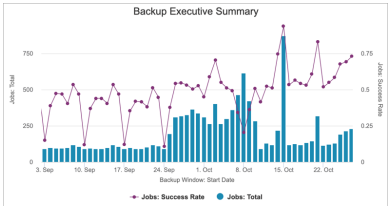
These charts are ideal to view data trends. Usually, the category data is assigned to the horizontal axis and values are assigned to the vertical axis. Data is distributed evenly on both the axes.

- Bar Chart



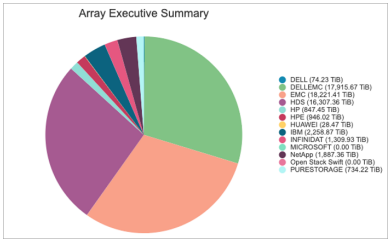
These charts are used compare individual items, either side-by-side or in a stacked manner. Typically, the category data is assigned to the vertical axis, while numeric data values are plotted evenly on the horizontal axis.

- Bar and Line Combo Chart



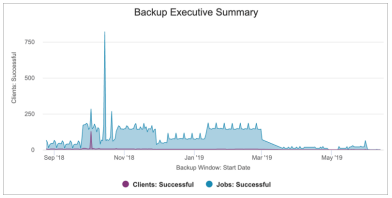
These charts are used for easier data representation by combining bar and line charts. One set of data is displayed in a series of bars while the other is represented as a line.

- Pie Chart:



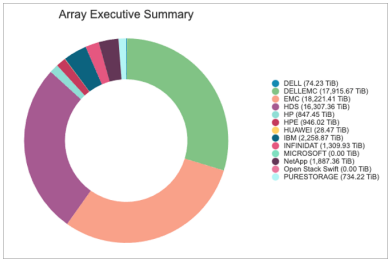
The portions displayed are in proportion to the sum of the values of the selected field and are represented as percentage. Ensure you do not have zero or negative data values while using a pie chart.

- Area Chart



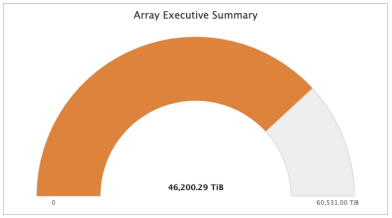
This chart is used to highlight the total value across a trend. It shows the relationship of the sum of one part of the plotted values to the whole. You can also view this chart in stacked form.

- Donut Chart



The portions displayed are in proportion to the sum of the values of the selected field and are represented as percentage. Ensure you do not have zero or negative data values while using a donut chart.

- Solid Gauge Chart



This chart is used to compare two values of the same category, such as total storage against consumed or free storage. The chart displays a comparison of the sum of values from one column against the other.

Convert a chart into a tabular report

The option to convert a chart to tabular report is available for native chart reports and for the tabular reports converted to charts.

To convert a chart to its tabular report:

1. Login to the IT Analytics Portal and open the chart report that you want to revert to its tabular form. Since such reports are saved under the **My Reports** folder, you need to access it from the **Reports** tab > **My Reports** folder to open the report.
2. Go to **Actions** menu > **Convert to Table**.

   The chart will revert to its original tabular report and display all the default columns with the updated data.

Add and remove custom header and footer to a report or dashboard

By adding a custom header or footer to a report or a dashboard, you can provide captions such as confidential, for internal use, or classified for the report viewer.

To be able to add or remove custom header and footer of a report, configure the **Classified Header** and **Classified Footer** values in the first place on the **Admin** > **Advanced** > **System configuration** > **Portal** tab > **Report** section and restart the portal Tomcat service. Restarting the portal Tomcat service is essential to be able to add or remove custom header and footer from the Scope Selector pop-up of the reports.

Add custom header and footer

To add a custom header and footer to a report or dashboard:

1. Login to the IT Analytics Portal and find the report on which you want to add the custom header and footer. You can use the **Search** or navigate to the report from the **Reports** tab to get to the report.
2. Select the report and click **Run**. The Scope Selector for the report is displayed.
3. Click **Advanced** and select **Include Custom Header and Footer**.

   The **Advanced** button is displayed, provided you have configured the **Classified Header** and **Classified Footer** values under System Configuration and restarted the portal Tomcat service as described above.
4. Click **OK** and click **Generate**. In case you are in the process of editing the scope of the report, you need to click **Regenerate**.

   The report or dashboard is displayed with the custom header and footer specified under the **System configuration** > **Portal** tab.

To view the custom header and footer every time you generate the report, you must save the report (under **My Reports**). Otherwise, the report is reset to its default appearance when you close it.



Remove custom header and footer

To remove the custom header and footer from the report or dashboard:

1. Login to the IT Analytics Portal and find the report on which you want to remove the custom header and footer. You can use the **Search** or navigate to the report from the **Reports** tab.

   If you have the report open, skip to step 3.
2. Select the report and click **Run**.
3. Go to the **Actions** menu > **Edit Scope**. The Scope Selector pop-up displays an **Advanced** button.

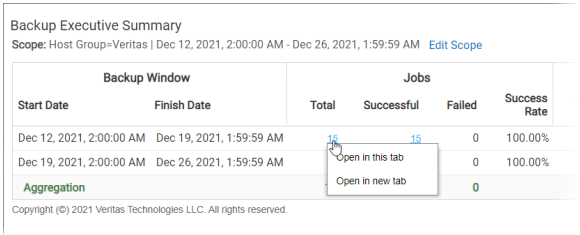
4 Click **Advanced** and clear **Include Custom Header and Footer**.

5 Click **OK** and click **Regenerate**.

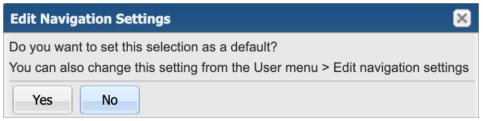
The custom header and footer are removed from the report or dashboard. You must save the report to retain the change.

Set or edit the drill down preference

Several IT Analytics reports, both tabular and chart, provide a drill down view of their respective report components. You can choose to open the drill down view either in a new tab or within the same tab through the popup menu displayed when you click either a chart component or the drill down link in a tabular report.



If you choose the same drill down option three times in succession from the popup menu, the following dialog is displayed from which you can set a default drill down preference for your account.

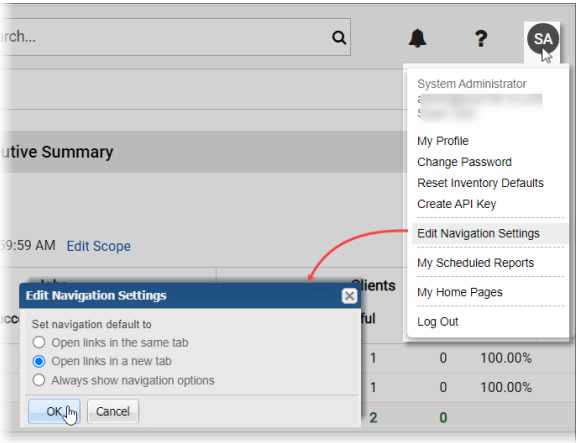


This procedure below describes how you can set or edit the report drill down preference for your account.

To set the drill down behavior:

1 Login to the IT Analytics Portal and open the user menu.

2 Select **Edit Navigation Settings**.



The **Edit Navigation Settings** popup is displayed.

3 Select the drill down behavior on the popup and click **OK**.

Your drill down navigation preference is set permanently for your account until you change it again. The drill down popup menu will cease to appear unless your preference is set to **Always show navigation options**.

Access NetBackup web interface from the IT Analytics Portal

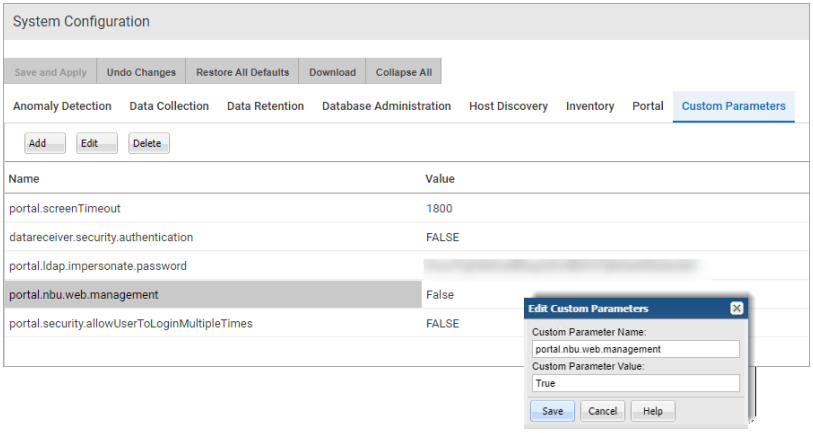
With the configuration described below, you can access the NetBackup web UI from the Inventory view and from Backup Policy Details report of the IT Analytics Portal. This provision allows you to manage the NetBackup policies directly from the portal, provided your NetBackup version is 9.0 or higher.

Before you attempt to access the NetBackup web UI, ensure the NetBackup Data Collector policy has completed at least one successful data collection and the **Clients Detail** probe on the Cohesity NetBackup Data Collector policy configuration screen is selected during the collection. This prerequisite enables the portal to verify the NetBackup version as the **Clients Detail** probe connects directly to each NetBackup client to collect and persist environmental details.

Enable NetBackup web UI access from the portal

To enable NetBackup web UI access from the IT Analytics Portal:

1. Login to the IT Analytics Portal and open the **Admin** tab.
2. In the left pane, go to **Advanced**>**System Configuration** and then click the **Custom Parameters** tab.



3. Select **portal.nbu.web.management** parameter and click **Edit**.
4. On the **Edit Custom Parameters** popup, set **Custom Parameter Value** as **True** and click **Save**.
5. Click **Save and Apply** and accept the confirmation message.
6. Restart the IT Analytics Portal service.

   On Linux, you can run `/opt/aptare/bin/tomcat-portal restart` and on Windows, you can restart from the Services list.

NetBackup web UI access is now enabled.

Access NetBackup web UI from the Inventory view

Ensure you have the NetBackup access credentials when you perform this procedure.

On the IT Analytics Portal:

1. Go to **Inventory** > **Backup Servers** or **Hosts** > **NetBackup** and select one of the NetBackup Primary Servers.
2. Click the **Actions** menu and select **Manage** from the menu.

   NetBackup web UI is launched in a separate tab and its login page is displayed. Use your access credentials to login and perform the further operations. If you are already logged in, you will be directed to the dashboard of the NetBackup UI.

Access NetBackup policy web UI from Backup Policy Details report

Ensure you have the NetBackup access credentials when you perform this procedure.

On the IT Analytics Portal:

1. Go to **Reports** > **Backup Manager** > **Backup Policies** and generate the NetBackup Policies report.

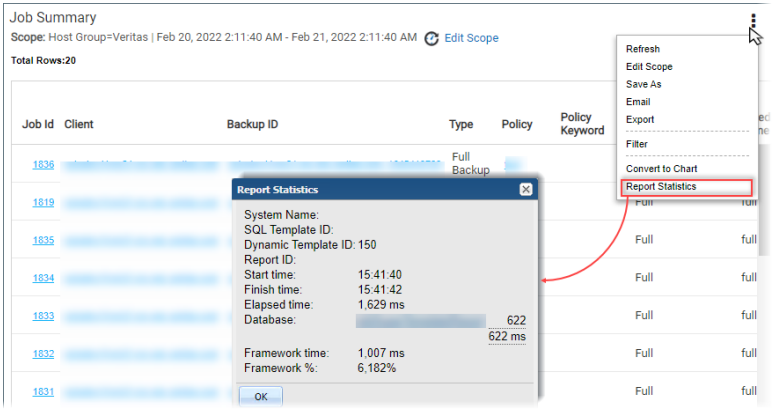
   The report displays the policy names as clickable links.
2. Click a NetBackup policy link to drill down to the Backup Policy Details report.
3. Click the **Actions** menu and select **Open in NetBackup Administration Console**.

   NetBackup web UI opens in a new tab and its login page is displayed. Use your access credentials to login and manage the NetBackup policy. In case you are already logged in, the NetBackup policy page is displayed on the tab.

View report statistics

For each report, you can view the report statistics using the Actions menu. The report statistics provide details such as System name, SQL and Dynamic Template IDs, Report ID, Database accessed, and so on. This information is essential for Cohesity Support while resolving user queries.

To view report statistics, go to Actions menu > **Report Statistics** on any report page.



Interact with Reports

Many interactive features are available depending on the report type:

- Time series - zoom and reset zoom: On a series of time-based data points, using the mouse, drag select a section of the chart to zoom. Click the **Reset zoom** button to return to your original settings.
- Time series - zoom and pan: Once you have zoomed on a series of time-based data points, you can scroll left and right through the chart. Hold down the shift key when panning.
- Expand to full screen: Charts displayed as components of a dashboard may be hard to read. Click the **Expand** icon to pop the chart out of the dashboard and have it displayed in the full browser window.
- Configure Portlets: Choose custom attributes for filters and switch positions for portlets on dashboard style displays.
- Show/hide data series: The legend displays the data series in a chart with a symbol and the name of the series. Click a value in the legend to add or remove the data point from the chart. The chart automatically updates based on the addition or removal.
- Column sorting: Click any column header on a tabular chart to sort the table by that column.
- Hyperlink drill downs: Drill down to a more granular level by clicking links in the report. This feature is available for all chart types.
  The drilldown also provides an option to open the report in the same or new window.

- Bar chart drill downs: Click the bar of the chart to the a more granular report specific to the bar data source. This feature is available for all charts.
- Time Period Quick Filters: Change the time period represented in the chart without formally editing the scope. Choose from 5 days, 2 days, 1 day, 12 hours, 4 hours or 1 hour.
- Expanding sectors: Click parts of a pie and donut charts to have sections emphasized without changing the data.

Table 5-6 Interactive features across report types

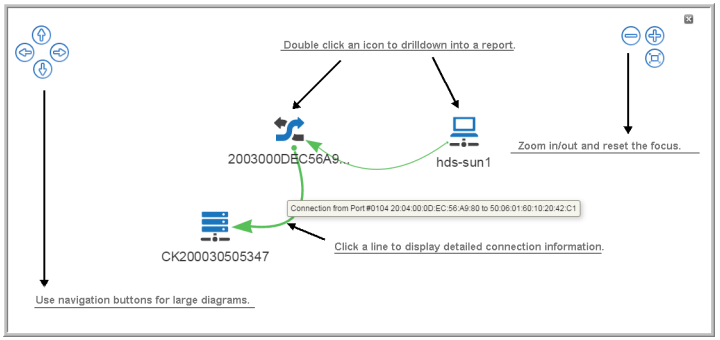
Interactive feature	Available in report type
Hyperlink drill downs	■ Tabular ■ Bar chart (all types) ■ Line chart ■ Area chart ■ Pie/Donut chart
Column sorting	Tabular
Expand to full screen	■ Tabular ■ Bar chart (all types) ■ Line chart ■ Pie/Donut chart ■ Gauge
Time series - zoom and reset zoom	■ Bar chart (all types) ■ Line chart ■ Area chart
Time series - zoom and pan	■ Bar chart (all types) ■ Line chart ■ Area chart
Show/hide data series	■ Bar chart (all types) ■ Line chart ■ Area chart ■ Pie/Donut chart
Expand sectors	Pie/Donut chart

See “Work with dashboards” on page 196.

Work with Topology

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.



Add a note to a Backup Job Details Report

Backup Manager reports provide details about the status of backup/restore jobs. Enterprise environments must monitor these reports to mitigate risk, ensure service-level agreement (SLA) compliance, and meet auditing goals. Toward that end, the Add a Note feature in the Job Details report can be used to annotate specific jobs so that details and problem resolutions can be documented.

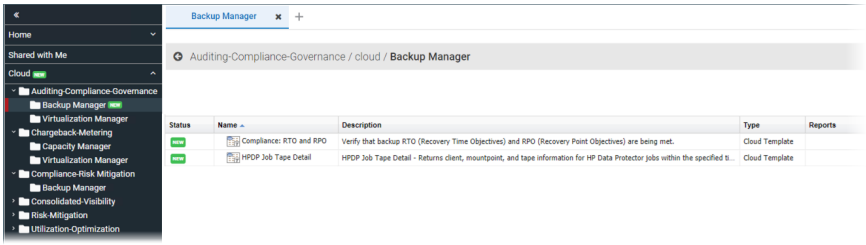
Use new and updated User Report Templates

New or updated user report templates are periodically published and automatically made available. These templates are displayed in the Cloud section for all users with Cloud privileges enabled. User and Cloud folders are badged with a **NEW** flag to indicate that new or updated report templates are available in your system.

When new or updated user templates are released the following actions are required:

- View the relevant user and **Cloud** folders badged with a **NEW** flag. These badges serve as an alert to new or updated templates.

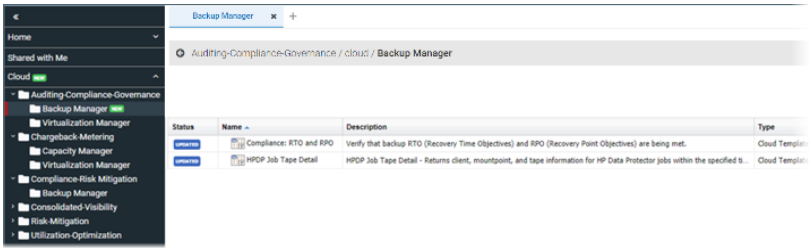
- Review the changes and then choose to update the template based on your requirements.



To use a new user report template from the Cloud

In addition to the **NEW** badge on a folder, **NEW** is displayed beside the individual report template.

1. Click the **Cloud** folder with the **NEW** badge. The reports are displayed on the view panel.

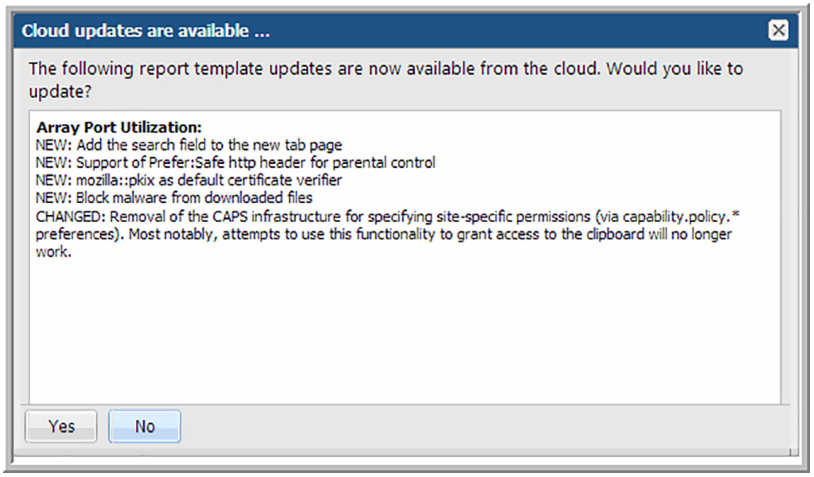


2. Select the template and run it to preview the data.
3. Click Copy (Ctrl+C) to move a copy of the template into your user folders. Choose a folder or create a new folder within your **My Reports** folder. You can rename the **My Reports** folder. If a rename has occurred, the path in this dialog box will reflect that change.

To apply updates to user report templates

1. Click the user folder with the **NEW** badge. The reports are displayed on the view panel.
2. Select the updated report template. The **UPDATED** badge is displayed in the **Status** column.

3 Click **Update** on the Action bar. A list of updates is displayed.



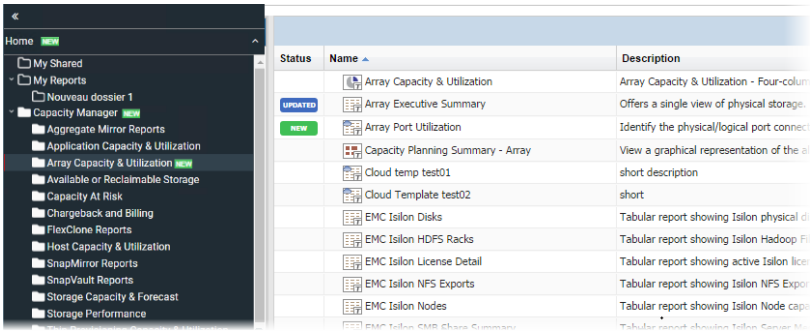
4 Review the updates and click **Yes** to accept the template changes.

Use new and updated System Report Templates (Super User only)

New or updated system report templates are periodically published and automatically made available to Super Users. These templates are displayed for the licensed products installed on your systems. System folders are badged with a **NEW** flag to indicate that new or updated report templates are available in your system. This badging is only displayed to Super Users.

When new or updated user templates are released the following actions are required:

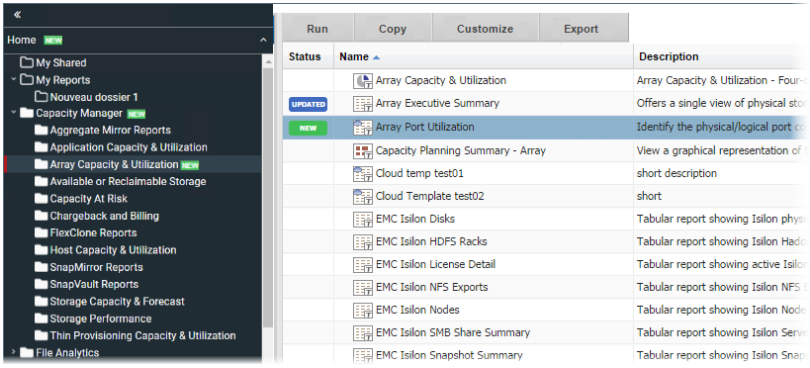
- View relevant folders badged with a **NEW** flag. These badges serve as an alert to new or updated templates.
- Review the changes and then choose to update the template based on your requirements.
- Enable templates for individual users or user groups: **Admin > Users and Privileges**.



To use new system reports

In addition to the **NEW** badge on a folder, **NEW** is displayed beside the individual report template. New report templates automatically are displayed in your system folders. Updated templates require acceptance before applying changes to your existing templates.

1. Click the system folder with the **NEW** badge. The reports are displayed on the view panel.
2. Select the new report.



3. Run the new template to preview the report.
4. Enable the template for individual users or user groups: **Admin>Users and Privileges**.

   See “Enabling new product report templates” on page 586.

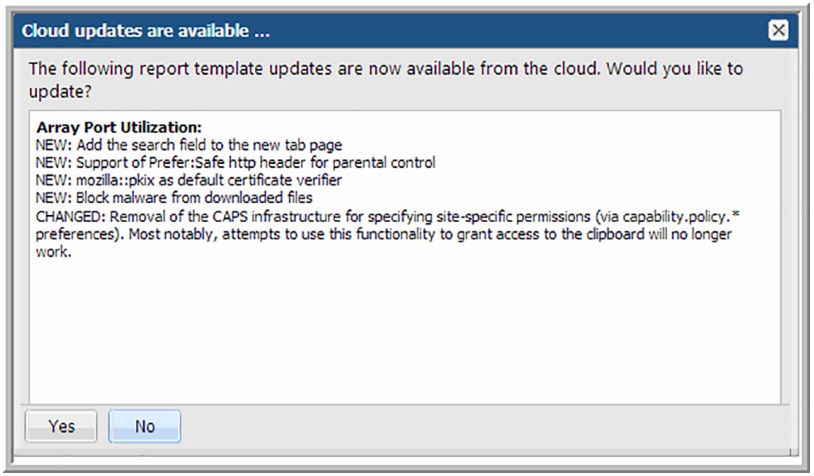
To apply updates to report templates

New report templates automatically are displayed in your system folders. Updated report templates require acceptance before applying changes to your existing report templates.

1 Click the system folder with the **NEW** badge. The reports are displayed on the view panel.

2 Select the updated report template. The **UPDATED** badge is displayed in the Status column.

3 Click **Update** on the Action bar. A list of updates is displayed.



4 Review the updates and click **Yes** to accept the template changes.

Save Reports

When you save a report, you are saving the definition and parameters, not the report output. Each time you launch a saved report, the report's data refreshes. If you intend to access this report often, consider creating a dashboard and setting it as your home page.

See "Work with dashboards" on page 196.

See "Manage My Home Pages" on page 211.

Save Cloud Reports

In general, cloud user report templates and their reports function the same as system templates and reports. There are a few cases where they behave differently. All templates, system, cloud and custom, are linked to their reports. This means that changes made to the template are inherited by the linked report. This rule is true for cloud linked reports, with one addition, when a cloud template is updated in cloud, you are prompted to accept or refuse the update when you save the report.

For example, if you run a cloud template to create a report and save it, Cloud One, the cloud template and your report, Cloud One, are saved locally in your portal. Now let's say the cloud template has been modified by customer support and a new version has been promoted to the cloud. The next time you save Cloud One, you are prompted to accept the template modifications or refuse them.

When improvements are introduced, the **Cloud** section is automatically refreshed and your local Cloud-linked templates can also be updated if you choose to do so. However, if you saved a Cloud template with a different name or made a copy, that saved template is detached from the Cloud template and will not automatically inherit the updates.

See "Use new and updated User Report Templates" on page 161.

Note: To capture report output into a saved file, use the export feature.

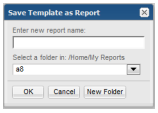
See "Export reports and dashboards" on page 169.

To save a report

1. Generate the report.

   See "Generate reports" on page 137.

2. Click **Actions** and select **Save**.



3. Enter a descriptive report name.
4. Choose a folder or create a new folder within your **My Reports** folder. You can rename the **My Reports** folder. If a rename has occurred, the path in this dialog box will reflect that change.

   Note: Reports must be saved into folders.

To save a report with a different name

1. Generate a saved report.
2. Click **Actions** and select **Save As**.
3. Enter a descriptive report name.
4. Choose a folder or create a new folder within your **My Reports** folder.

Note: You can change the name at any time by selecting the report while viewing the contents of the **My Reports** folder and clicking **Rename**.

Delete Reports

You cannot delete system report templates, but you can delete reports that you created either by saving from a system report or by building it with a template designer. When you delete a report, you permanently delete it from the reporting database and remove it from all areas in the portal.

When you delete a report you must consider where the report is being used. A single report may be:

- Shared with other users or user groups
- Used as a homepage for users or user groups
- Used in the **Inventory** as a custom report associated with an object or host group

To delete a saved report

1. Locate and select the report on the **Reports** tab.
2. Click **Delete (Del)**.

Chapter 6

Distribute, share, schedule, and alert

This chapter includes the following topics:

- Export reports and dashboards
- PDF report format
- Known Limitations for Donut Charts
- Schedule export of reports and dashboards
- Email reports and dashboards
- Schedule email of reports and dashboards
- Viewing and Managing Your Scheduled Reports
- Schedule Types
- Deleting Scheduled Reports and Dashboards
- Setting Up Alerts for Tabular Reports
- Configure an alert
- Deleting Report-Based Alerts
- Share Reports, Dashboards, and Folders
- Working with Shared Reports

Export reports and dashboards

You can export reports to external applications, such as Microsoft Excel or if you'd like to generate a hardcopy, you can export them to a file, such as a PDF. The date on the report will be the date and time that the report was generated. Date and time are derived from the Portal Server's time zone, which was determined when your administrator installed the Portal.

Exported HTML and PDF reports and dashboards match the Portal display. Report and table text uses the same font as the live view. Panel and chart sizes reflect the onscreen layout, including panel width and the proportions of pie, bar, and line charts. Reports displayed in Chinese, Japanese, or Korean use the appropriate language font when exported. Available from 11.8.2610.

Exporting data is limited to the first 20,000 rows. Reports longer than 20,000 rows are truncated when exported. The report output (PDF, HTML, CSV, etc.) displays the message "Data in this report has been truncated." You can revise default number of rows by modifying the `portal.properties` file.

On a Linux Portal server, to ensure proper rendering of reports that are emailed or exported as HTML images or PDF files, a graphics manager such as X Virtual Frame Buffer (XVFB) is required. Contact your IT organization to configure this capability, if you plan to export/email reports as HTML images or as PDF files.

Do not confuse the exporting procedures described in this section with exporting a Custom Report definition (created in the SQL Template Designer or dynamic Template Designer).

See "Export/Import SQL Templates" on page 418.

Caution: The export report feature of IT Analytics uses PhantomJS to export charts. Cohesity assures that the known vulnerability of PhantomJS (https://cve.report/vendor/phantomjs) fails to apply to the exported reports. However, if you choose to remove PhantomJS from IT Analytics Portal server, charts will not be available in your exported reports. For more information, see the Remove "PhantomJS" from IT Analytics Portal tech note.

Note: The download status of the report will be displayed in the **Downloads** popup. To see all the download performed during the session, go to the user profile menu > **Downloads**.

PDF report format

When you export a report or dashboard to PDF, IT Analytics formats the file so that it is easy to read, print, and distribute, for example when a printed or archived copy is required for compliance and audit review.

A PDF report includes the following elements:

- Report title and generation timestamp. Displayed at the top of the first page. The date and time shown reflect when the report was generated, based on the Portal Server's time zone.
- Column headers. Displayed at the top of the data table. For reports that span multiple pages, the column headers automatically repeat at the top of every page, so column context is retained no matter which page you are viewing.
- Report data. The tabular data for the report, presented in the same column order as it appears on screen.
- Page footer. Displayed at the bottom of every page and shows the current page number together with the total page count, in the format Page X of Y. For example, a ten-page report displays Page 1 of 10 on the first page and Page 10 of 10 on the last page. The total page count is calculated automatically based on the number of pages in the report.

Note: Repeating column headers and page numbering are added to the page layout only. They do not alter, obscure, or remove any report data.

This layout is designed to support common review scenarios, such as printing multi-page reports, archiving reports to meet record-retention requirements, and distributing reports to auditors, compliance teams, or other stakeholders. Because column headers and page numbers are retained on every page, reviewers can accurately reference and validate data even after a report has been printed or removed from its original digital format.

See “Export reports and dashboards” on page 169.

Known Limitations for Donut Charts

When exporting and emailing Donut Charts there are known limitations including:

- Chart size may not in be in the same proportion as in the web browser.
- Mouse overs are missing.
- Totals are missing from the center of the chart.

To export report or dashboards

1. Generate the report or dashboard.

   See "Generate reports" on page 137.

2. Click **Actions** and select **Export** with the report/dashboard displayed.

3. In the **Export Report/Dashboard** dialog box, choose the format. The report type dictates the available export options. Not all export types are available for all report types.

   The following lists all potential options:

   - CSV(**.csv**). Exports data to a spreadsheet or a database application such as Microsoft Access.
   - Excel (.xlsx). Similar to the CSV export, this file supports MS Excel-specific features.
   - XML (.xml). Accesses the web page components.

   Note: XML is not available as an export type for dashboards.

   - PDF (.pdf). This read-only file lends itself to easy distribution and printing.
   - Text. A text file of the report. This selection provides the option to set a delimiter and use quotes in the exported file.
   - HTML (.html). Exports this data as a web page, along with supporting files, in a zipped file.

   The Export-to-HTML process produces a **.zip** file with all the necessary files. Extract the files and then click **report.html** to display the report. Depending on your browser's settings, the file will be saved to your default location (typically, Desktop), or a dialog box launches and requires that you specify the location for the file.

4. Click **Export Now**.

5. **Open** or **Save** the file.

Schedule export of reports and dashboards

By scheduling the export of a report or dashboard, you configure it to run and then export it at the specified schedule. You can schedule export for saved reports and dashboards. To schedule the export of a report or a dashboard template you must generate and save its instance that can run on your prescribed schedule.

Note: Exporting data is limited to the first 20,000 rows. Reports longer than 20,000 rows are truncated when exported. The report output (PDF, HTML, CSV, etc.) displays the message “Data in this report has been truncated.”

See “Known Limitations for Donut Charts” on page 170.

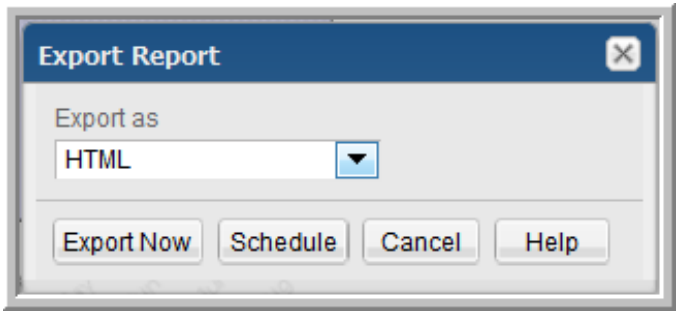
To schedule an export of report or dashboard

1. If you are scheduling an export of a report or dashboard template, generate the report or dashboard. To schedule export of a saved report or a dashboard, skip to the next step.

   See “Generate reports” on page 137.

   See “Save Cloud Reports” on page 166.

2. Click **Actions** and select **Export** with the report/dashboard displayed.
3. On the **Export Report** dialog box, choose the format. The report type dictates the available export options. Not all export types are available for all report types.



The following lists all potential options:

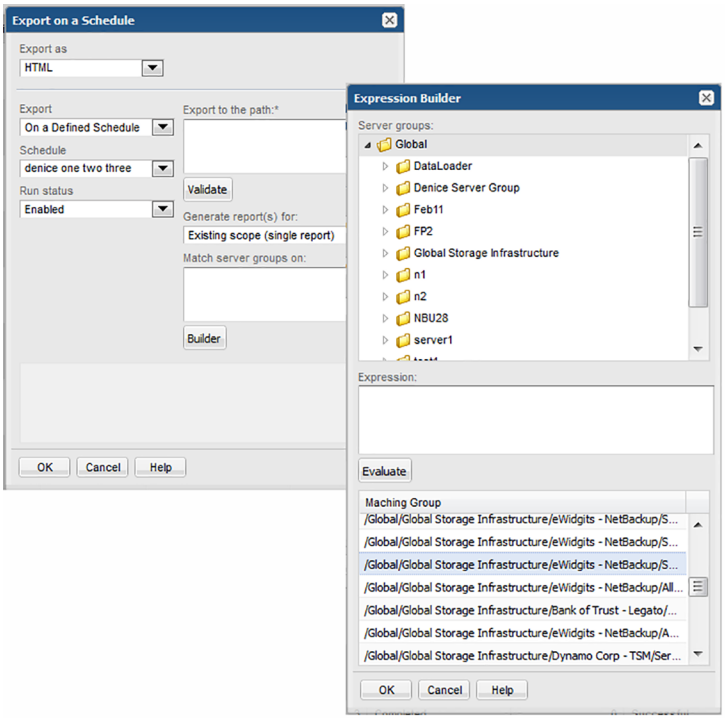
- CSV(**.csv**). Exports data to a spreadsheet or a database application such as Microsoft Access.
- Excel (.xlsx). Similar to the CSV export, this file supports MS Excel-specific features.
- XML (.xml). Accesses the web page components.

Note: XML is not available as an export type for dashboards.

- PDF (.pdf). This read-only file lends itself to easy distribution and printing.

- Text. A text file of the report/dashboard. This selection provides the option to set a delimiter and use quotes in the exported file.
- HTML (.html). Exports this data as a web page, along with supporting files, in a zipped file.

4. Click **Schedule**. When you schedule a report/dashboard to be exported on a regular basis, you must first configure a number of settings.



5. Define the schedule. Use the following table to configure a schedule:

Action	Steps
Export	■ **On a defined schedule**- Primary schedules can be configured and then applied to reports. Modifications to a primary schedule are applied to all reports associated with that schedule. ■ **Frequency in Minutes** - Select 5, 15 or 30 minutes. ■ **Hourly** - Select 1, 2, 3, 4, 6, or 12 hours. ■ **Daily** - **At: hour/minute**. Select a specific run time. ■ **Weekly**: ■ **On every**. Check the day(s) on which the report will run. ■ **At: hour/minute**. Select a specific run time. ■ **Monthly**: ■ **On the**. Select the day on which the report will run. ■ **At: hour/minute**. Select a specific run time. ■ **Cron Expression** - Enter a CRON expression to fine tune the export schedule.

Action	Steps
Export to the path*	Specify a sub-folder where the output will be saved. This sub-folder is relative to the default folder and it will be created if it does not already exist. If this field is left blank, the report will be saved in the base folder. The default base folder is set to: Linux: **/opt/aptare/export** Windows: **C:\opt\aptare\export** ■ For security reasons, the base folder for file exports is configured in the Portal under **Admin** > **Advanced** > **System Configuration** > **Portal** > **Report** > **Report export path**. Exporting to a network share requires that the share is already mounted on the Portal server and also configured in the Portal UI under **Admin** > **Advanced** > **System Configuration** > **Portal** > **Report** > **Report export path**. Note that this default base folder may have been changed by a Portal administrator. ■ The sub-folder name that you enter is appended to the base folder path and the report is saved with a filename that matches the report name. Invalid characters will be replaced by an underscore. ■ To create a unique filename that will not be overwritten, use variables: ${month}-${year}-${hour}-${minute}. This will create a file in the base folder with a name similar to: 12-2015-21-1. ■ To separate the filename from the folder, use double slashes, similar to: myfolder//JobSummary-12-2015-21-1. FOR ADVANCED USERS: In some cases (for example, Managed Services Partners with hundreds of managed host groups), you may want to export to specific host group or report directories. Rather than configure each of these separately, use the following variables in the path specification to generate an individual report for each instance of the variable: ${hostGroup} For example: /opt/aptare/export/${hostGroup}/doc ${reportName} To overwrite/replace files each time the reports are exported, use the variables: ${day}, ${month}, ${year}, ${hour}, ${minute} where the day, month, year, hour, and minute are represented as 2-digits, such as 02.
Run Status	Select **Enabled** or **Disabled**. This selection enables or disables the schedule for the report or dashboard to be exported.

Action	Steps
Generate reports for	Choose one of the following: ■ **Existing scope** - Creates a single report/dashboard based on the report's scope. ■ **Matching host groups** - Creates multiple reports/dashboards, based on the list that results from the Expression Builder. See "Using the Expression Builder" on page 176.

6. Click **OK**. If you are scheduling export of a saved report, your configuration is saved. Instead, if you are scheduling export of a report or dashboard template, the **Save Template as Report** pop-up is displayed.
7. Enter the report name and select the folder (default **/Home/My Reports**) where the report instance can be saved.

   Use **New Folder** to save the report instance in a new folder.
8. Click **OK** to save the export schedule.

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 IT Analytics Portal tech note for further details.

Using the Expression Builder

1. Click **Builder** to create a regular expression to be used for searching for host group names with pattern matching.
2. Click **Evaluate** to view which host groups are selected.

   Examples:

   - .*NetBackup.* - the period specifies any character; the asterisk specifies zero or more of the preceding element.
   - [abc] - simple alpha character match
   - [a-zA-Z] - any alpha character in upper or lowercase
   - [^abc] - any character except a, b, or c
   - The Expression Builder is case-sensitive.

Email reports and dashboards

Note: On a Linux Portal server, to ensure proper rendering of reports that are emailed or exported as HTML images or PDF files, a graphics manager such as X Virtual Frame Buffer (XVFB) is required. Contact your IT organization to configure this capability, if you plan to export/email reports as HTML images or as PDF files.

After you generate a report and it renders, you can choose to instantly email the report. You can email a report to yourself, other individuals, or a distribution list. You can also schedule a report to be mailed regularly.

See "Schedule email of reports and dashboards" on page 180.

Emailing exported data is limited to the first 20,000 rows. Reports longer than 20,000 rows are truncated when exported for email. The report output (PDF, HTML, CSV, etc.) displays the message "Data in this report has been truncated." You can revise default number of rows by modifying the portal.properties file.

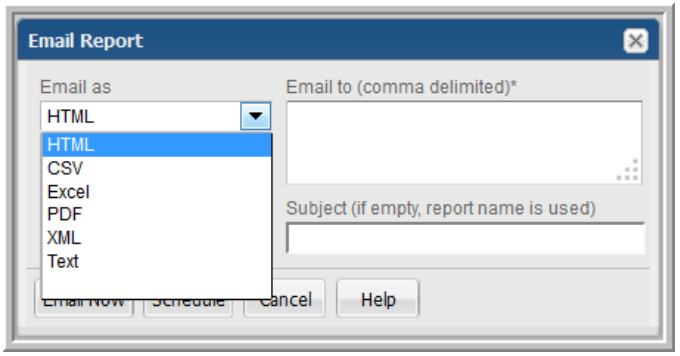
See "Known Limitations for Donut Charts" on page 170.

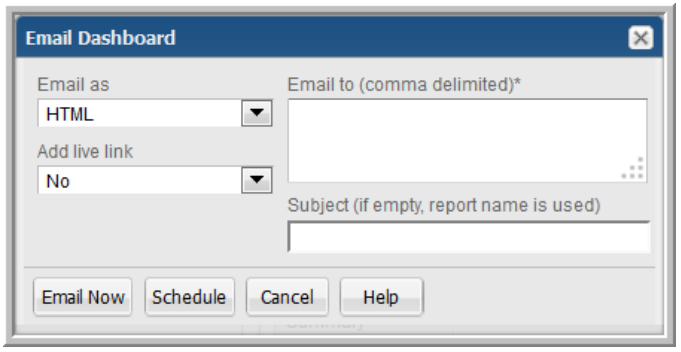
To email reports or dashboards

1 Generate the report or dashboard.

See “Generate reports” on page 137.

2 Click **Actions** and select **Email**. The **Email Report** or **Email Dashboard** dialog box is displayed.





You can email the report or dashboard immediately.

3 In the **Email Report** dialog box, specify your delivery settings:

Email as	■ HTML (.html). Use this format if you want to export this data as a web page. ■ CSV(**.csv**). Use this format to export tabular data to a spreadsheet or a database application such as Microsoft Access. ■ Excel (.xlsx). Similar to the CSV export, this file supports MS Excel-specific features. ■ PDF (.pdf). This read-only file lends itself to distribution and printing. ■ XML (.xml). Accesses the web page components. **Note:** XML is not available as an export type for dashboards. ■ Text (.txt). Produces a basic text file of the report. A text file of the report/dashboard. This selection provides the option to set a delimiter and use quotes in the emailed file.
Email to	Provide a comma-separated list of email addresses. Verify all email addresses are valid.
Subject	Specify an email subject if you want to override the report title. To use date variables in the subject, see the *Using date variables in the subject* section below.
Add live link	You can add a link to the report within the Portal. After logging in, a user can view the report in the Portal. In environments where SSL is enabled, a configuration change is required to ensure that this link is secure. See the System Administrator's Guide for details on modifying the **portal.properties** file.

Using date variables in the subject

You can include the report's run date and time in the email subject using variables. This is useful for quickly identifying scheduled emails (and their attachments) by date.

Enter variables in the **Subject** field, for example: Backup DSR - ${day}${month}${year}

When the report runs on 25 January 2024, the email subject becomes: Backup DSR - 25-01-2024

Available variables:

Table 6-1 Available date variables

Variable	Description	Example
${day}	Day of month, 2 digits	25
${month}	Month, 2 digits	01
${year}	Year, 4 digits	2024
${hour}	Hour, 24-hour format, 2 digits	21
${minute}	Minute, 2 digits	05

The values reflect the date and time the report is run (Portal server time).

Attachment file names

When the email includes the report as an attachment, the attachment file name is derived from the subject, so the same date values appear in the file name (invalid file-name characters are replaced with an underscore character). Example attachment: `Backup_DSR-25-01-2024.csv`.

In-dialog help

A hint listing the available variables appears at the bottom of the dialog when you click in the **Subject** field. Hovering the field also shows a tooltip.

Schedule email of reports and dashboards

By scheduling the email of a report or dashboard, you configure it to run and then email it at a specified frequency. You can schedule emailing of saved reports and dashboards. To schedule an email of a report or a dashboard template, you must generate and save its instance that can generate the report and email it based on the prescribed schedule.

Note: Emailing exported data is limited to the first 20,000 rows. Reports longer than 20,000 rows are truncated when exported for email. The report output (PDF, HTML, CSV, etc.) displays the message “Data in this report has been truncated.”

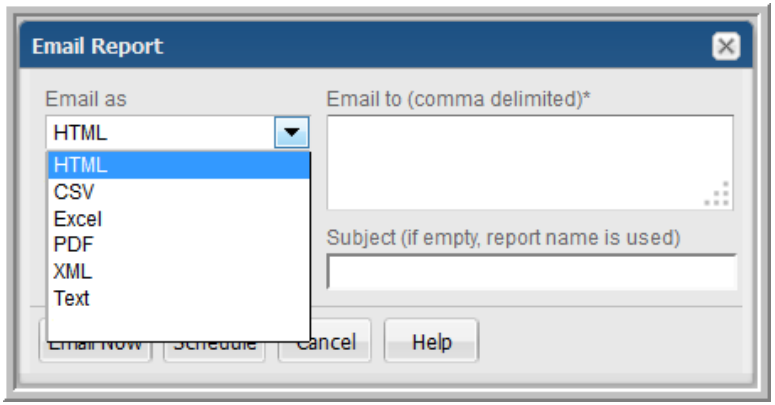
See “Known Limitations for Donut Charts” on page 170.

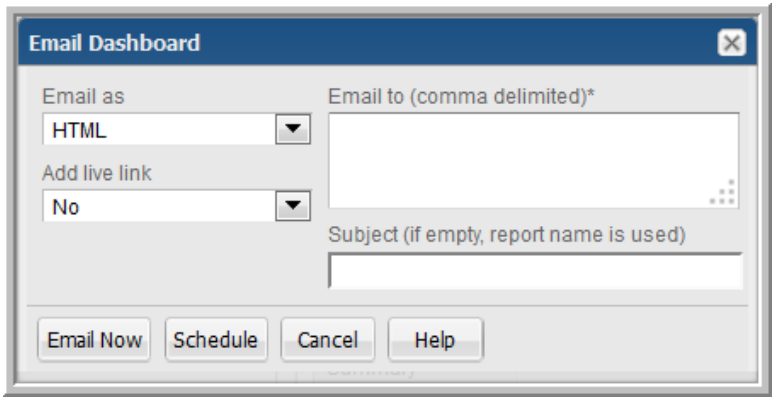
To email a report or dashboard at regularly scheduled intervals

1 Generate a report/dashboard. You can save it if required. To schedule saved reports/dashboards by email, skip to the next step.

See "Generate reports" on page 137.

2 Click **Actions** and select **Email**. The **Email Report** dialog box is displayed.

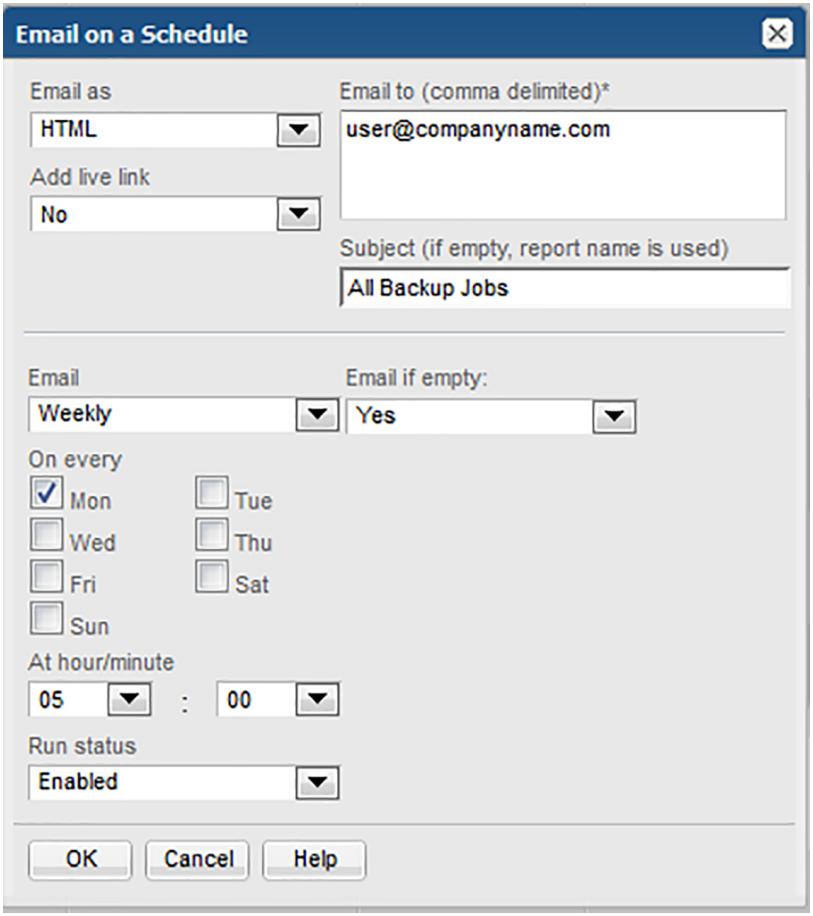




See "Email reports and dashboards" on page 177.

for email parameter descriptions.

3 Click **Schedule**. The **Email on a Schedule** dialog box is displayed.



4 Specify your delivery settings:

Email as	■ Choose the preferred output format: HTML, CSV, Excel, Text, HTML image, PDF, or XML (.xml). Selecting Text provides the option to set a delimiter and use quotes in the emailed file. **Note:** XML is not available as an export type for dashboards.
Email to	Provide a comma-separated list of email addresses.
Add live link	You can add a link to the report within the Portal. After logging in, a user can view the report within the Portal.

Subject	Specify an email subject if you want to override the report title.
Email	■ **On a defined schedule**- Master schedules can be configured and then applied to reports. Modifications to a master schedule will automatically be applied to all the reports associated with that master schedule. See "Configure primary schedules" on page 606. ■ **Frequency in Minutes** - Select from 5, 15 or 30 minutes. ■ **Hourly** - Select 1, 2, 3, 4, 6, or 12 hours. ■ **Daily** - **At: hour/minute**. Select a specific run time. ■ **Weekly** - ■ **On every.** Check the day(s) on which the report will run. ■ **At: hour/minute**. Select a specific run time. ■ **Monthly** - ■ **On the**. Select the day on which the report will run. ■ **At: hour/minute**. Select a specific run time. ■ **Cron Expression** - Enter a CRON expression to fine tune the email schedule. for details about working with CRON expressions.
Email if empty	This option is available only when you are scheduling a tabular report to be emailed at regular intervals. Sometimes reports might not have any data. If you don't want to email empty reports, choose **No**.
Run Status	Select **Enabled** or **Disabled**. This selection enables or disables the schedule for emailing the report.

5 Click **OK**. If you are scheduling email of a saved report, your configuration is saved. Instead, if you are scheduling email of a report or dashboard template, the **Save Template as Report** pop-up is displayed.

6 Enter the report name and select the folder (default **/Home/My Reports**) where the report instance can be saved.

Use **New Folder** to save the report instance in a new folder.

7 Click **OK** to save the email schedule.

Viewing and Managing Your Scheduled Reports

User type defines the access to scheduled reports:

- Non-Administrators can view, reschedule, delete and modify the export format for their own scheduled reports and those that belong to other users in their home group.

- As an Administrator, if you need to take an action on another user's schedule reports, use the

  where you can view, reschedule, delete and modify the export format for allusers' scheduled reports. For example, if a user leaves the company, but her scheduled reports continue to be emailed, you can delete the report's schedule.

1. Select a report and click **Edit** to modify the export format or the schedule.

Schedule Types

The Schedule column in the Scheduled Reports Administration window lists the type of schedule for each scheduled report:

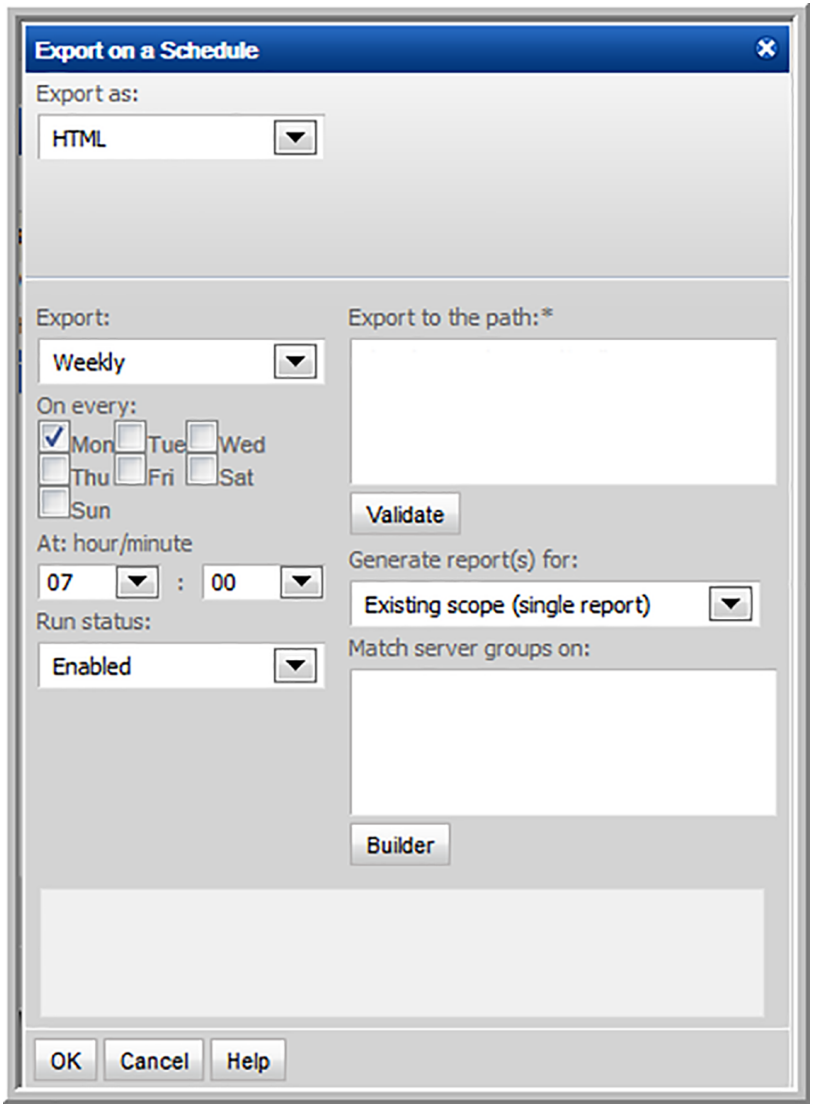
- Primary - Use primary schedules for your common report scheduling requirements. These schedules can be easily referenced when scheduling a report by selecting On a defined schedule.

  See "Export reports and dashboards" on page 169.

  See "Email reports and dashboards" on page 177.

  To create a primary schedule, go to: **Admin > Reports > Schedules**.

- One off - This schedule type refers to scheduling options other than a primary schedule. Examples of this schedule type include Frequency in Minutes, Hourly, Daily, Weekly, and Monthly.



Note: The Scheduled Reports Administration options will differ depending on how the report was scheduled, either as an email or a file export.

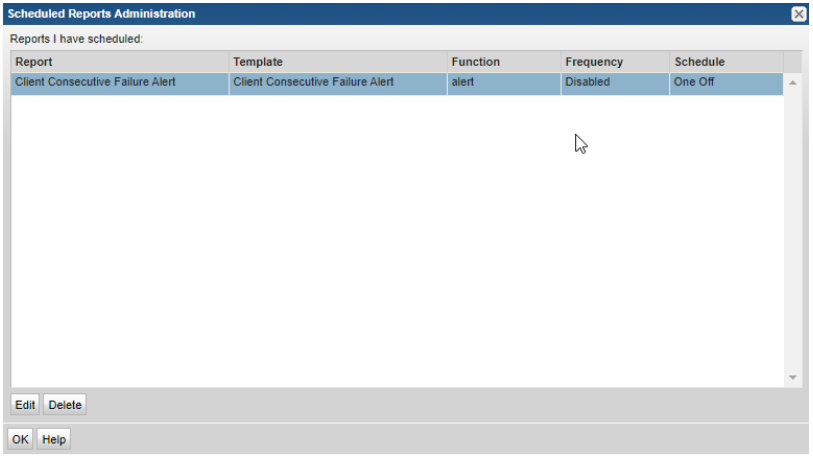
See "Schedule email of reports and dashboards" on page 180.

See "Schedule export of reports and dashboards" on page 171.

Deleting Scheduled Reports and Dashboards

To delete a scheduled report or dashboard

1 Click the User Account menu and select **My Scheduled Reports**.



Your scheduled reports are displayed in the **Scheduled Reports Administration** window.

2 Select reports or dashboard and click **Delete**.

Setting Up Alerts for Tabular Reports

Use the report-based alerting feature to **notify you when a tabular report has been populated with data**. For example, save a Job Summary report for Failed Events and then configure an alert for this report. The Portal will check, according to the schedule you select, for a report that contains data.

An alert can be delivered via the following mechanisms, described in detail later in this section:

- Email
- Script
- SNMP
- Native log

Note: Alerting is configured at the Domain level.

See "Add/Configure a domain" on page 766.

Configure an alert

You can configure alerting for any report that contains a single table. A report must be saved before you can set an alert.

1. Generate a saved report.

   See "Generate reports" on page 137.

2. Click **Actions** and select **Alert**.
3. Use the following table to configure an alert:

Check for alerts (schedule)	To schedule a report to be checked for alerts on a regular basis, you must select one of the following options from the drop-down list: ■ **On a defined schedule**- Master schedules can be configured and then applied to reports. Modifications to a master schedule will automatically be applied to all the reports associated with that master schedule. See "Configure primary schedules" on page 606. ■ **Frequency in Minutes** - Select from every 5, 15, or 30 minutes. ■ **Hourly** - Select 1, 2, 3, 4, 6, or 12 hours. ■ **Daily** - **At: hour/minute**. Select a specific time. ■ **Weekly** - ■ **On every**. Check the day(s) on which the report will be checked. ■ **At: hour/minute**. Select a specific time. ■ **Monthly** - ■ **On the**. Select the day on which the report will be checked. ■ **At: hour/minute**. Select a specific time. The Portal can check the same report multiple times in a single day. ■ **Cron Expression** - Enter a CRON expression to fine tune the alert schedule. for details about working with CRON expressions.

Run Status	Select **Enabled** or **Disabled**. This selection enables or disables the schedule for the report to be checked for alerts.
Time Zone	Select the time zone from the list with which the schedule can align. If you do not select a time zone, the schedule will run based on the Portal server time zone.
Email	Select the checkbox and enter comma-separated email addresses.
Subject	Enter the subject. The report name is used if the field is left blank.
Script	The user-created script needs to reside in: `/opt/aptare/portal/user_scripts` Check the box and enter a shell script name (Linux). If a path name is included, it will be appended to the above path. For example, filter a report to include only the columns of information that you need. When the alert is triggered, a .csv file of the report is generated and the path to that file is made available to the script to take whatever actions you want with this report data. Administrators: To enable Script delivery, refer to the following. See "Add/Configure a domain" on page 766.
SNMP	When you check this box, the Port, Community string, and Management servers fields will be populated from the configured defaults. To override the defaults, overwrite any or all of the three SNMP fields. Administrators: To enable SNMP delivery and to configure SNMP default values refer to the following. See "Add/Configure a domain" on page 766.

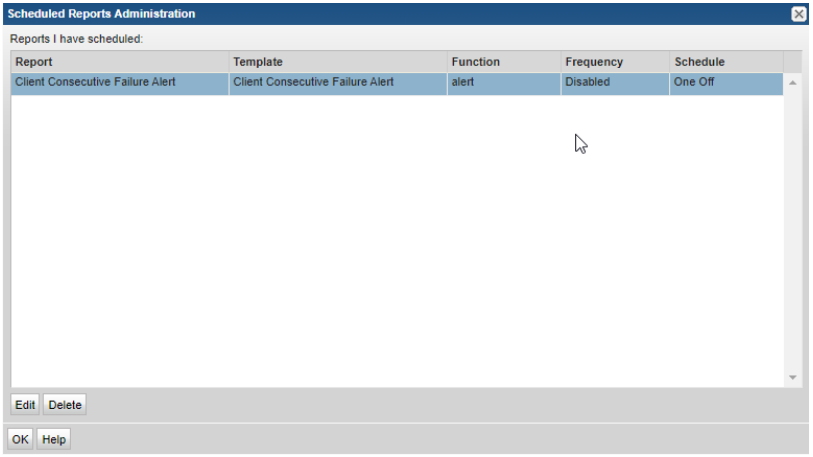
Native Log	When this box is checked, a log entry is written to the OS-specific log: either the Windows event log or the Linux syslog. Administrators: To enable Native Log delivery refer to the following. See “Add/Configure a domain” on page 766.

4. Click **OK** to save the alert configuration.
5. Click the User Account menu and select **My Scheduled Reports** to view the alerts. For details on managing alerts in the Scheduled Reports list refer to the following.

   See “Viewing and Managing Your Scheduled Reports” on page 183.

Deleting Report-Based Alerts

1. Click the User Account menu and select **My Scheduled Reports**. The **Scheduled Reports Administration** dialog is displayed.



2. Select the required alert.
3. Click **Delete**.

Share Reports, Dashboards, and Folders

Sharing is a privilege based on ownership. Reports and dashboards must be saved before they can be shared. You can share a report or dashboard with any number

of users or user groups. Reports in a shared folder inherit the sharing properties of the folder. You do not need to set the sharing properties on each report. The folder properties will append any properties set on the individual reports.

As an example, report A is shared with users, B and C. Folder A is shared with users, D, E and F. If you add report A to folder A, it becomes shared with users B, C, D, E, and F. You can add reports to a folder at any time.

Use the **Reports** tab to view sharing status associated with reports and folders. For report and folder creators, this information can be useful to determine the impact of stopping the sharing, modifications to the report, or deletion. Shortcuts cannot be shared.

Note: If you have many reports to share, you can add these reports to a user folder and share the entire folder. All reports within the folder inherit the assigned sharing properties.

Share a report, dashboard, or folder

The **Share** dialog allows to share and unshare. A report/folder shared with others will have the check boxes selected when the dialog opens.

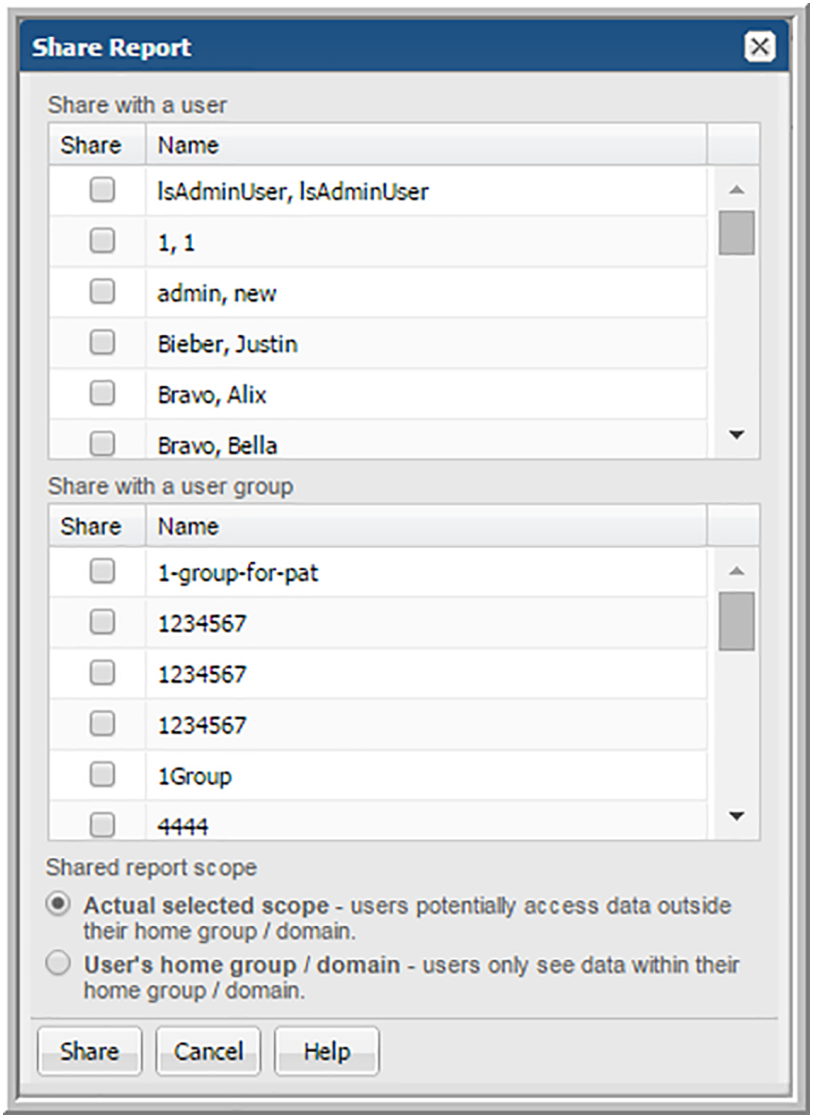
To share a report, dashboard, or folder:

1 Select the saved report or dashboard in the **Reports** tab, click **Share** on the Action bar.

OR

From the navigation panel, select a user folder and click **Share**.

The dialog displays a list of users and user groups.



2 Check any number of users or user groups.

3 Check the shared report scope:

- **Actual selected scope** - When this option is selected, users may be able to access data that is not in their home group.
- **User's home group/domain** - This option limits access to data within the user's home group and domain. For example, a Managed Services Provider may create a report and share it with multiple clients in different home groups. These users will only be able to access the data for their specific home group.

Note: Custom reports, created with either the SQL Template Designer or Dynamic Template Designer are always restricted to the user's home group.

---

4 Click **Share**.

Add users or groups to shared report, dashboard, or folder

You can edit by adding or removing users/groups (by clearing the checkbox).

To add users or groups to shared reports, dashboards or folders:

1 Select the shared report or dashboard in the **Reports** tab, click **Share** on the Action bar.

OR

From the navigation panel, select the shared user folder and click **Share**.

The dialog displays a list of users and groups with those already shared with selected.

2 Check any number of users or user groups.

3 Check the shared report scope:

- **Actual selected scope** - When this option is selected, users may be able to access data that is not in their home group.
- **User's home group/domain** - This option limits access to data within the user's home group and domain. For example, a Managed Services Provider may create a report and share it with multiple clients in different home groups. These users will only be able to access the data for their specific home group.

Note: Custom reports, created with the SQL Template Designer or Dynamic Template Designer, are always restricted to the user's home group.

4 Click **Share**.

The shared folder and /or the report start appearing under the **Shared by You** folder in the navigation panel.

Working with Shared Reports

You can share saved reports and user folders with other users or user groups. Similarly, other users can share reports and user folders with you.

View reports, dashboards, and folders shared with others

As the report/folder creator, you can modify and delete your shared reports and folders. Use the **Reports** tab to view properties associated with shared elements. This information can be useful to determine the impact of modifications or a deletion. Reports listed in the **Reports** tab indicate if sharing is active, and identifies who the report is shared with.

To view a list of reports, dashboards and user folders shared with others:

1 Click **Reports**.

2 Go to **Home** > **Shared by You** in the navigation panel.

View reports and folders that other users shared with you

As the report/folder recipient, you are limited in the actions you can perform on the report or folder. Shared reports can be emailed, exported or can trigger an alert. You cannot edit or delete shared reports or folders.

If the report creator updates the report, it is automatically updated in your **Shared with You** section.

When Home Pages are defined for you through User Groups, those reports and dashboards are displayed as shared reports under the **Shared with You** section. Once reports are assigned as Home Pages by a system administrator, they can be removed or resorted by a user. When a report has been removed as a user's Home Page, the user can still view it, run it, and save a new version of it using the **Shared with You** section.

See "User Home Pages and User Group Home Pages" on page 212.

To view a list of reports and user folders shared with you:

1. Click **Reports**.
2. Go to **Home** > **Shared with You** in the navigation panel.

Stop sharing/Unshare reports and dashboards

Use the **Reports** tab to view properties associated with shared reports. This information can be useful to determine the impact of stopping the sharing, modifications to the report, or a deletion.

To stop sharing/unshare reports and dashboards:

1. Open the **Reports** tab.
2. Go to **Home** > **My Reports** in the navigation panel.
3. Select the report or folder you want to stop sharing.
4. Click **Share** on the Action bar and clear the checkbox against the user or user group with whom you want to stop sharing the report or dashboard.

Chapter 7

Organize reports

This chapter includes the following topics:

- Work with dashboards
- Understand the dashboard Icons
- About Cached Reports on Dashboards
- Create new dashboards and add reports
- Add reports to an existing dashboard
- Save a dashboard with a different name
- Customize a Dashboard
- Edit the scope for a report on a dashboard
- Edit the dashboard scope
- Expand a report to full size on a dashboard
- Set the dashboard layout
- Delete a dashboard
- Remove reports from dashboards
- Create a custom report folder
- Create shortcuts to reports and templates
- Manage My Home Pages
- User Home Pages and User Group Home Pages

Work with dashboards

You can drag and drop individual reports, to create a consolidated view to contain reports you access regularly, or view multiple scenarios simultaneously. Dashboards can also help you troubleshoot a particular problem by gathering data from multiple reports into one view.

- See “Understand the dashboard Icons” on page 196.
- See “About Cached Reports on Dashboards” on page 197.
- See “Create new dashboards and add reports” on page 197.
- See “Add reports to an existing dashboard” on page 200.
- See “Save a dashboard with a different name” on page 201.
- See “Edit the scope for a report on a dashboard” on page 204.
- See “Set the dashboard layout” on page 208.
- See “Delete a dashboard” on page 209.
- See “Remove reports from dashboards” on page 209.

Understand the dashboard Icons

Once placed on a dashboard, you can revise the scope or remove the report entirely. These functions are displayed when you hover your cursor over the report title. The placement varies depending on the report type, but icons are displayed in the top right corner of your report portlet. You can change a report location by simply dragging it to a new position.

See “Interact with Reports” on page 159.

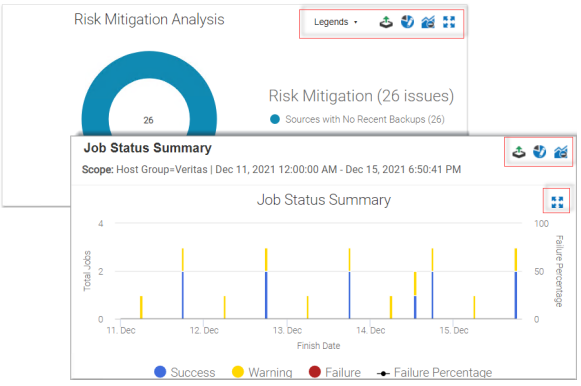


Table 7-1 Dashboard report icons

	Edit Scope - Allows you edit the scope of the individual report portlet and rerun the report without impacting the entire dashboard. See "Edit the scope for a report on a dashboard" on page 204. for details.
	Remove Report from Dashboard - As you use your dashboards, you may find that some reports are no longer relevant. These reports can be quickly removed from your dashboard. See "Remove reports from dashboards" on page 209.
	Expand - Charts displayed as components of a dashboard may be hard to read. Click the Expand icon to pop the chart out of the dashboard and have it displayed in the full browser window. See "Expand a report to full size on a dashboard" on page 207.
Legends menu	The menu provides options to define pagination with preset or custom values.

About Cached Reports on Dashboards

When you run a report, the console takes the scope of the report, and checks if the cache contains the same report, for the same scope. If it does, the results are displayed from the cache. If the combination does not exist, the report is run from the database, saved in the cache, and then sent to the user interface.

When a report is served from the cache, an indicator icon is displayed on reports and dashboards. You can roll over the indicator to show the age of the report from the cache. Click the icon to rerun the report from the database.

See "How reports and caching work together?" on page 839.

Create new dashboards and add reports

Dashboards provide a custom, at-a-glance overview by displaying reports you choose on a single page.

See "Add an External Reference" on page 413.

Dashboards created using this procedure can be modified using the **Actions > Customize** menu. See "Customize a Dashboard" on page 203.

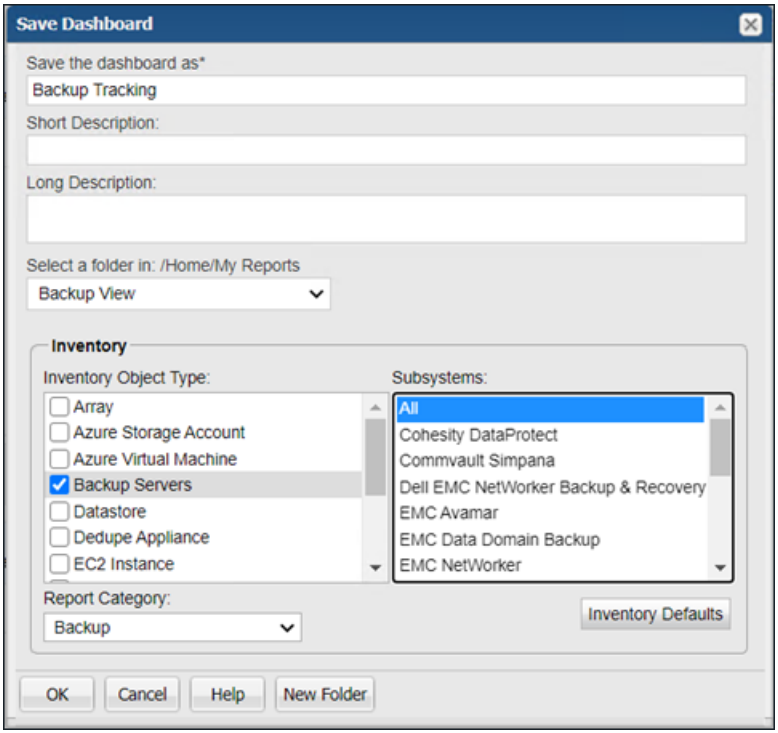
To create a dashboard and add reports:

1. Search and generate each report you want on your dashboard. Each report will display in a separate tab.
2. Click **Create Dashboard**. A new blank dashboard is displayed.
3. Click and drag the tab of each report you want to place on the dashboard. You can add multiple reports to a custom dashboard, but you are limited to four horizontal columns. You can resize row heights at any time using Set **Layout**.

   Note: A report that has been generated as a result of a drill down in another report cannot be saved or added to a dashboard because of the dependency with the parent report.

   ---

4. As you drag the report on to the dashboard, areas are highlighted indicating the potential positions of the report. You can change the report's location by dragging it to a new position.
5. Click **Actions** > **Save**.
6. Specify the details on the **Save Dashboard** pop-up as described below:



- **Save the dashboard as**: Assign a name to the dashboard. The maximum number of characters is 80.
- **Short Description**: Provide a description that can help in searching the dashboard. The description can help to determine the intent of dashboard. This description is limited to 512 characters and is optional.
- **Long Description**: Enter any additional details that you wish to specify regarding the dashboard. The details may included vendor-specific names or acronyms that may help in searching the dashboard. This description is optional.
- **Select a folder in: /home/My Reports**: Dashboards are saved folders created inside `/home/My Reports`. Select a folder from the list.
- **Inventory Object Type**: The inventory objects associated with the reports added on the dashboard are selected by default. You can add or remove the objects to override the default selection.
- **Subsystems**: Specify one or more subsystems from which you want to view data on the dashboard. By default, the appropriate subsystem is selected.

- **Report Category**: Assign a report category from the list. If you select a report category, you must set the **Inventory Object Type** and **Subsystems** values as well.
- **Inventory Defaults**: Click to reset the default **Inventory** settings. This resets the default values under **Inventory Object Type** and **Subsystems** fields.
- **New Folder**: Click to add a new folder under `/home/My Reports` and save the dashboard in it.

Note: Reports on dashboards are treated as stand-alone instances and are not linked to the original report. For example, when a name or scope change is made in the original report, those changes are not reflected in the dashboard version.

Add reports to an existing dashboard

You can quickly add new reports to your dashboard as your business needs change.

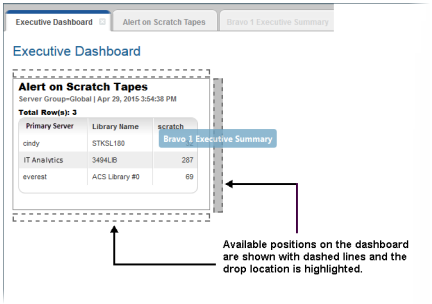
To add reports to an existing dashboard

1. Search for the dashboard name using **Search**.
2. Run the dashboard directly from the search results. The updated dashboard is displayed.
3. Search and generate each report you want to add to your dashboard. Each report will display in a separate tab.
4. Click and drag the tab of each report you want to place on the dashboard. You can add an unlimited number of reports to a dashboard, but you are limited to four horizontal columns. You can resize row heights at any time using **Set Layout**.

Note: A report that has been generated as a result of a drill down in another report cannot be saved or added to a dashboard because of the dependency with the parent report.

---

5 As you drag the report on to the dashboard, areas are highlighted indicating the potential positions of the report. You can change their location by dragging them by the title to a new position.



6 Click **Actions** and select **Save**.

Save a dashboard with a different name

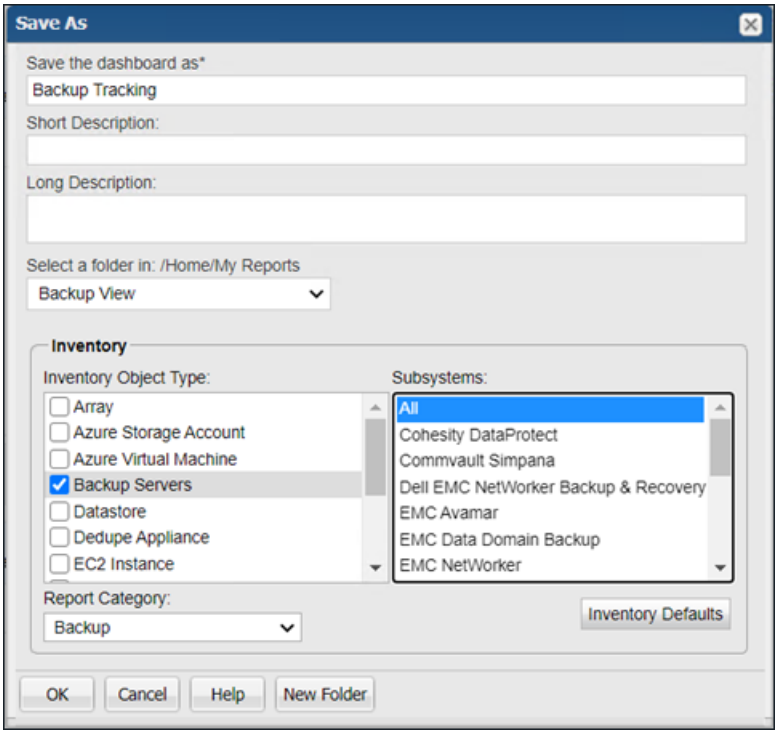
Once you have created a dashboard you like, you can make modifications and save the dashboard with a different name.

To save a dashboard with a different name

1 Select and run a saved dashboard.

2 Click the **Actions** menu and select **Save As**.

3 Change the details on the **Save As** pop-up as required based on the descriptions below:



- **Save the dashboard as**: Assign a new name to the dashboard if required. The maximum number of characters is 80.
- **Short Description**: Add or edit the description that can help in searching the dashboard. The description can help to determine the intent of dashboard. This description is limited to 512 characters and is optional.
- **Long Description**: Add or update any additional details that you wish to specify regarding the dashboard. The details may included vendor-specific names or acronyms that may help in searching the dashboard. This description is optional.
- **Select a folder in: /home/My Reports**: Dashboards are saved folders created inside `/home/My Reports`. Select a folder from the list if required.
- **Inventory Object Type**: The inventory objects associated with the reports added on the dashboard are selected by default. You can edit the selection to override the default settings.
- **Subsystems**: Specify one or more subsystems from which you want to view data on the dashboard. By default, the appropriate subsystem is selected.

- **Inventory Defaults**: Click to reset the default **Inventory** settings. This resets the default values under **Inventory Object Type** and **Subsystems** fields.
- **Report Category**: Change the report category from the list if required. This field is mandatory.
- **New Folder**: Click to add a new folder under `/home/My Reports` and save the dashboard in it.

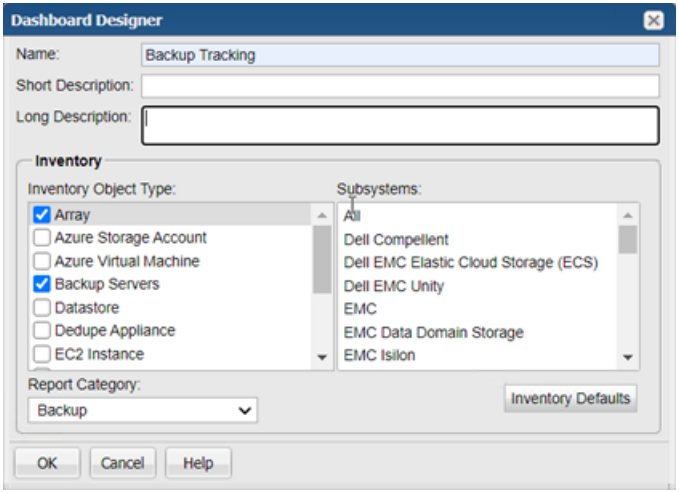
Note: Reports on dashboards are treated as stand-alone instances and are not linked to the original report. For example, when a name or scope change is made in the original report, those changes are not reflected in the dashboard version.

Customize a Dashboard

You dashboard may require modifications over a period of time. This procedure describes how to customize your dashboard to suite your requirement.

To customize a dashboard:

1. Select and run a saved dashboard.
2. Click the **Object Actions > Customize**.
3. Change the details on the **Dashboard Designer** pop-up as required based on the descriptions below:



- **Name**: Assign a new name to the dashboard if required. The maximum number of characters is 80.
- **Short Description**: Add or edit the description that can help in searching the dashboard. The description can help to determine the intent of dashboard. This description is limited to 512 characters and is optional.
- **Long Description**: Add or update any additional details that you wish to specify regarding the dashboard. The details may included vendor-specific names or acronyms that may help in searching the dashboard. This description is optional.
- **Inventory Object Type**: The inventory objects associated with the reports added on the dashboard are selected by default. You can edit the selection to override the default settings.
- **Subsystems**: Specify one or more subsystems from which you want to view data on the dashboard. By default, the appropriate subsystem is selected.
- **Inventory Defaults**: Click to reset the default **Inventory** settings. This resets the default values under **Inventory Object Type** and **Subsystems** fields.
- **Report Category**: Change the report category from the list if required. This field is mandatory.

Note: Reports on dashboards are treated as stand-alone instances and are not linked to the original report. For example, when a name or scope change is made in the original report, those changes are not reflected in the dashboard version.

Edit the scope for a report on a dashboard

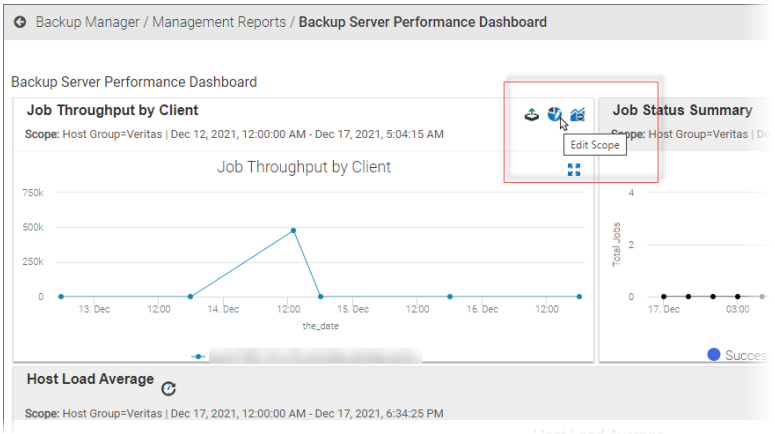
You can edit the scope for an individual report on a dashboard. This allows you to create what-if analyses or troubleshoot an issue by changing the data set and creating a preview report. Once you've changed the scope, regenerate the report. The report title is displayed in red italic to indicate a change has occurred. You can save your changes or just close the dashboard to retain the original scope.

See "Interact with Reports" on page 159.

To edit the scope for a report on a dashboard

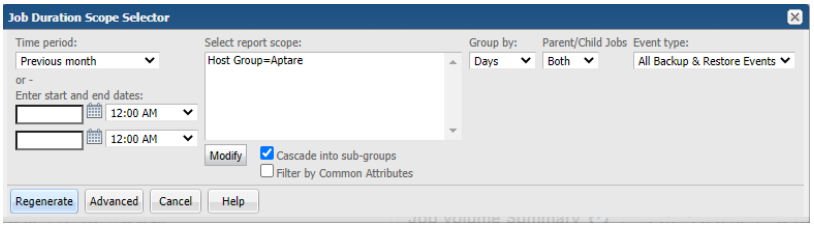
1. Search for the dashboard name using **Search**.
2. Run the dashboard directly from the search results. The updated dashboard is displayed.

3 Roll your cursor over the corner of a report on the dashboard to display the **Edit Scope** icon.

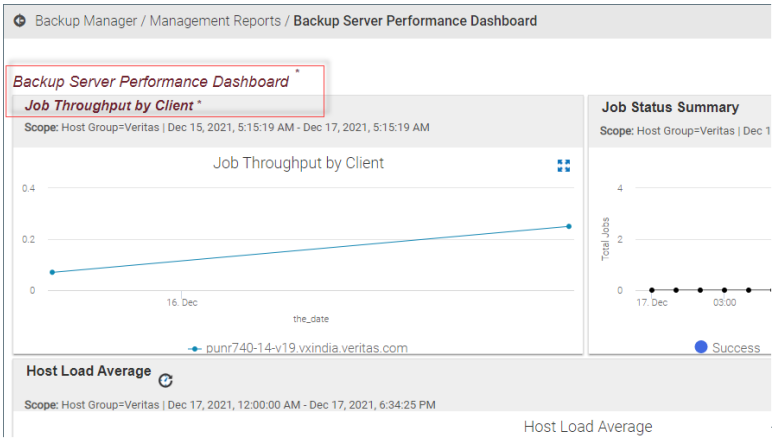


4 Click the **Edit Scope** icon to display the Scope Selector for the report.

5 You can apply three common designer components **Scope**, **Time Period**, and **Group by** to all portlets in the dashboard at dashboard level. If any portlet does not have that component, the setting will not be applicable to that specific portlet.



6 Regenerate the report. The report title and the dashboard name are shown in red italics to indicate a change has been made to the dashboard.



7 Save the changes to update the dashboard with the new scope. You can also close the dashboard to discard the scope changes and retain the original scope.

Edit the dashboard scope

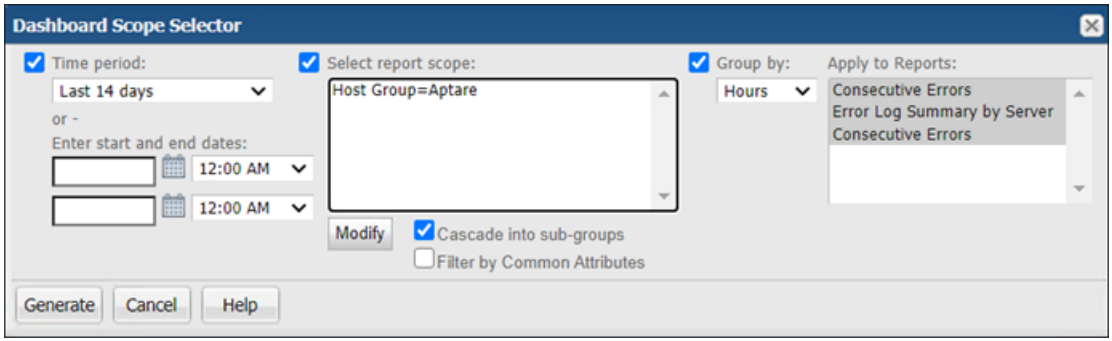
You can edit the scope of the report portlets added to your dashboard collectively through the Dashboard Scope Selector. This scope editing is can be most effective if your dashboard portlets contain data from the same or related sources, such as hosts, drives, arrays, VM Servers or Guests. If a scope component appears disabled, it means it is not applicable to the selected report.

To edit the dashboard scope:

1 Select and run a saved dashboard.

2 Click the **Actions > Edit Scope** menu.

3 Edit the settings on the **Dashboard Scope Selector** pop-up as required based on the descriptions below.



- **Time period**: The reports will display the data from the time period set in this component. You can specify number of days or a data range.
- **Select report scope**: Click **Modify** to edit the scope of the reports. You can add or remove objects such as hosts, arrays, drives, and datastores as applicable to the reports on your dashboard.
  - **Cascade into sub-groups**: Cascades the report scope settings to all the child sub-groups while generating the report when selected. To view the scope settings on the host group, clear this checkbox.
  - **Filter by Common Attributes**: Select to have the report scope display attributes using "AND" Logic. If left unchecked, the report displays 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.
- **Group by**: Groups the report results based on the time line specified in this field.
- **Apply to Reports**: Allows you to select the reports to which you want to apply the modified scope. By default, all reports are selected. For example, if you select a single report, the scope modifications will apply only on the selected report.

Expand a report to full size on a dashboard

The **Expand** icon pops the chart out of the dashboard and displays in the full browser window.

To expand a report on a dashboard

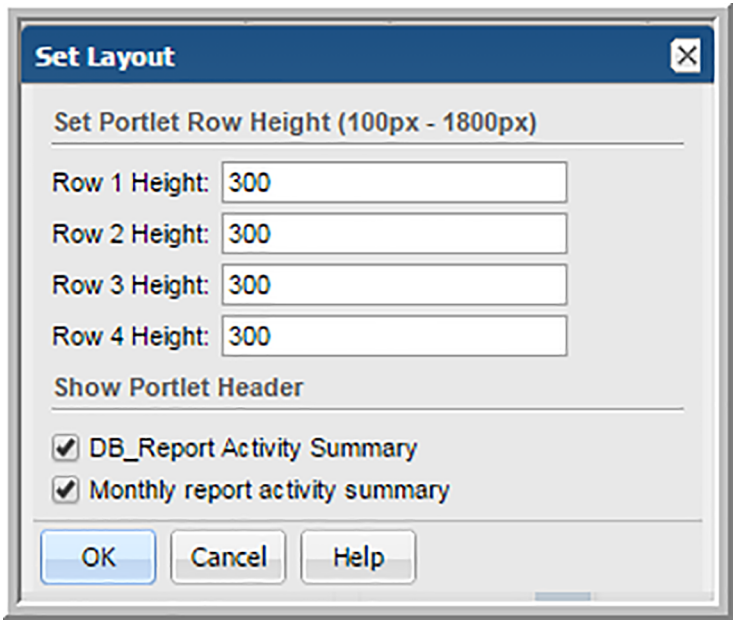
1. Search for the dashboard name using **Search**.
2. Run the dashboard directly from the search results. The updated dashboard is displayed.
3. Hover your cursor over the report title. The icons are displayed.
4. Click the **Expand** icon. The report portlet is displayed in full screen mode. Drill downs are still accessible.
5. Click the **Close** icon to return to the dashboard.

Set the dashboard layout

The **Set Layout** function allows you to resize row heights on the dashboard and to toggle the report headers on and off for an optimal visualization of your data.

To change the portlet row height on a dashboard

1. Select **Set Layout** from the **Actions** menu with your dashboard active.



2. Set the row height for reports grouped horizontally on the dashboard. Each row is controlled individually.
3. Enter a value per row. The range is 100 to 1800 pixels. The default is 300 pixels.

To show/hide portlet headers on a dashboard

1. Select **Set Layout** from the **Actions** menu with your dashboard active.
2. Deselect the reports to remove the header. By default, the header on the report is displayed.

Delete a dashboard

Use the information displayed in the **Reports** tab to help assess the impact of deleting a dashboard. For example, you can determine if you've shared your dashboard with other users.

To delete a dashboard

1. Locate and select the dashboard on the **Reports** tab.
2. Click **Delete** on the Action bar.

Remove reports from dashboards

As you use your dashboards, you may find that some reports are no longer relevant. These reports can be quickly removed from your dashboard.

To remove a report from a dashboard

1. Search for the dashboard name using **Search**.
2. Run the dashboard directly from the search results. The updated dashboard is displayed.
3. Hover your cursor over the report title. The icons are displayed.
4. Click the **Remove Report from Dashboard** icon. You can reposition the remaining reports by dragging them by the title to new locations.

Create a custom report folder

Once you accumulate several saved reports, you can better organize them by creating custom folders within the **My Reports** group. This user-defined folder enables you to access specific and unrelated reports quickly by creating a folder structure relevant to your environment. Because sharing reports is available from the folder level, you can also create a folder structure to share certain groups of reports instead of sharing individually.

See "Share Reports, Dashboards, and Folders" on page 189.

You can rename the **My Reports** folder if required.

To create a custom report folder

1. Right-click the **My Reports** folder and select **New**.
2. Enter the new folder name and description.
3. Click **Save**.
4. Drag saved reports into the new folder.

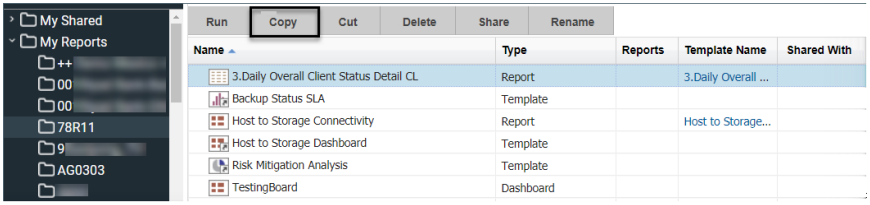
Create shortcuts to reports and templates

A shortcut is a link to a report or a template in a different folder.You can copy a report or template and then paste them as shortcuts into any user folder. This allows you to easily access the item associated with the shortcut. Shortcuts can be distinguished from the original file by the arrow that appears on the icon. After you paste the shortcut into its new location, you can rename it if required.

Note: Shortcuts cannot be shared.

To create a shortcut

1. Select the report or template.
2. Click **Copy**.
3. Navigate to the user folder you want to paste it into.



4. Click **Paste as Shortcut**.

   Once you paste the link, you can right-click and select **Rename** if required.

To delete a shortcut

1. Select the report or template shortcut.
2. Click Delete.

Manage My Home Pages

Home Pages are associated with login credentials and display when you log in to the product. When Home Pages are selected or assigned, they will always launch as the default display. Only the first five reports in your list of Home Pages on the **My Home Page Administration** dialog will launch automatically at login.

These designated home pages launch automatically each time you log in allowing you to:

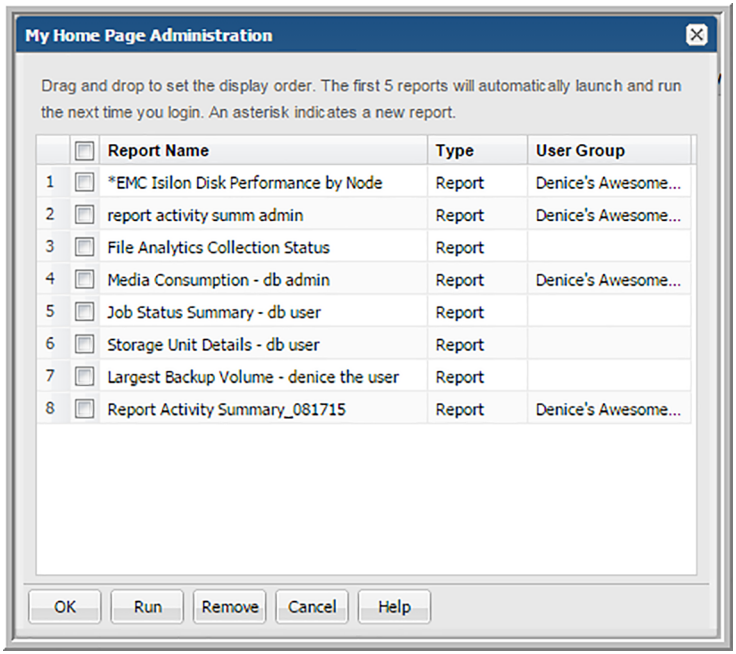
- Quickly access the reports you run most often.
- Easily revise or remove reports as your Home Pages.
- Set the display order to your own preference.
- Add an unlimited number of reports as your Home Pages

Note: If no Home Pages are selected or assigned, and you have privileges, the **Inventory** view is displayed.

The **My Home Page Administration** dialog allows you to reorder your home pages, remove your home pages and run selected reports.

See “Managing user group home pages (Administrator)” on page 599.

Note: Home pages assigned by Administrators through User Groups cannot be removed by individual users.



User Home Pages and User Group Home Pages

Users can create and add Home Pages. Home Pages can also be assigned with User Groups. When Home Pages are selected or assigned, they will always launch as the default display. When Home Pages are assigned through User Groups:

- Users can reorder the display order.
- Users cannot delete the reports.
- Reports are automatically added for each User Group member when updated.
- Reports are automatically removed for each User Group member when deleted.
- Reports are removed for each User Group member when the User Group is deleted.
- The first 5 reports are automatically added to the top of the user's Home Page list.
- New reports are flagged with an asterisk.

See "Managing user group home pages (Administrator)" on page 599.

for additional information.

To set reports or dashboards as your home pages

You can designate one or multiple home pages. You can add an unlimited number of reports as your Home Pages, however, only the first five reports listed on the **My Home Page Administration** will automatically launch at login.

1 Generate a saved report or dashboard.

2 Right-click and select **Add to My Home Pages**.

To remove a report or dashboard as your home page

If all home pages are removed, and you have privileges, the Inventory view is displayed. If you do not have privileges, the Reports view is displayed. The Remove action only removes the designation as a Home Page. The reports still exist within your Portal.

Note: Home pages assigned by Administrators through User Groups cannot be removed by individual users.

---

1 Click the User Account menu and select **My Home Pages**.

2 Select the reports to remove as your home pages.

3 Click **Remove**.

To set the order for your home page display

With Home Pages, if you have multiple reports, you can assign the order in which they are displayed. By default they are displayed in the order in which they were added as your home pages. If Home Pages are added through User Groups, they are automatically added to the top of the list.

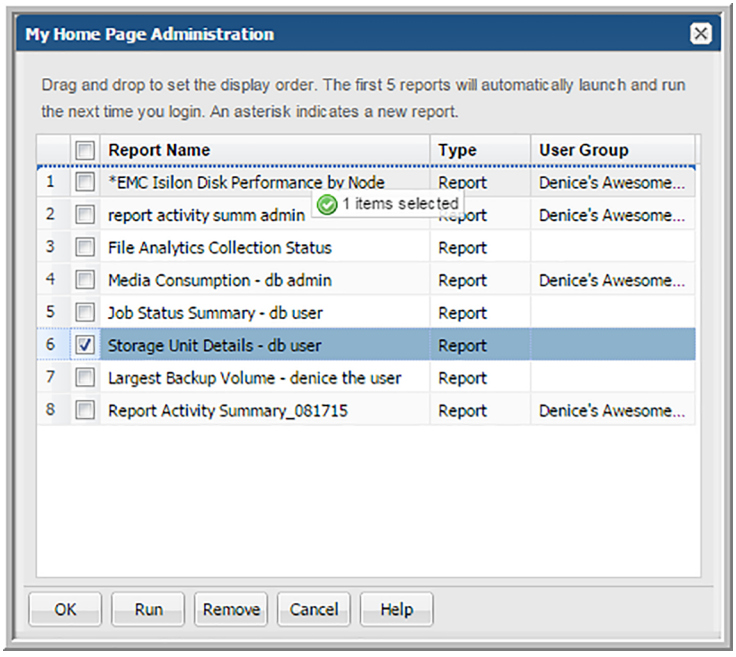
Note: You can add an unlimited number of reports as home pages, however, only the first five reports list on the **My Home Page Administration** dialog will automatically launch at login.

---

1 Click the User Account menu.

2 Select **My Home Pages**. On the **My Home Page Administration** dialog, you can drag and drop reports/ dashboards you've designated as home pages into the order you require.

3 Select the report, and drag and drop it into position. The top of the list displays the report in the first position (far left) in your home page display.



4 Click **OK**.

Chapter 8

Work with the dynamic template designer

This chapter includes the following topics:

- Dynamic Template Designer Overview
- Dynamic Template Designer Cheat Sheet
- Modify an Out-of-the-Box Dynamic Template
- Dynamic Template Designer Quick Start
- Steps to Create Dynamic Templates
- Working with Enterprise Objects and Fields
- Select an Enterprise Object
- Configure General Dynamic Template Designer Components
- Converting to a Homogeneous, Product-Specific Template
- Add Fields and Methods to a Dynamic Template
- Alias Names for Fields
- Conditions When Hidden Fields are Included by Default
- Dynamic Template Field Configuration
- Valid Formatter Patterns
- Examples of Dynamic Templates Containing Graphics Tool Tips
- Dynamic Template Drilldown Configuration

- Custom Drilldowns and Examples
- Drilldown Parameters
- Examples of Dynamic Templates Containing Drilldowns
- Using Groups in Dynamic Templates
- Add a Group to Create Separate Line Charts in a Report
- Add a Group to Create a Double Header in a Tabular Report
- Dynamic Template Function Configurations
- Examples of Dynamic Templates Containing Functions
- Create User-Defined Fields with the Field Builder
- Create Fields with the Field Builder
- Configure a Static Filter
- Filter on Date Fields
- Static Filter vs. Tabular Report Filter
- Configure the Field Sort Order
- Return Unique Results
- Dynamic Template Scope Selector Components
- Scope Selector Component - Custom Filter
- Data Domain Enterprise Object Scope Selector Components
- Host Enterprise Object Scope Selector Components
- Job Enterprise Object Scope Selector Components
- Storage Array Enterprise Object Scope Selector Components
- Customize and Export Dynamic Templates
- Export/Import Dynamic Templates Containing Custom Attributes
- Export a dynamic template that contains attributes
- Import a dynamic template that contains attributes
- Save a Dynamic Template After Edits
- Format the Dynamic Template Output

- Configure a Bar Chart Dynamic Template
- Steps to Create a Bar Chart Dynamic Template
- Examples of Bar Chart Dynamic Templates
- Configure an Area/Stacked Area Chart Dynamic Template
- Configure a Donut Chart Dynamic Template
- Example of a Donut Chart Dynamic Template
- Configure a Horizontal Bar Chart Dynamic Template
- Example of a Horizontal Bar Chart Dynamic Template
- Configure a Horizontal Stacked Bar Chart Dynamic Template
- Example of a Horizontal Stacked Bar Chart Dynamic Template
- Configure a Line Chart Dynamic Template
- Line Charts for Performance Metrics
- Steps to Create a Line Chart Dynamic Template
- Line Chart Field Requirements
- One Object Per Line Chart, One or More Metrics Per Chart
- Multiple Objects Per Line Chart, One Metric Per Chart
- Format Line Chart Fields
- Line Chart Scope Selections
- Examples of Line Chart Dynamic Templates
- Configure a Pie Chart Dynamic Template
- Examples of Pie Chart Dynamic Templates
- Configure a Stacked Bar Chart in a Dynamic Template
- Example of a Stacked Bar Chart Dynamic Template
- Configure a Table Dynamic Template
- Examples of Tabular Dynamic Templates
- Create a Sparkline Chart in a Tabular Dynamic Template
- Configure Chart Axes

- Using Methods in Dynamic Templates
- Overview of Method Creation
- System Methods
- Stand-Alone Methods
- Viewing Methods
- Adding or Editing Methods
- Tokens for Methods
- Exception Tables for a Method's Join Clause
- Validate and Save a Method
- Delete a Method
- Troubleshooting Dynamic Templates

Dynamic Template Designer Overview

The Dynamic Template Designer is a tool that does not require Structured Query Language (SQL) knowledge in order to create custom report templates. You can easily assemble a simple report template by dragging and dropping fields into the template. The SQL database query is generated dynamically in the background, based on the configuration selections you make.

See "Working with Enterprise Objects and Fields" on page 231.

Another report template designer, the SQL Template Designer, is available for anyone with SQL knowledge.

The Dynamic Template Designer provides an inventory of **Enterprise Objects**, such as Job, Data Domain, Host, Storage Array, and Amazon Web Services (AWS) objects, each designed for a specific type of data. The enterprise object is an abstraction of the physical implementation of the relationships of the collected data in the database.

These enterprise objects provide the basis for creating and generating reports on collected data, to satisfy business use cases; for example, determining if your backup environment is providing sufficient data protection. Using the Dynamic Template Designer, you select an enterprise object as the starting point for your template. As you develop your template, you'll select fields and functions required to report on a specific enterprise object. You can assign a category which enables it to be grouped in the **Inventory**. The Dynamic Template Designer also presents

a variety of elements to help you design reports in the way that you'd like to view the data--as bar charts (in a variety of styles), donut charts, line charts, pie charts, or tables. In addition, you'll use this designer to define the elements required to select the scope of the data to be included in a report during report generation.

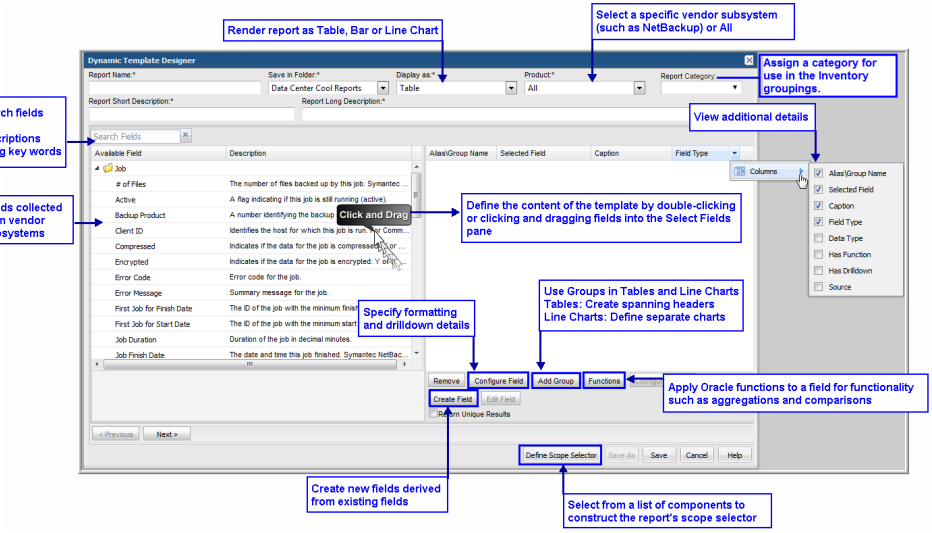
The final output from the Dynamic Template Designer is a dynamic report template. There are two types of report templates: Dynamic Report Template (SQL knowledge not required) and SQL Report Template (SQL knowledge required for database query creation). A report template is the basis for all reports. When you generate a report, you start by selecting a report template and then selecting the scope of the data to be included in the report's output. Often, the terms report template and report are used interchangeably, but it is important to note that the report template is where the designer elements reside.

To become acquainted with the components and initial steps required to create a report template:

- Review the Quick Start:
  See "Dynamic Template Designer Quick Start" on page 223.
- Read the Overview of Steps:
  See "Steps to Create Dynamic Templates" on page 230.
- Learn by Example: View advanced capabilities that have been incorporated into existing, out-of-the-box reports:
  See "Modify an Out-of-the-Box Dynamic Template" on page 222.
- Reference a Visual Overview:
  See "Dynamic Template Designer Cheat Sheet" on page 221.

Dynamic Template Designer Cheat Sheet

This diagram provides a visual overview of certain key features of the Dynamic Template Designer.

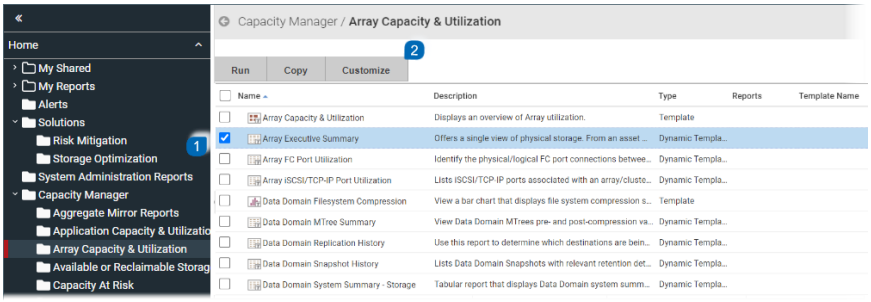


Modify an Out-of-the-Box Dynamic Template

Note: Tip: As a best practice, choose an existing dynamic template that is provided as an out-of-the-box report template and customize it.

To modify and save an out-of-the-box Dynamic Template

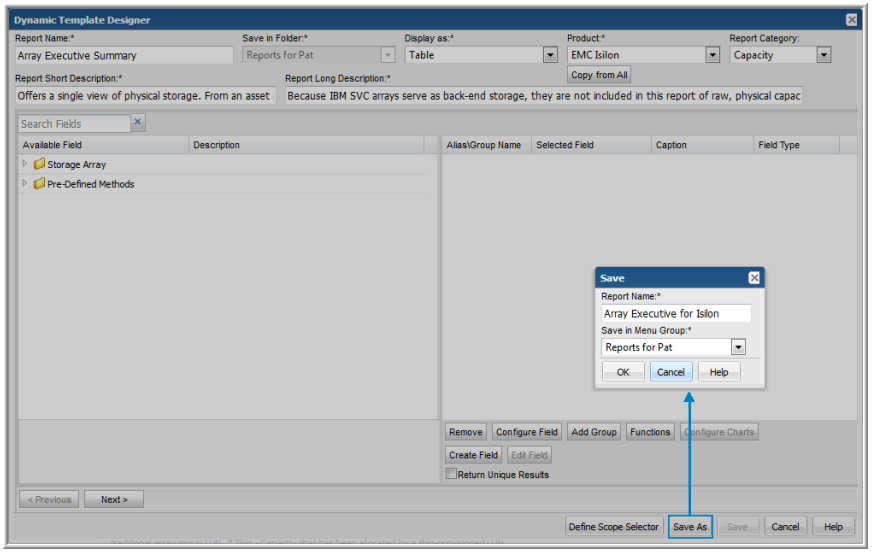
1. Search or navigate the report folders to find the report template to be modified and select it.
2. Click **Customize** to launch the Dynamic Template Designer.



3. Make the required changes in the Dynamic Template Designer.

4 Click **Save As**.

5 Enter a report name and select a user-defined menu group.



6 Click **OK** to save the revised template.

Dynamic Template Designer Quick Start

A basic use case simply can be to list values for various entities of an enterprise object. For example, list the allocated and available capacity of arrays. This is best represented in a tabular report.

The following quick-start example requires the following tasks to create a tabular report:

Quick Start Step 1: Create a Basic Table Dynamic Template

Prerequisite:

See "Select an Enterprise Object" on page 232.

1. Start by opening the folder in the left panel and drag fields into the Selected Field panel at the right.

   See "Configure General Dynamic Template Designer Components" on page 233.

   See "Dynamic Template Field Configuration" on page 239.

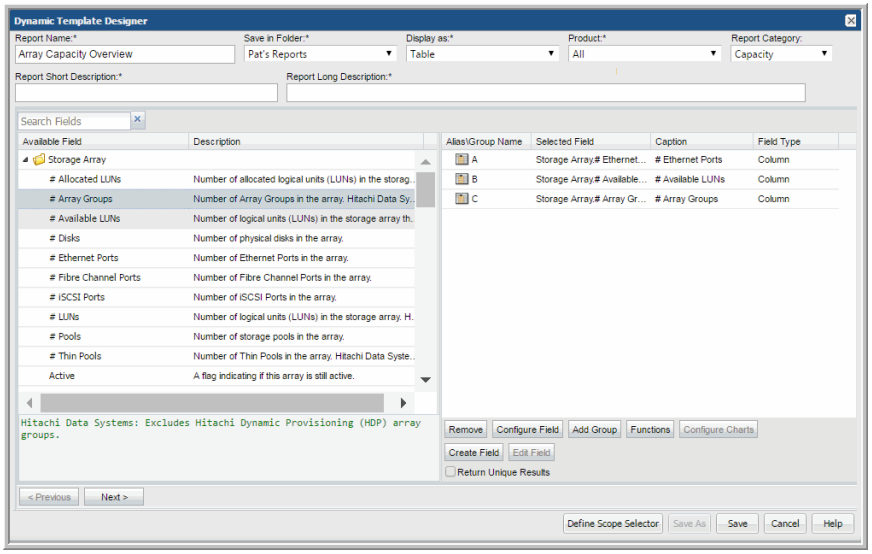
Note that when you drag fields into the right panel, an Alias name is assigned.

See “Alias Names for Fields” on page 239.

2. Enter or select the mandatory/optional fields shown at the top of the window and save the report template. **Category** refers to grouping assignment within the Inventory. When displayed in the Inventory, templates are sorted into information categories such as performance, storage or forecasting. If a category has not been selected, these templates are displayed under the heading Uncategorized. You can always assign a category, by customizing an existing template.

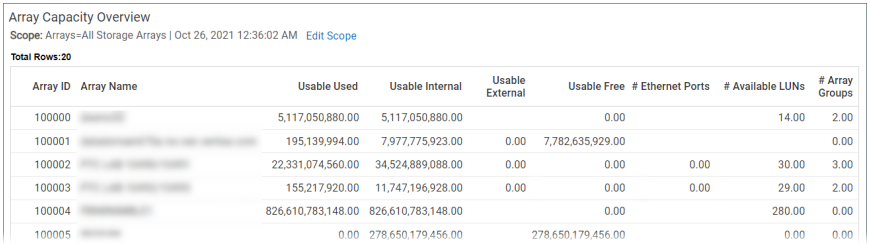
   - A dynamic template categorized for an object type and All products are displayed for all groups and individual servers of this object type.
   - A dynamic template categorized for an object type and specific products are only displayed for groups and individual servers of this object type, subsystem and not others.

   This step alone can produce a very basic report.



3. Click **Save** to save this initial version of the report template.

   This example produces a report similar to the following:



Array Capacity Overview

Scope: Arrays=All Storage Arrays | Oct 26, 2021 12:36:02 AM Edit Scope

Total Rows:20

Array ID	Array Name	Usable Used	Usable Internal	Usable External	Usable Free	# Ethernet Ports	# Available LUNs	# Array Groups
100000	[redacted]	5,117,050,880.00	5,117,050,880.00		0.00		14.00	2.00
100001	[redacted]	195,139,994.00	7,977,775,923.00	0.00	7,782,635,929.00			0.00
100002	[redacted]	22,331,074,560.00	34,524,889,088.00	0.00	0.00	0.00	30.00	3.00
100003	[redacted]	155,217,920.00	11,747,196,928.00	0.00	0.00	0.00	29.00	2.00
100004	[redacted]	826,610,783,148.00	826,610,783,148.00		0.00		280.00	0.00
100005	[redacted]	0.00	278,650,179,456.00		278,650,179,456.00		0.00	0.00

As you can see, these results would benefit from additional formatting, hyperlinks to other reports, totals at the bottom of the numeric columns, and perhaps even additional table columns with values calculated from other fields. For example, totals would be useful at the bottom of the table and perhaps an additional column with calculated values would improve this report.

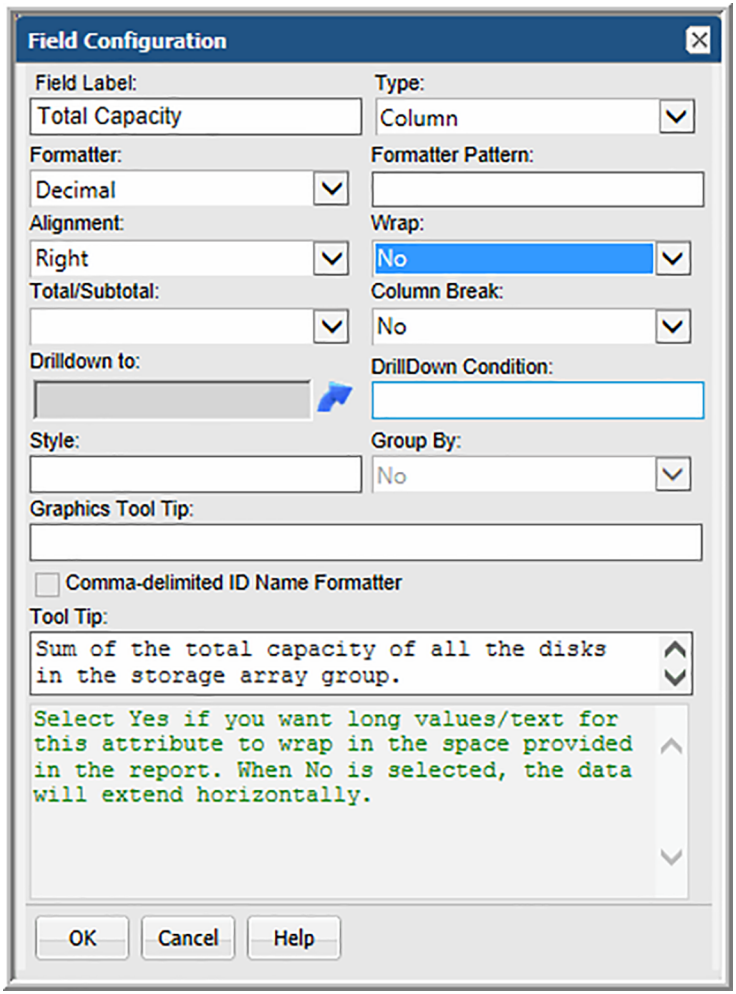
Basic formatting will be covered in this example.

Quick Start Step 2: Format Fields

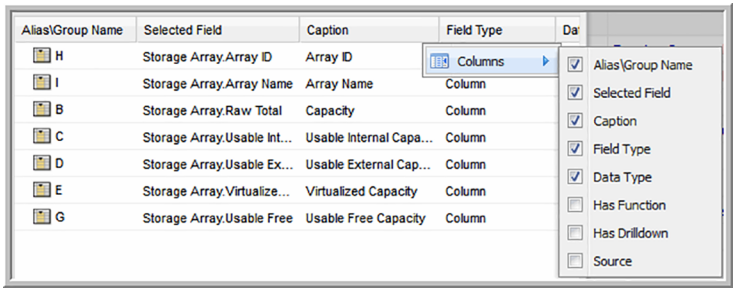
Double-click a numeric field. For this example, double-click the column that has the Capacity label.

1. Change the Caption from Capacity to **Array Capacity**.
2. For the Formatter, select **Decimal** and make sure the **Formatter Pattern** is blank.
3. Later, you may want to configure this field to get a total of the capacity of all the arrays in the report.

   See “Dynamic Template Field Configuration” on page 239.



Note: Tip: If you don't know if a field is numeric, hover over a header in the Selected Field panel to view the list of columns to select. For this example, select **Data Type**.

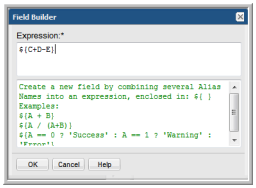


Quick Start Step 3: Create a Calculated Field

For this example, it would be useful to have a Total Usable column, derived from fields we have already included in the template. The calculation for total usable capacity will be defined as:

Usable Internal + Usable External - Virtualized Capacity = Total Usable Capacity

1. Click **Create Field** at the bottom of the Dynamic Template Designer window.



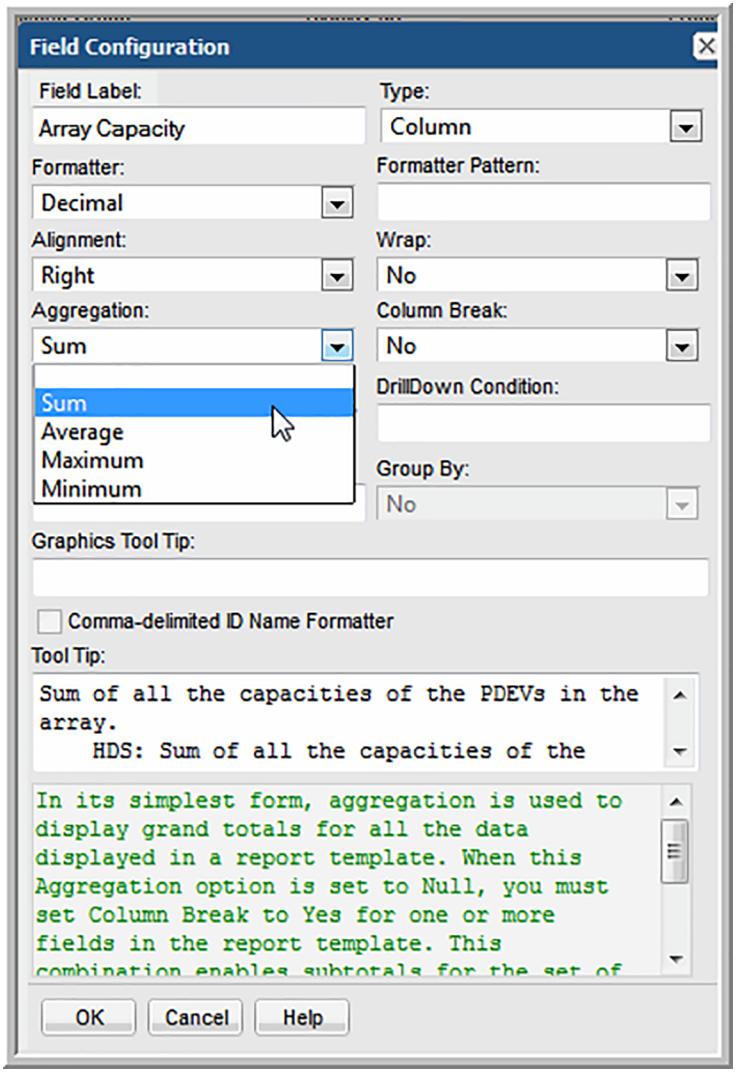
2. In the Expression box, enter the calculation for total usable capacity: ${C+D-E}

   Note that all expressions need to be enclosed in braces, with the following format: ${}

3. Click **OK** to launch the Field Configuration window where you can configure a Caption that will display **Total Usable** in the rendered report.

Quick Start Step 4: Configure Table Column Totals

1. Double-click a capacity field, such as Array Capacity.
2. In the Field Configuration window, select **Sum** from the **Total/Subtotal** drop-down list.



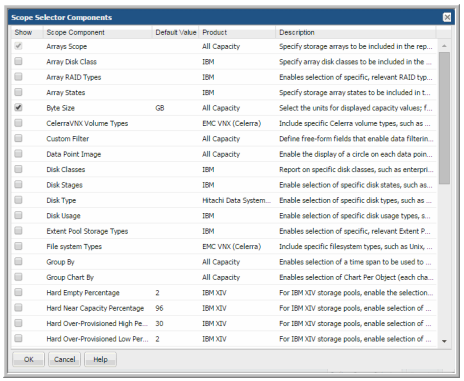
3. Click **OK**.
4. Repeat these steps for any numeric column where you want to see a total at the bottom of the tabular report.

Quick Start Step 5: Define the Scope Selector

This set of steps lets you select items that will appear in the Scope Selector when the report is run.

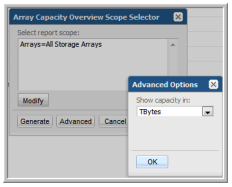
At the bottom of the Dynamic Template Designer window, click **Define Scope Selector**.

Scope selectors are specific to an enterprise object.



Note that some entries may be already selected and cannot be deselected. This is because those scope elements are required for the selected enterprise object.

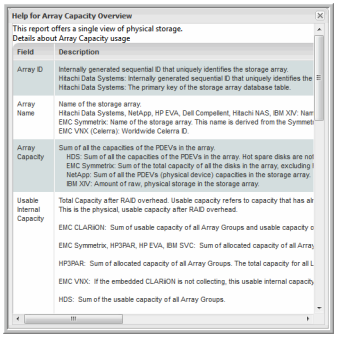
In this example, we've selected Byte Size. This lets the user select the units in the Advanced Options of the Scope Selector, as shown in the following screen.



Quick Start Step 6: Generate a Report and View the Automatically Generated Help

There is an underlying data dictionary that contains detailed descriptions of not only enterprise objects, but also the characteristics of those objects. When you create a report template, the help is automatically derived from the data dictionary.

1. Use **Search** at the top of the window and enter the name of the report template: Array Capacity Overview.
2. Select the report template from the search results and the report will render.
3. Click **Help** at the top right of the console window. The auto-generated help window appears.



Steps to Create Dynamic Templates

1. See “Select an Enterprise Object” on page 232.

   In the Dynamic Template Designer, all report templates are specific to an enterprise object: storage array, job, data domain, or host.

2. See “Configure General Dynamic Template Designer Components” on page 233.

   See “Working with Enterprise Objects and Fields” on page 231.

3. See “Add Fields and Methods to a Dynamic Template” on page 237.

   Drag and drop the fields and methods to be included in a report template.

4. See “Dynamic Template Field Configuration” on page 239.

   Customize fields by applying formatting and adding functionality, such as a drilldown to a detail report.

5. See “Dynamic Template Drilldown Configuration” on page 253.

   Access details and other reports using an html drilldown link.

   See “Examples of Dynamic Templates Containing Drilldowns” on page 260.

6. See “Using Groups in Dynamic Templates” on page 262.

   For examples of report templates that have double, spanning headers, see the Backup Executive Summary or the Chargeback Array Capacity report templates.

7. See “Dynamic Template Function Configurations” on page 264.

   A variety of functions can be applied: aggregation, comparison, numeric, character, date, character string to return numeric, null, and decode.

   See “Examples of Dynamic Templates Containing Functions” on page 272.

8 See “Create User-Defined Fields with the Field Builder” on page 275.

Create fields that are calculated from other fields, such as converting dates to numbers.

9 See “Configure a Bar Chart Dynamic Template” on page 295.

See “Examples of Bar Chart Dynamic Templates” on page 302.

10 See “Configure a Line Chart Dynamic Template” on page 312.

See “Examples of Line Chart Dynamic Templates” on page 323.

11 See “Configure an Area/Stacked Area Chart Dynamic Template” on page 303.

12 See “Configure a Pie Chart Dynamic Template” on page 324.

Create a chart that includes a sector field and a caption field.

See “Examples of Pie Chart Dynamic Templates” on page 326.

13 See “Configure a Static Filter” on page 277.

Specify a field, operation, and value.

14 See “Configure the Field Sort Order” on page 278.

Control the order that the fields will appear in a report template.

15 See “Dynamic Template Scope Selector Components” on page 279.

Define what the user can specify at run time.

16 See “Customize and Export Dynamic Templates” on page 292.

Modify existing Dynamic Templates and export them.

17 See “Save a Dynamic Template After Edits” on page 294.

Report templates that are shipped with the product or that are downloaded from the Cloud cannot be overwritten and will only allow save as to a different name.

Working with Enterprise Objects and Fields

Data collection can gather a significant amount of data from a variety of subsystems, such as storage arrays and backup systems. The Dynamic Template Designer can report on the following sets of data, represented by enterprise objects.

- Data Domain - Use this enterprise object to create templates that report on data collected from EMC Data Domain systems.
- Host - Hosts can be collected from any number of subsystems, such as storage arrays or backup systems. This host enterprise object enables you to identify the subsystem from which host data was collected, along with a number of other

characteristics, such as total filesystem capacity, that can help you make business decisions.

- Job - The Job enterprise object can be used to report on collected backup and restore jobs.
- Storage Array - Storage metrics can be collected from a number of storage systems, such as Hitachi, EMC, and IBM arrays. Use the Storage Array enterprise object to report on these systems.
- Other Enterprise Objects - As new data systems are introduced for data collection, typically via the SDK, enterprise objects will be listed in the Dynamic Template Designer. For example, Amazon Web Services objects may be listed, including AWS Billing Record, AWS EC2 Instance, and AWS S3 Bucket.

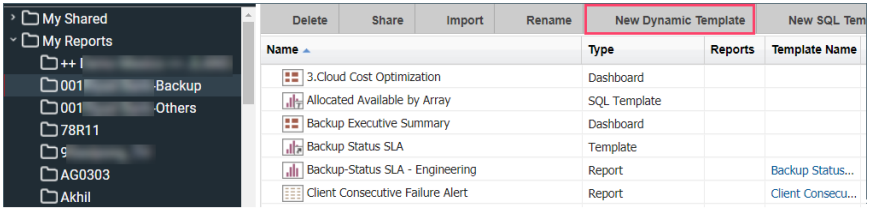
Collected data is available as fields--that is, object-specific details--to be included in report templates and reports. These fields represent the rich set of data that has been collected from a specific environment. For example, you could select the number of files for a backup job and the size of those files. Both “number of files” and “size” are characteristics or fields that will be displayed in a report for backup jobs.

Select an Enterprise Object

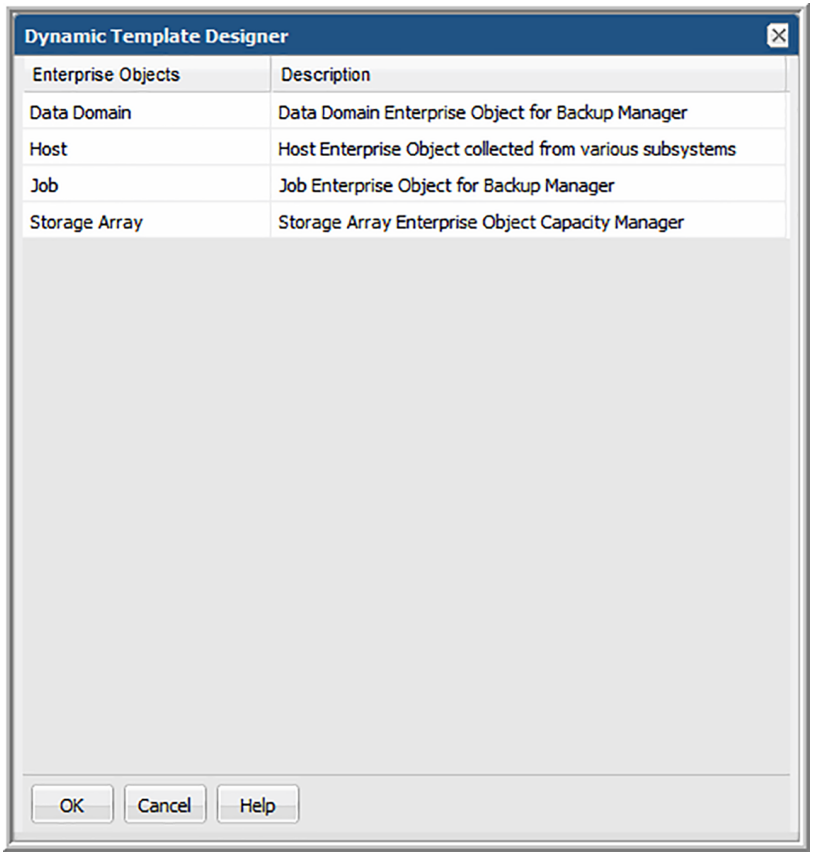
See “Working with Enterprise Objects and Fields” on page 231.

To select the enterprise object on which a report template will be based, take the following steps:

1. Select **Reports > My Reports> Your Custom Name Folder**.
2. Click the **New Dynamic Template** button.



3. In the Dynamic Template Designer dialog, select the object for which you want to design a report template and click **OK**. Simply double-clicking the enterprise object name achieves the same results--launching the designer window.



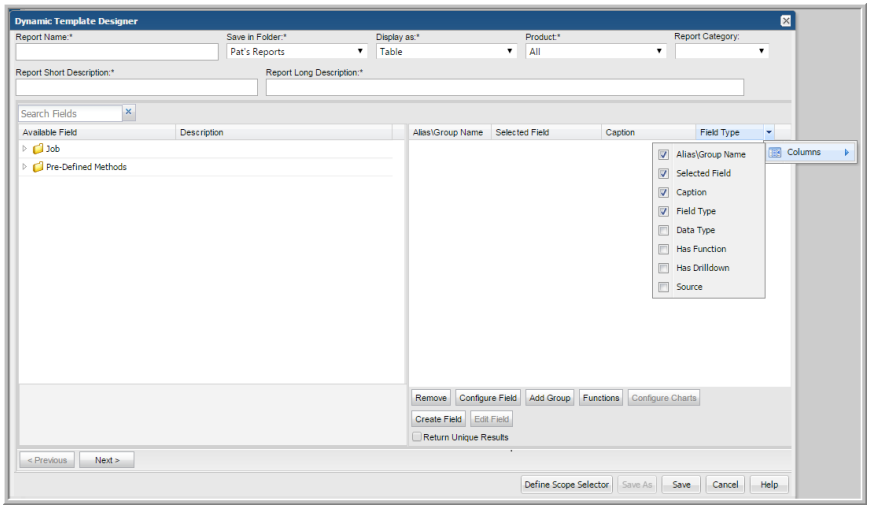
The Dynamic Template Designer window provides the elements required to design a report template.

Configure General Dynamic Template Designer Components

Prior to configuring the various elements of a report template, you must first select the enterprise object on which you want to report.

See “Select an Enterprise Object” on page 232.

The Dynamic Template Designer window enables the configuration of both the content and the display elements to be included in a report template.



The following components comprise the Dynamic Template Designer, with mandatory elements denoted with an asterisk (*):

- **Report Name*** - User-defined report name.
- **Save in Folder*** - Select a user-defined menu group in which to save the report template.
- **Display as*** - Select the format for the report's output:
  - **Area Chart** - Requires at least one field of the type caption and one or more fields of the type Area.
  - **Bar Chart** - Requires at least one field of the type caption and one or more fields of the type bar; or one field of the type caption, one legend field, and one bar field.
  - **Donut Chart** - Requires at least one field of the type caption and one of the type sector; or two or more fields of the type sector.
  - **Horizontal Bar Chart** - Requires at least one field of the type caption and one or more fields of the type bar; or one field of the type caption, one legend field, and one bar field.
  - **Horizontal Stacked Bar Chart** - Requires at least one field of the type caption and one or more fields of the type bar; or one field of the type caption, one legend field, and one bar field.
  - **Line Chart** - Requires at least one field of the type caption and one or more fields of the type line.

  - **Pie Chart** - Requires at least one field of the type caption and one of the type sector; or two or more fields of the type sector.
  - **Stacked Area Chart** - Requires at least one field of the type caption and one or more fields of the type Area.
  - **Stacked Bar Chart** - Requires at least one field of the type caption and one or more fields of the type bar; or one field of the type caption, one legend field, and one bar field.
  - **Table** - Requires at least one visible field.

- **Product*** - Select **All** for heterogeneous reports--that is, reports that display data collected from multiple vendor products or subsystems. Otherwise, from the drop-down list, select a specific vendor.
- **Category** - Select a category for the report template. These correspond with the groupings in the Inventory. Select from: **Administration**, **Backup**, **Billing**, **Capacity**, **Forecasting**, **Management**, **Overview**, or **Performance**. If a category is not selected, report templates are displayed in the **Inventory** under the heading **Uncategorized Report Templates**.
- **Copy from All** - This button appears only when a specific product is selected; for example, Veritas NetBackup. This option enables you to copy the configuration that has been set up for heterogeneous reports so that you can modify the configuration for your specific homogeneous report requirements. This starting point provides characteristics common to all vendor products, to which you can add fields specific to the product you selected.
  See "Converting to a Homogeneous, Product-Specific Template" on page 236.
- **Report Short Description*** - Provide a description that will appear in search results, enabling users to determine the intent of the report template and saved reports. This description is limited to 512 characters.
- **Report Long Description*** - Use this field to supply additional details such as vendor-specific names and acronyms so that this report will be included in search results. For example, Hitachi NAS reports have HNAS in the long descriptions to enable a Search on the acronym. This description is limited to 4,000 characters.
- **Search Fields** - Use this search box to find specific available fields to include in a report template. The search will return fields and pre-defined methods that have the string in either the name or description. For example, search for **error** to find both fields and pre-defined methods that can be included in a report template. This search is case-insensitive. Clear this box to reset the list to all available fields and pre-defined methods. If the object name matches the search string, all the related fields will be displayed.
- **Available/Selected Fields***-

See “Add Fields and Methods to a Dynamic Template” on page 237.

- **Column Selector** - Hover your mouse over any column header in the Selected Fields section at the right of the Dynamic Template Designer window. This reveals an arrow, which when selected, provides a list of columns that can be added or removed from this portion of the window. For example, you may want to add **Has Function** to the window so that you can easily identify which fields have had a function applied.

Alias/Group Name	**Alias**: The name of the alias for the selected field. An alias is assigned when a field is dragged into the Selected Fields panel. This enables ease-of-use when creating custom fields so that a simple alias can be referenced rather than a full field name. **Group Name**: A group name enables aggregated groupings for tabular reports.
Selected Field	Indicates the field that will be appear in the report. Note that the field is pre-pended with either its enterprise object or the prepackaged method that is applied to the field.
Caption	The caption originally is derived from the data dictionary, but it can be overridden when the field is configured.
Field Type	This indicates how the field will be used in the report template.
Data Type	Indicates the field’s data type: date, decimal, or string.
Has Function	Indicates if the field has a function applied to it.
Has Drilldown	Indicates if this field has been configured to be a link to another report.
Source	**Custom**: Indicates that the field was created from at least one enterprise object, such as a numeric calculation or text appended to a string value. **Enterprise**: Indicates that the field is derived from an enterprise object. **Method**: Indicates that the field was derived from a prepackaged method, listed in the left panel. These methods enable extended functionality for enterprise objects.

Next Step:

See “Add Fields and Methods to a Dynamic Template” on page 237.

Converting to a Homogeneous, Product-Specific Template

The Dynamic Template Designer can be used to create two different versions, starting with the same template:

- Heterogeneous: The scope selector and fields in the template are designed to select data across multiple vendor products for an enterprise object, such as EMC Symmetrix, Hitachi Data Systems, and IBM SVC arrays.
- Homogeneous: The scope selector and fields in the template are designed to select data for one specific vendor product, such as EMC Symmetrix arrays.

Copy from All

Once you have created a report template that can be used for all vendor products for the enterprise object (a heterogeneous template), you then can use the **Copy from All** feature to create a vendor-specific (homogeneous) version of this template, using the same fields that you've already selected and configured. Then, you can add vendor-specific fields to the new report template.

The Copy from All operation overwrites your homogeneous configuration, so if you modify this once and copy from all again, it will overwrite any modifications you've made. Also, you cannot go in the reverse direction--vendor-specific template to a heterogeneous, "all vendor products" template.

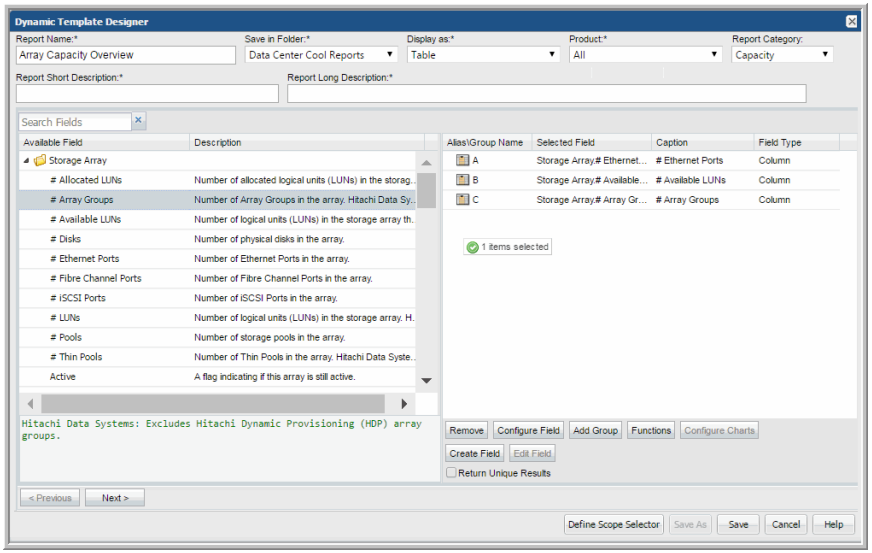
Note: There are cases for the Array enterprise object where Copy from All does not work. For example, a heterogeneous report template (a template representing all vendor products) may have a LUN object, but the vendor of the destination report template may not have that LUN object.

Add Fields and Methods to a Dynamic Template

When you design a report template, you'll take a number of steps to include fields and configure the way fields will be treated in a report's output. You control not only how the fields' values will be displayed, but you also can include pre-defined methods and totals/subtotals to make the data more meaningful.

1. Expand the folders at the left of the Dynamic Template Designer window.
2. Drag fields and pre-defined methods into the right window pane. You can also select and double-click a field to include it in the template.
    - Alternatively, **search** for a field based on a characteristic on which you want to report. For example, search on job to find backup job-related fields. A Search box at the top left of the pane enables searches of the available enterprise object fields and methods.
    - Double-clicking a field/method is an alternative way to add it to the template definition.

- Note that when you click on a field in the left pane, an additional description displays in green at the bottom of the pane.
- When fields and methods are dragged into the Selected Fields pane, they receive an Alias Name.



3. Pre-defined methods can be dragged into a report template to be treated as a field for which you can apply additional formatting.
4. In the right pane of the Dynamic Template Designer window, double-click a field to configure custom settings.
5. Once the fields have been configured, click **Next** at the bottom of the Dynamic Template Designer window to configure sort order and filtering.
6. To define additional filtering--that is, a static filter that cannot be changed at report run time--click Add at the bottom of the Filter pane.

   See “Configure a Static Filter” on page 277.
7. To define the sort order:

   See “Configure the Field Sort Order” on page 278.
8. To ensure that duplicate sets of data do not appear in a report:

   See “Return Unique Results” on page 279.
9. Click **Define Scope Selector** to select the elements that will be presented to the user when generating a report from the report template.

   See “Dynamic Template Scope Selector Components” on page 279.

Next Step:

See “Dynamic Template Field Configuration” on page 239.

Alias Names for Fields

To enable easier manipulation of fields in the Dynamic Template Designer, a simpler name is assigned when a field is dragged into the Selected Fields panel in the Report Designer. This name, typically a single alpha character, can be used to perform operations, such as mathematical calculations, with other fields.

Note that as you add and remove fields from the selected list, the alphabetic sequence for the alias names is not retained. That is, when a field is removed, its alias name will not be reused when you add fields to the selected list.

Conditions When Hidden Fields are Included by Default

Under certain conditions, when a field is dragged into a report template and then configured for a specific action, another hidden field will also appear in the list. This hidden field is mandatory and cannot be removed, as the system knows that it is required for the operation that you are configuring. These hidden fields are included to mitigate performance issues.

Example: If a report template is based on the summary status description field and you want to configure a drilldown in this report, the report will perform much better if the summary status was added as a hidden field. Automatic addition of a hidden field is performed in situations when the report meaning will not change.

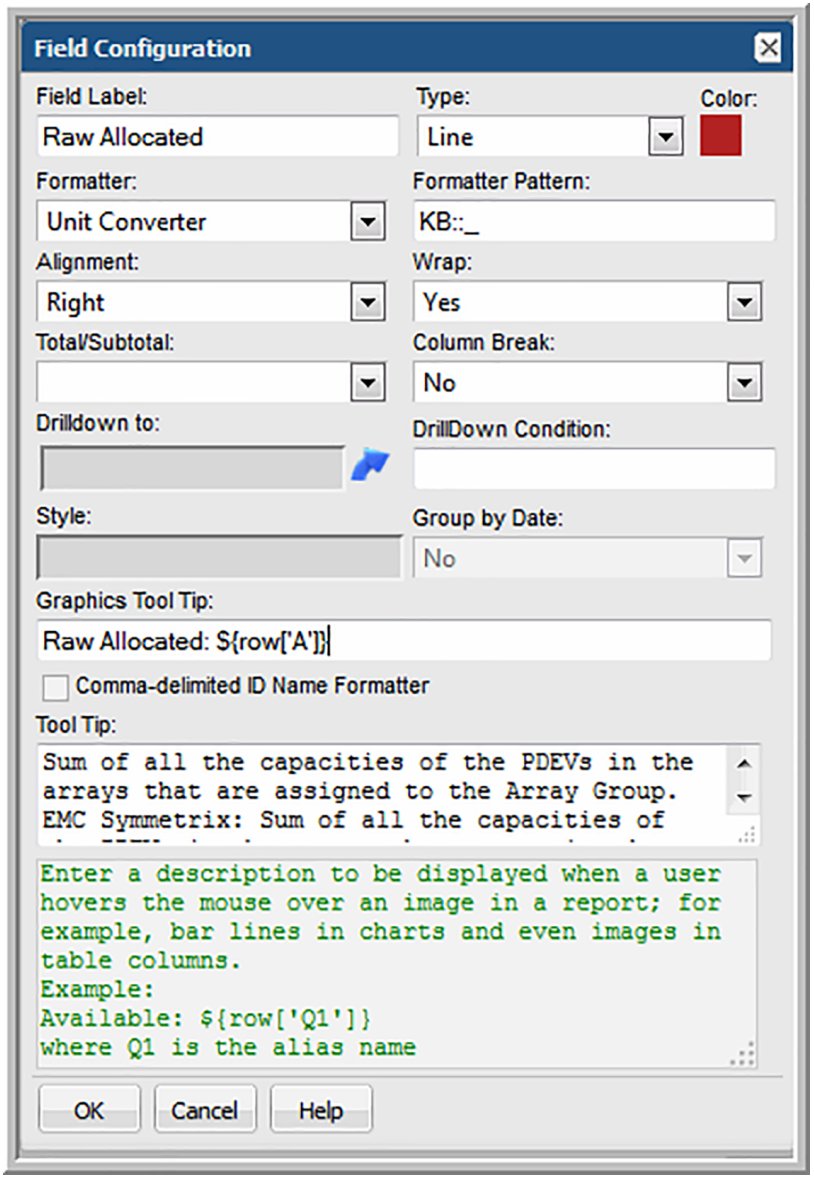
Dynamic Template Field Configuration

Prerequisite:

Add fields and methods to a dynamic template. Once fields have been added to a report template, a number of special treatments can be configured to customize the way the data will be represented in a report.

1. Double-click a selected field to access the Field Configuration window.

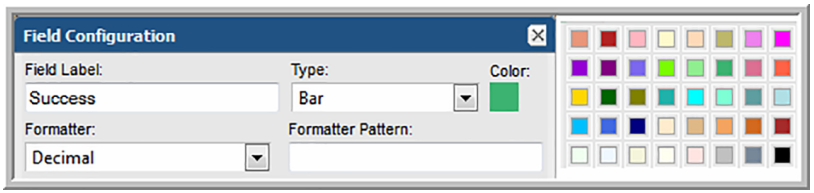
   This Field Configuration window presents lists of values that can be selected to configure formatting. Only values relevant for the selected field will be listed and some options may be grayed out if they are not relevant to the selected field.


Field Configuration
Field Label:
Raw Allocated
Type:
Line
Color:
Formatter:
Unit Converter
Formatter Pattern:
KB::_
Alignment:
Right
Wrap:
Yes
Total/Subtotal:
Column Break:
No
Drilldown to:
DrillDown Condition:
Style:
Group by Date:
No
Graphics Tool Tip:
Raw Allocated: S{row['A']}
Comma-delimited ID Name Formatter
Tool Tip:
Sum of all the capacities of the PDEVs in the arrays that are assigned to the Array Group. EMC Symmetrix: Sum of all the capacities of
Enter a description to be displayed when a user hovers the mouse over an image in a report; for example, bar lines in charts and even images in table columns.
Example:
Available: ${row['Q1']}
where Q1 is the alias name
OK
Cancel
Help

Label	Description
Field Label	Field Labels in Tables This is the text that will be displayed as a field label or a table column header in a tabular report. To replace the default caption with one customized for your reports, simply overwrite the default text. Field Labels in Charts The caption field becomes the x-axis label. For Line Charts, the label can be any object, such as a primary server. When charting performance metrics for an object, a date field should be configured as a label. See “Line Charts for Performance Metrics” on page 313.
Type: Display as Table	When the report template has the “Display as” component set to Table, the following Type options are available. ■ **Hidden** - Use this option when configuring a drilldown to another report template. The field will not be displayed in the report, but is required to access the necessary details to link to the sub-report template. There will be cases where the UI will automatically add hidden fields to the report template, to ensure satisfactory performance. ■ **Column** - Lists a column of values in a tabular report. ■ **Checkbox** - Include/exclude selectors to be displayed in a column in a tabular report. ■ **Exclude** - There will be cases when a field will be needed only for filtering purposes and you do not want this field to be included in the selected fields for the report.

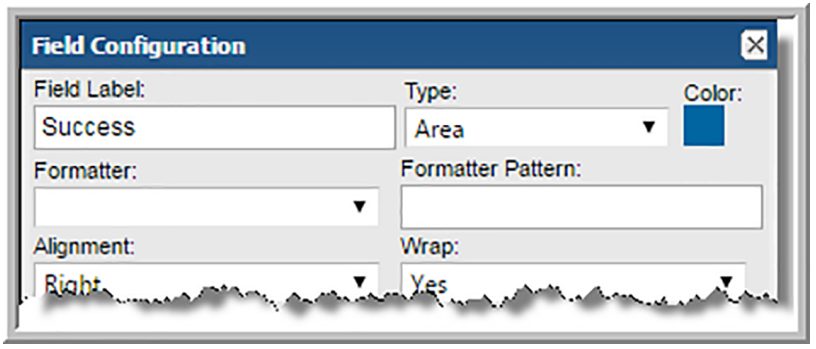
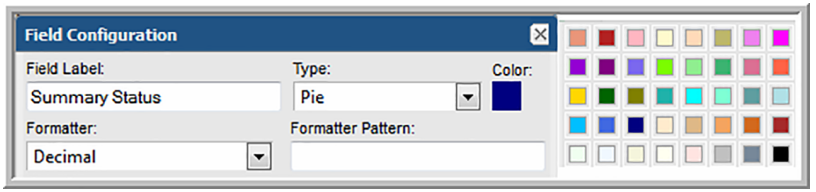
Label	Description
Type: Display as Bar Chart	When the report template has the “Display as” component set to Bar Chart, the following Type options are available.

- **Caption** - Define this field to be the x-axis values.
- **Bar** - Configure the field to be represented as a bar and set its color.



- **Line** - Configure the field to be represented as a line and select the shape to denote its value in the line: rectangle, circle, diamond, or none.
- **% Line** - Configure the field to be represented as a % line and select the shape to denote its value in the line: rectangle, circle, diamond, or none.
- **Legend** - Define the labels and colors to be displayed in the bar chart’s legend.
- **Hidden** - Use this option when configuring a drilldown to another report template. The field will not be displayed in the report, but is required to access the necessary details to link to the sub-report. There will be cases where the UI will automatically add hidden fields to the report template, to ensure satisfactory performance.
- **% Bar** - Define the field to be a % bar in the chart.
- **Exclude** - There will be cases when a field will be needed only for filtering purposes and you do not want this field to be included in the selected fields for the template.

Label	Description
Type: Display as Line Chart	When the report template has the “Display as” component set to Line Chart, the following Type options are available. ■ **Caption Field**- This required field type supplies the category caption for the x-axis. Typically, a date field becomes the caption because line chart data is best represented over a time line. When you configure a date field as a caption, the field configuration automatically sets **Group By Date** to **Yes**. ■ **Line Field** - A numeric field can be configured as a line. One or more line fields can be included in a line chart report template. Select **Line** for the field type. In a “multiple objects per chart” scenario, a line represents a specific object, for example, an array. Line colors do not apply in this case. Colors will be determined dynamically by the system, since the number of objects represented by the lines is unknown until run time. In a “one object per chart” scenario, a line represents a metric for that object; for example, both used and available capacity could be lines in each array chart. ■ **Group by Object** ■ **Hidden** ■ **Exclude** - There will be cases when a field will be needed only for filtering purposes and you do not want this field to be included in the selected fields for the report template. For example, to filter on failed and partial backup jobs, you would drag in the job status field, make it an Exclude field type, and **Note:** Use the **Add Group** option to create folders when you need to group data fields that have disparate units of measure. This Group functionality is used in tabular reports to create spanning headers, but for line charts, it renders separate charts. For example, you might have a capacity value in KiB, but a performance value in Kbps. This requires two charts with different metric scales to plot the values.

Label	Description
Type:Display as Area	When the report template has the “Display as” component set to Area/Stacked Area Chart, the following Type options are available. ■ **Caption** - Define this field to be the x-axis values. ■ **Area** - Configure the field to be represented as an area and set its color.  ■ **Legend** - Define the labels and colors to be displayed in the area/stacked area chart’s legend. ■ **Hidden** - Use this option when configuring a drilldown to another report template. The field will not be displayed in the report, but is required to access the necessary details to link to the sub-report. There will be cases where the UI will automatically add hidden fields to the report template, to ensure satisfactory performance. ■ **Exclude** - There will be cases when a field will be needed only for filtering purposes and you do not want this field to be included in the selected fields for the template.
Type: Display as Pie/Donut Chart	When the report template has the “Display as” component set to Pie Chart or Donut Chart, the following Type options are available. ■ **Caption** - Enter the caption associated with the pie/donut sector. ■ **Sector** - Configure the field to be a pie or donut sector and set its color.  ■ **Hidden** - Use this option when configuring a drilldown to another report template. The field will not be displayed in the report, but is required to access the necessary details to link to the sub-report template. There will be cases where the UI will automatically add hidden fields to the report template, to ensure satisfactory performance. ■ **Exclude** - There may be cases when a field is needed only for filtering purposes and you do not want this field to be included in the selected fields for the report template.

Label	Description
Formatter	The formatter list includes only values relevant for the field. For example, numeric fields have only numeric formatters listed, such as Number, Currency, and File Size. Valid formatters include: Custom Number, Numeric String, Date, Date + TimeZone, Group by TimePeriod Start/End Formatter, Date Group By, Numeric, Decimal, Speed, Currency, File Size, Time Duration, Char Substitute, Unit Converter, File Size No Label, Null Substitute, Time Elapsed, Cron Schedule, Yes/No, Percentage, Status Icon, Full Name, Area Sparkline, Column Sparkline, Line Sparkline.
Formatter Pattern	Enter a specific pattern for the formatter to customize the default format. Example: KiB::GB or KB::_
Alignment	The default value provides the most common usage. For example, numeric values typically should be right-aligned so that decimal points line up.
Wrap	Select **Yes** if you want long values/text to wrap in the space provided in the report. When **No** is selected, the data will extend horizontally.
Total/Subtotal	Total/Subtotal Options: Null, Sum, Average, Minimum, Maximum. Use this configuration to aggregate the values, such as all the backup job sizes for a backup server. In its simplest form, this aggregation is used to display grand totals for all the data displayed in a tabular report. To enable subtotals for a set of fields: ■ Configure **Total/Subtotal** with a Null value and set **Column Break**to **Yes** for one or more fields in the report template. ■ Once you configure a subtotal, you can no longer create grand total rows at the end of the tabular report template. * Never configure **Total/Subtotal** and **Column Break = Yes** for the same field.

Label	Description
Column Break	**Column Break for Tabular Report Templates** Select **Yes** if you want to list subtotals at the end of a group of related fields. Note that when you use subtotals, the sort order is automatically determined. All fields with a Column Break selection are sorted in the relevant ascending order and this sort order cannot be changed. Be sure not to remove or modify this sort order. Example: You could create a tabular report for client, server, and kilobytes. Then, you could subtotal on server to aggregate the kilobytes per server. * Never configure **Total/Subtotal** and **Column Break = Yes** for the same field. * Once you configure a subtotal, you can no longer create grand total rows at the end of the tabular report template. --- Column Break for Line Chart Report Templates For line charts, the column break is required if the selected object is potentially not unique. To ensure uniqueness, you may need to include a unique ID field with the Column Break set to Yes. For example, an array name may not be unique, however, the array ID is sure to be unique. Include both the name and ID in the template, but make the ID a hidden field with the Column Break set to Yes.
Drill Down to	Select the report to which this field will link. This takes you to the Drilldown Configuration window. For a list of out-of-the-box report templates that have drilldowns that you can use as examples, See “Examples of Dynamic Templates Containing Drilldowns” on page 260.

Drilldown Condition	A drilldown condition lets you control cases where a drilldown doesn’t make sense to be displayed, for example, null or 0 values for a field. For example, you wouldn’t want to drill down on a Job ID for a row if the value of Job ID returned from the database is Null or 0 (zero).
	
	Example: ${row['B'] != '0'}
	
	where B is the alias name
	
	In this example, the drilldown link will appear only when the returned value is not zero.
	
	Example: Configure a drilldown on field D and enter a drilldown condition as shown in the following example. In this example, if the row where the value of the field aliased as A is NULL, the value in column D will not have a drilldown URL enabled.
	
	`${row[‘A’] != ''}`
	
	Example: Multiple conditions can be configured.
	
	`${row[‘C'] != '' && row['E'] == ''}`

Label	Description
Style	Use this option to configure the background color of a cell in a tabular report. Example: `${row['C'].data == 'Partial' ? 'yellowBackground' : row['C'].data == 'Failed' ? 'redBackground' : '' }` where C is the alias of the selected field.
Group by Date	This is applicable only for **Date** fields in **Bar, Line,** and **Pie Chart** report templates to aggregate data for charting. If only one date field is specified, it is used as the default for the Group By Date configuration. If you configure a date field as a caption, Group by Date is automatically set to Yes. If there are multiple date fields in a template, you must choose only one for this setting. This Group by Date setting identifies the field that the scope selector will use with the Group By scope component. When this Group by Date element is set to Yes, the Group By component in the Scope Selector is automatically selected.
Time Period Date	Use the Time Period Date selection in report templates to designate which date field will be used for the Time Period scope. If only one date field is included in the template, that field is used as the default. If more than one date field is defined in a template, only one date field can have this Time Period Date set to Yes. The Time Period Date must be set to Yes for one of the date fields in the template when a Time Period, but no Group By is selected in the scope selector. This configuration enables grouping of raw, non-aggregated data. Time Period Date can be set for displayed, hidden, or excluded date fields. Currently, this option is available only for storage arrays and objects created via the SDK.
Graphics Tool Tip	Enter a description to be displayed when a user hovers the mouse over an image in a report; for example, use this setting to enable mouse-overs on lines in bar charts, lines in line charts, and even images in table columns. Example: Available: ${row['Q1']} where Q1 is the alias name.
Comma-delimited ID Name Formatter	For text fields only. Use this option if data returned for a field has multiple values that need to be separated by commas. Only the names will be shown in the rendered report and the IDs will be used for drilldowns, if drilldowns are configured.
Tool Tip	This is the description supplied by the Data Dictionary for this field. You can overwrite or supplement this text. This description will be displayed when a user hovers the mouse over the field or table column label in a report. It will also be incorporated into the context-sensitive help that is displayed from the Help button.

Next Step: Configure dynamic template drilldowns.

See “Dynamic Template Drilldown Configuration” on page 253.

Valid Formatter Patterns

In the Field Configuration window, the Formatter selector requires additional details provided in the Formatter Pattern box. Note that only Formatters relevant to the selected report template field will be made available for selection.

Use the following table to determine valid formatting patterns.

Formatter Name	Data Type	Valid Patterns	Description
Date	Date	N/A	Date formats are derived from the user's profile.
Date+TimeZone	Date	N/A	Date formats are derived from the user's profile.
Group by TimePeriod Start/End Formatter	Date	[S|E]::[D|H|M|S|0|Y|T] S is used to calculate start date when End date is provided based on "group by". E is used to calculate End date when Start date is provided based on "group by". D-> Day H-> Hour M-> Minute S-> Second O- Month Y- Year T- Epoch Time in Milliseconds	Allows the framework to calculate the finish time based on the start time and “group by” selected.
Date Group By	Date	N/A	Used in charts and graphs to format the legend, based on the “group by” that is selected in the scope selector at run time. This formatter determines the optimum date format for that group by.

Formatter Name	Data Type	Valid Patterns	Description
Time Elapsed	Number	(SEC)|(MIN)|(HR)|(DAY)	Displays time in a more readable format, converting minutes to hours, where relevant. For example, 120 min would be displayed as 2 hrs.
Time Duration	Number	N/A	Formats a time value that is in milliseconds to the following format: hh:mm:ss Applicable only for fields that denote time duration; for example: finish_date minus start_date
Cron Schedule	Number Varchar2	N/A	Presents a standard cron expression in a user-readable form. For example, */30**** will display as “Every 30 minutes.”
Numeric	Number	N/A	Format derived from the user’s profile.
Numeric String	String	N/A	Treats numbers as a string without formatting. This allows sorting to be done numerically.
Custom Number	Number		Enables users to define their own number format which is not tied to the setting in their user profile.
Decimal	Number	% converts to % format Otherwise, leave as null.	Formats the value with the number of decimal places specified in the user’s profile; can also be used for percentages.
Speed	Number	N/A	Formats speed values. N/A if speed = 0.
Currency	Number	N/A	Format derived from the user’s profile.

Formatter Name	Data Type	Valid Patterns	Description
File Size	Number	(KiB)|(MiB)|(GiB) is the input data type	Formats the file size unit of measure to what is most applicable to the selected field value. For example, 2048 MB would become 2 GB
File Size No Label	Number	Units::Units	For Backup Manager only. Formats the file size to the most relevant unit of measure, but omits the unit label. Applicable for byte size data.
Unit Converter	Number	**KB::TB** **KB::_** (for dynamic output)	Determines the value to be used when converting from one unit to another, for example, from KB to GB. The “division by” value is determined from the user’s profile settings, either 1000 or 1024. Note that for line charts, if you select Unit Converter for a line, it only formats the mouse-over, not the Y-axis label.
Char Substitute	Number Varchar2 Date	`<value_to_replace>::<replacement_value>`	Replaces specified characters in a string value. For example, if the database returns "Host1, Host2" this formatter would be replaced the "," with " "
Null Substitute	Number Varchar2 Date	<string>	Replaces null values with a user-supplied string.

Formatter Name	Data Type	Valid Patterns	Description
Percentage	Number	number1:number2:number3 Data values from 0 to number1 will render Green color in the percentage bar. Values between number1 and number2 will render Yellow color and values between number2 and 100 will render Red color. The values number1 and number2 should always be between 0 and 100. The number3 parameter determines the size of the % bar in pixels.	Display a % bar (thermometer) in a chart.
Full Name	Number	N/A	For File Analytics only. Converts a shortcut name to a full name for NetApp, CIFS, and WinFS.
Yes/No	Varchar2	N/A	Translates Y/N to Yes/No.
Status Icon	Varchar2 Number	Example: 0|Green,2|Red,1|Yellow,*|White If the value = 0, the circle icon will render in green.	Displays status icons depending on the data, which could render in red, yellow, blue, green and white.
Truncate	Varchar	Number	Truncates a string with ellipsis. The maximum size of the string is 28 characters. Use the pattern to override the maximum size. When using the Truncate formatter, it is recommended to complete the Tool Tip field.

Examples of Dynamic Templates Containing Graphics Tool Tips

The best way to learn how to configure graphics tool tips in a report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A

Report Template	Report Type	Product: Alias	Caption	Field Type
Reports				
Job Status Summary	Bar	All: T	Failure	Bar
		All: R	Success	Bar
		All: S	Warning	Bar
Job Summary	Table	EMC Avamar: M	Error Code	Column
		All: K	Exit Code	Column
		Veritas NetBackup: P	Exit Code	Column
Job Volume Summary	Bar	All: D	# of Files Backed Up or Restored	Line
		All: B	# of Jobs	Hidden
		All: C	Backup/Restore Volume	Bar
Job Duration	Bar	All: F	Job Duration	Bar
Reports				
Array Executive Summary	Table	All: J	Used	Column
NetApp Aggregate Summary	Table	NetApp: R	Usage	Column

Report Template	Report Type	Product: Alias	Caption	Field Type
NetApp Volume Summary	Table	NetApp: K	Usage	Column
NetApp Cluster-Mode Summary	Table	NetApp Cluster-Mode: H1	Usage	Column
Hitachi DP Pool Summary	Table	Hitachi Data Systems: H	Used	Column
EMC VNX (Celerra) Storage Pools Summary	Table	EMC VNX (Celerra): C1	Used	Column
EMC VNX (Celerra) Volume Summary	Table	EMC VNX (Celerra): B1	Usage	Column
Hitachi NAS Storage Pool Summary	Table	Hitachi NAS: P	Usage	Column
EMC Isilon CPU Performance by Cluster	Line	EMC Isilon: B	Max CPU %	Line
EMC Isilon CPU Performance by Node	Line	EMC Isilon: G	Avg CPU %	Line
NetApp Cluster-Mode Disk Performance by Cluster	Line	NetApp Cluster-Mode: C	Disk Busy %	Line

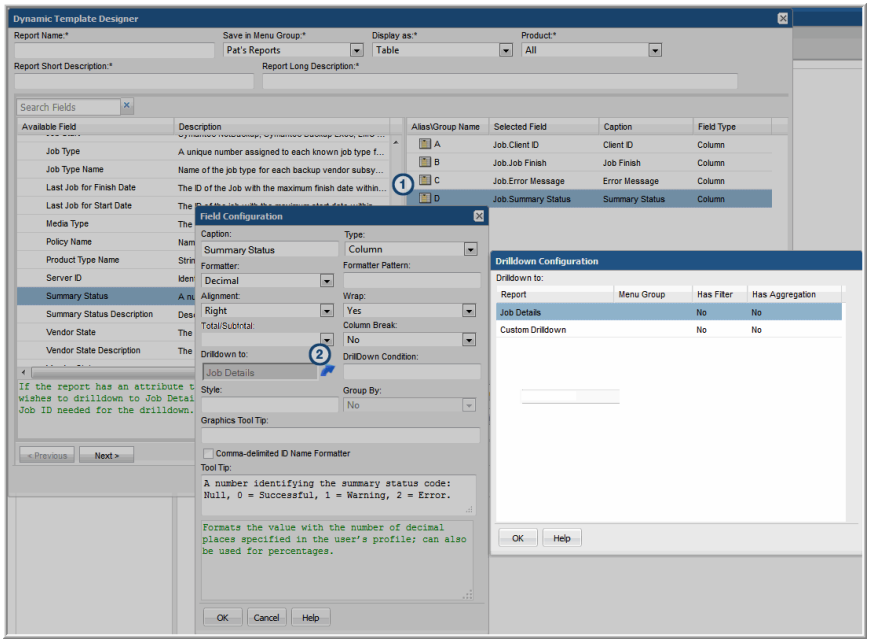
Dynamic Template Drilldown Configuration

For a list of out-of-the-box report templates that have drilldowns that you can use as examples,

See “Examples of Dynamic Templates Containing Drilldowns” on page 260.

Note: An ID field cannot have formatting if you want to use this field for a drilldown. Ideally, this field should be a hidden field. Or, you could include the ID twice, one would be hidden without formatting while the other could be displayed with formatting. The hidden field would then be used as the drilldown field. Note that some ID fields are added by the system and will therefore not have any formatting. For example, when a host name is dragged into a template, a related ID field is automatically added.

1. In the Dynamic Template Designer window, double-click a Selected Field to view the Field Configuration window.
2. In the Field Configuration window, click the **Drilldown to** icon to view the Drilldown Configuration window.



The Drilldown Configuration window lists reports that can be selected as a target for the drilldown. This list includes the following characteristics for each of the drilldown report templates:

- **Menu Group**: Displays the location of the drilldown report template.
- **Has Filter**: Indicates if the drilldown report template has filters applied. This is relevant because the data displayed in the drilldown report may differ from its parent report, due to the extra filtering. Therefore, the set of data may be different between the two reports.

- **Has Aggregation**: Indicates if the drilldown report has Oracle aggregation functions applied (See "Aggregation Functions" on page 267.). This may not be what you want, if the target of the drilldown has an aggregation that is different from its parent.

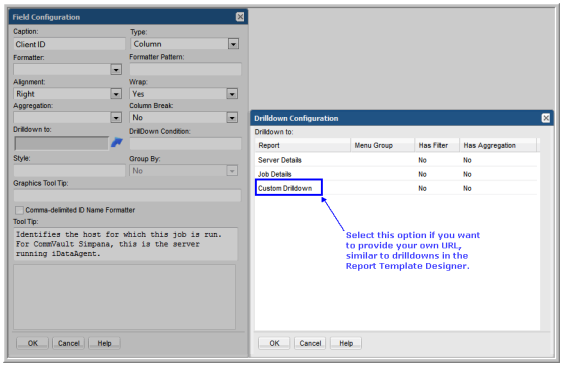
3. Select a report from the list that will be the target of the drilldown.

   To remove or de-select a drilldown report, press the **Ctrl** key and click the previously selected report.

4. Instead of selecting a report template from the drilldown reports list, you may want to specify a custom drilldown URL. In the Drilldown Configuration window, select **Custom Drilldown** and click **OK**.

5. In the bottom of the expanded Drilldown Configuration window, enter a URL expression.

   Note that there currently is no validation on the syntax of this expression in the Dynamic Template Designer.



Custom Drilldowns and Examples

The following examples illustrate how to specify a custom drilldown.

Example of a Custom Drilldown to the Host Details Report by Job ID

Be sure that the field used for the drilldown (in this example, S) does not have any formatting applied.

```
systemName=backupDetails&jobId=${row['S']}
```

Where:

systemName	systemName - The required prefix for drilldowns to out-of-the-box report templates. To identify the systemName, templateName, or templateInstanceID of an existing report template, generate that report and in the active browser window type: Ctrl-Alt-T templateName - The required prefix for drilldowns to custom report templates designed in the SQL Template Designer. templateInstanceID - The required prefix for drilldowns to report templates created in this Dynamic Template Designer.
<displayReportName>	Insert the specific report system name or template name or template instance ID, along with the variable that is to be supplied by the parent report. Note that a string prefaced with an ampersand (&) designates a parameter that will be passed to the report. To identify the systemName, templateName, or templateInstanceID of an existing report template, generate that report and in the active browser window type: Ctrl-Alt-T
=${row['<field_name>']}.data	The syntax required for the field name specification; substitute a value for <field_name>. In this example, S is the alias of the Job ID field.

Example of Custom Dual Drilldowns

To identify the systemName, templateName, or templateInstanceID of an existing report template, generate that report and in the active browser window type: Ctrl-Alt-T

This more advanced example illustrates how multiple drilldowns can be specified, with conditions.

```
${row['F'] == 'fb' ? 'templateInstanceId=25200' :

'systemName=listCKDVolumeUtilSummary'}&extendedPoolIds=$
{row['D']}&reportScope=&lunStatus=C
```

In this example, if the value of F is fb, then drill down to the report with report ID 25200 (the LUN Utilization Summary); otherwise, drill down to listCKDVolumeUtilSummary (the IBM CKD Volume Summary).

Example of a Custom Drilldown to the Array Detail Report

To identify the systemName, templateName, or templateInstanceID of an existing report template, generate that report and in the active browser window type: Ctrl-Alt-T

This example shows a drilldown to the Array Detail report.

systemName=arrayDetail&ignoreParent=true&arrayId=${drillDownParam}&arrayIds=${drillDownParam}&tryIndex=3&historyDays=3&dateRange.startDateTime=1359763535166&dateRange.finishDateTime=1390867535166&groupById=13

Where:

tryIndex	Used to control the top N rows displayed in aggregate and disk volume utilization: tryIndex=3&
historyDays	Used to control top N rows displayed in the volume utilization: historyDays=3&
dateRange	Times are seconds from epoch: dateRange.startDateTime=1359763535166&dateRange.finishDateTime=1390867535166

groupById

groupById = 13 means group by month

GROUP_BY_HOUR: 10

GROUP_BY_DAY: 11

GROUP_BY_WEEK: 12

GROUP_BY_MONTH: 13

GROUP_BY_QUARTER: 14

GROUP_BY_YEAR: 15

Note that if you are specifying this drilldown configuration in the SQL Template Designer, the syntax is:

```
systemName=arrayDetail&ignoreParent=true&arrayId=$
{row['storage_array_id'].data}&arrayIds=${row['
storage_array_id'].data}&tryIndex=3&historyDays=3
&dateRange.startDateTime=1359763535166&
dateRange.finishDateTime=1390867535166&groupById=13
```

Example of a Custom Drilldown to a Dynamic Template - Job Summary

To identify the systemName, templateName, or templateInstanceID of an existing report template, generate that report and in the active browser window type: Ctrl-Alt-T

This example shows the configuration of a drilldown to the Job Summary report.

```
templateInstanceId=150&supportsPagenation=true&jobStatusIds=1&dateRange.startDate=$
{row['start_date_char']}&dateRange.startHour=${row['start_hour_char']}&
dateRange.startMinute=0&dateRange.startSecond=0&dateRange.finishDate=${row['
finish_date_char']}&dateRange.finishHour=${row['finish_hour_char']}&
dateRange.finishMinute=59&dateRange.finishSecond=59&dateRange.useFinishTime=false&
backupServerIds=${row['server_id']}&parentJobOnly=false&ignoreRetries=true
```

The following parameters are supported for Job enterprise object drilldowns:

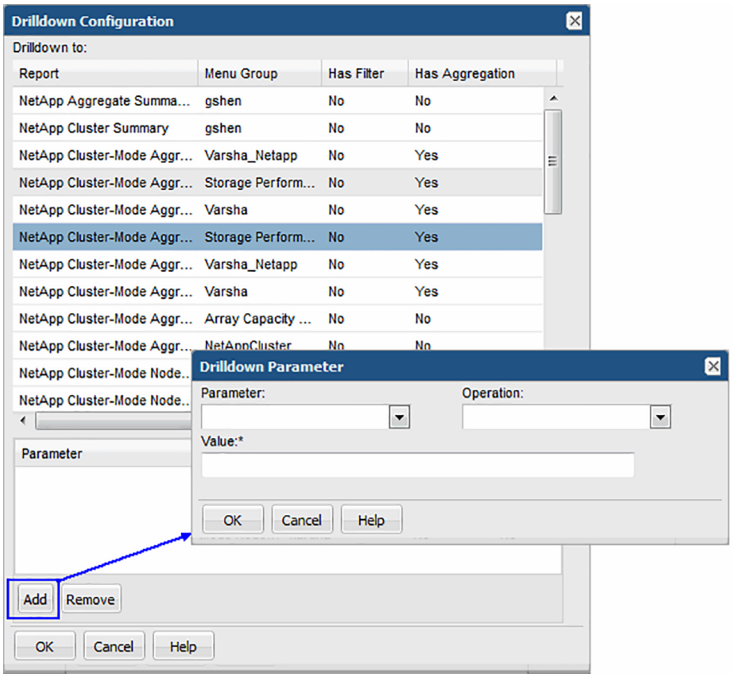
```
backupServerIds
clientIds
dateRange - dateRange is a complex Oracle type. Use the following
format to pass values
dateRange.startDate=${row['START_DATE']}&amp;dateRange.startHour=$
{row['START_HOUR']}&amp;dateRange.startMinute=${row['START_MIN']}
```

```
&amp;dateRange.startSecond=${row['START_SEC']}&amp;dateRange.finishDate=$
{row['END_DATE']}&amp;dateRange.finishHour=${row['END_HOUR']}
&amp;dateRange.finishMinute=${row['END_MIN']}&amp;dateRange.finishSecond=$
{row['END_SEC']}&amp;drilldownClientIds=${row['CLIENT_ID']}
&amp;dateRange.useFinishTime=true&amp
jobTypeIds
jobStatusIds
statusExcludeIds
includeFileList - 0 or 1
filePathName
policyNameFilter
mediaType
parentJobOnly - 0 or 1
backupWindowId
numberOfConsecutiveErrors
includeMasterServers - 0 or 1
drilldownClientIds - If this was passed in, then the original scope
 is ignored.
groupById
   Use the following values
     GROUP_BY_HOUR: 10
     GROUP_BY_DAY: 11
     GROUP_BY_WEEK: 12
     GROUP_BY_MONTH: 13
     GROUP_BY_QUARTER: 14
     GROUP_BY_YEAR: 15
```

Drilldown Parameters

You can configure additional conditions by selecting parameters, operations, and values to control when the drilldown will be active, as shown in the following example.

Note: This configuration is available only from a field that has an aggregated function and a drilldown to a report template that was configured with the Dynamic Template Designer.



Examples of Dynamic Templates Containing Drilldowns

The best way to learn how to define drilldowns in a report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A

Report Template	Report Type	Product: Alias	Caption	Field Type	Custom Drilldown or Drilldown Condition
Reports					
Backup Executive Summary	Table	All: F	Failed	Column	Condition
		All: I	Failed	Column	Custom with Condition
Job Status Summary	Bar	All: T	Failure	Bar	Drilldown
Job Summary	Table	Veritas NetBackup: P1	Client	Column	Custom
Job Volume Summary	Bar	All: D	# of Files Backed Up or Restored	Line	Drilldown
		All: C	Backup/Restore Volume	Bar	Drilldown
Job Duration	Bar	All: F	Job Duration	Bar	Drilldown
Error Log Summary by Server	Table	All: E	Error Occurrences	Column	Drilldown
		All: F	Last Error Date	Column	Custom
Error Log Summary by Client	Table	All: A	Client	Column	Drilldown
		All: F	Error Occurrences	Column	Drilldown
		All: G	Last Error Date	Column	Custom
Largest Backup Volume	Pie	All: C	Job Size	Sector	Condition

Report Template	Report Type	Product: Alias	Caption	Field Type	Custom Drilldown or Drilldown Condition
Job Summary by Server	Table	Veritas NetBackup: W	Errors	Sector	Custom with Condition
Data Domain Reports					
Data Domain Snapshot History	Table	EMC Data Domain: B	System	Column	Drilldown
Reports					
Array Executive Summary	Table	All: O	Frames	Column	Drilldown
LUN Utilization Summary	Table	All: O	# of Hosts	Column	Condition
Array Port Utilization	Table	All: O	# of Hosts	Column	Drilldown
NetApp Cluster-Mode Summary	Table	NetApp Cluster-Mode: G	# iSCSI Ports	Column	Custom with Condition
Array Utilization Summary	Table	All: N	# UnAllocated	Column	Condition

Using Groups in Dynamic Templates

Groups provide a way to logically group data in a tabular or line chart report. In a tabular report, it enables table column headings and subheadings. In a line chart report, each group will be rendered as a separate line chart and the group name will be used as the chart header in the report. Once a group is added to a template, every field in the template must reside within a group.

Groups are relevant for only Line Charts and Tables. See the following sections for details.

- See “Add a Group to Create Separate Line Charts in a Report” on page 263.
- See “Add a Group to Create a Double Header in a Tabular Report” on page 263.

Add a Group to Create Separate Line Charts in a Report

When you add a Group for a line chart report, it enables a separate line chart for each group.

Configure the following fields.

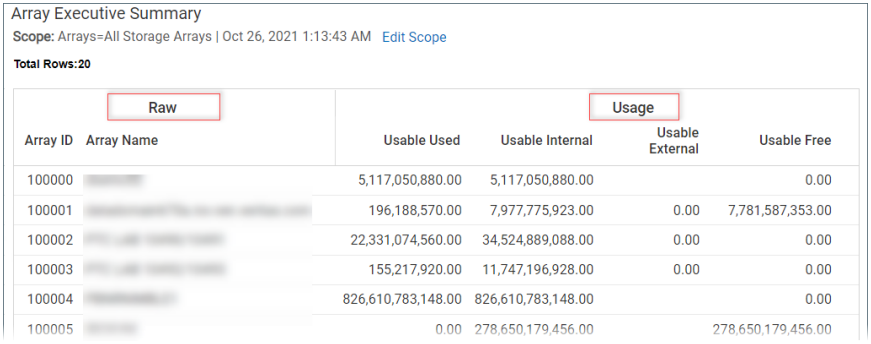
- **Group Name**: This name will appear as the header for the individual line chart.
- **Y-Axis Title**: Enter a title that will be displayed as the Y-axis label for the group's line chart.

For the **Add Group** steps,

See "Add a Group to Create a Double Header in a Tabular Report" on page 263.

Add a Group to Create a Double Header in a Tabular Report

Often, data can be presented more logically in tabular reports, by grouping the data with spanning headers and subheadings, as shown in the following example.



Array Executive Summary

Scope: Arrays=All Storage Arrays | Oct 26, 2021 1:13:43 AM Edit Scope

Total Rows:20

	Raw		Usage		
Array ID	Array Name	Usable Used	Usable Internal	Usable External	Usable Free
100000	[redacted]	5,117,050,880.00	5,117,050,880.00		0.00
100001	[redacted]	196,188,570.00	7,977,775,923.00	0.00	7,781,587,353.00
100002	[redacted]	22,331,074,560.00	34,524,889,088.00	0.00	0.00
100003	[redacted]	155,217,920.00	11,747,196,928.00	0.00	0.00
100004	[redacted]	826,610,783,148.00	826,610,783,148.00		0.00
100005	[redacted]	0.00	278,650,179,456.00		278,650,179,456.00

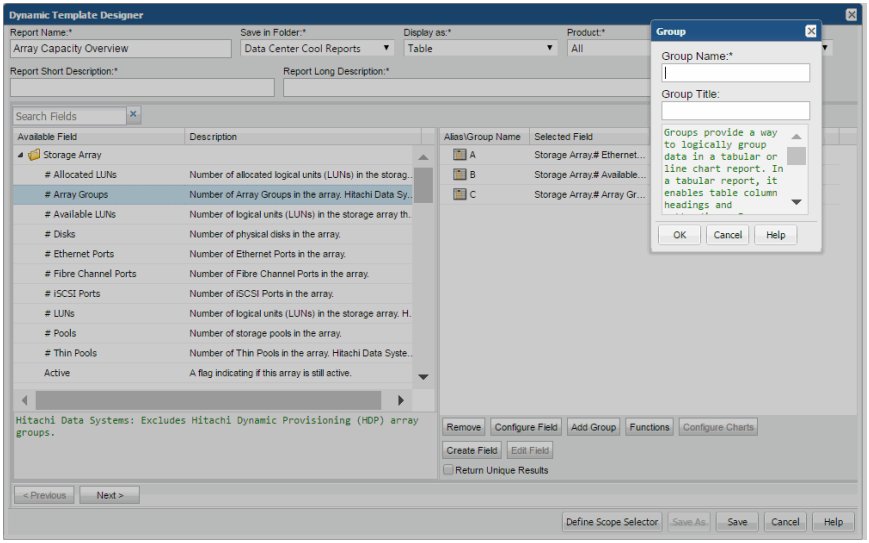
To group fields with headings and subheadings, take the following steps:

Note: The steps to add a group for a Line Chart are similar to these steps. Only the fields are different. For a description of the group fields that are relevant for Line Charts,

See "Add a Group to Create Separate Line Charts in a Report" on page 263.

---

1. In the Dynamic Template Designer window, in the Selected Fields pane, click **Add Group**.



2. Configure the following fields.
   - **Group Name**: The group name is used internally by the Dynamic Template Designer and also provides a title that will be displayed as a column heading in a tabular report. Note that in a Line Chart, the Group Name is used as the header of the individual line chart.
   - **Group Title**: Enter a title that will be displayed as a column heading in a tabular report. The group title becomes the table column heading and the field captions become the subheadings. In most cases, a group will have a title, although there may be cases where a single column of data requires no grouping. Therefore, Group Title is an optional setting.
   - Once you create a group, every selected field within the template must reside within a group, so you will likely need to create additional groups that will include the remaining fields.
   - In most cases, a group will have a title, although as shown in the following example a title is not necessary. Therefore, Group Title is an optional setting.
   - The fields within a group become the subheadings in the report.

Dynamic Template Function Configurations

Prerequisite:

See “Add Fields and Methods to a Dynamic Template” on page 237.

A variety of Oracle built-in functions can be applied to report template fields. These functions are grouped into the following sub-categories.

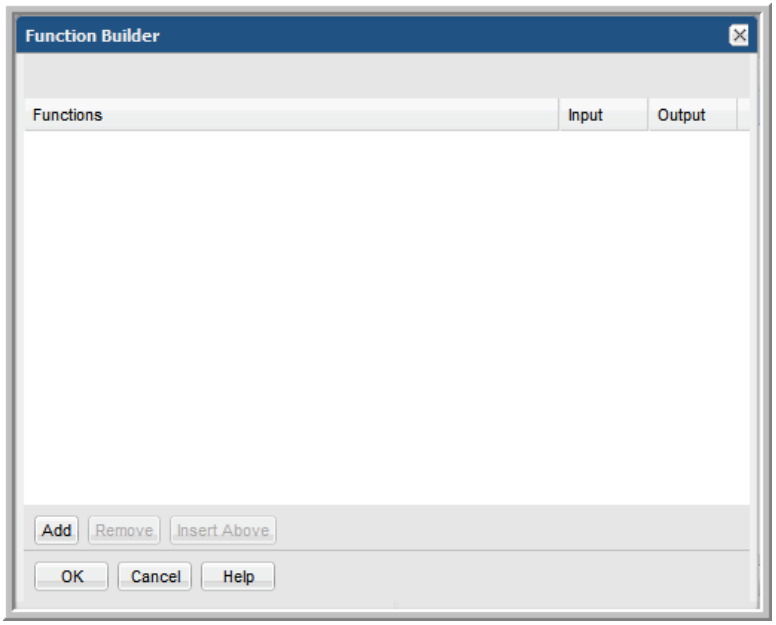
- See “Examples of Dynamic Templates Containing Functions” on page 272.
- See “Aggregation Functions” on page 267.
- See “Comparison Functions” on page 267.
- See “Numeric Functions” on page 268.
- See “Character Functions” on page 268.
- See “Date Functions” on page 270.
- See “Character String Returning Numeric Values” on page 271.
- See “Null Function” on page 271.
- See “Decode Function” on page 271.

Note: Nested functions are only supported within a function sub-category. Nesting of functions across sub-categories is only supported for the NVL function.

To apply a function to a field, take the following steps.

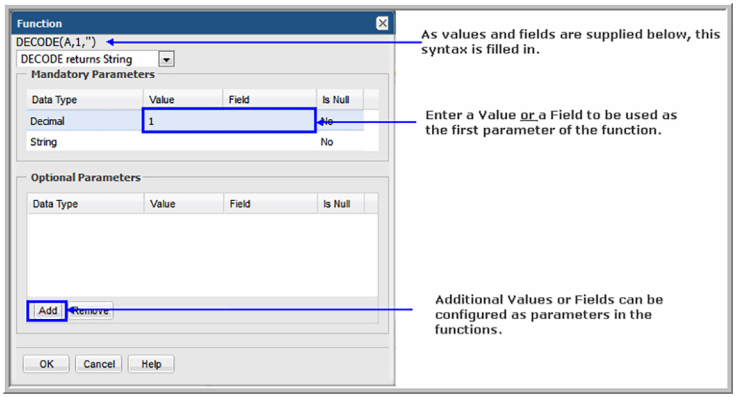
1. In the Dynamic Template Designer, select a field in the Selected Fields list at the right of the Dynamic Template Designer window.
2. Click **Functions** at the bottom of the Dynamic Template Designer, to launch the Function Builder window.

   A drop-down list of functions lets you select the function to be applied to the field.



3. Click **Add** to view the drop-down list of available functions.
4. Select a function from the list.

   Certain functions, such as DECODE, require parameters. A configuration window will display when parameters need to be configured.
5. Enter values or fields in the Function Details window, as shown in the following DECODE example.



   Note that as you configure function parameters, the syntax will auto-complete at the top of the Function Details window so that you can view how it will be implemented.

Aggregation Functions

Function	Description	Examples
AVG	Calculates the average of the values from the selected column.	Display the average hours worked for a department.
COLLECT	Takes a database column and creates a nested table from the selected rows; associates a list of data with a specific value, enabling aggregation of data into a collection.	Create a list of employees within a department.
COUNT	Lists the number of rows returned from the database for the selected column.	Determine the number of employee records returned from a database query.
MAX	Returns the maximum value of the selected column's values.	Display the maximum hours worked.
MIN	Returns the minimum value of the selected column's values.	Display the minimum hours worked.
STDDEV	Returns the standard deviation of the selected column's values.	Apply standard deviation to values, for example, for forecasting.
SUM	Calculates the sum of the values from the selected column.	Display the total hours worked for a department.

Comparison Functions

Function and Description	Examples
GREATEST From a list of expressions, Oracle determines which has the highest numeric precedence.	The following query results in BELIZE: SELECT GREATEST('BELIZE','BELGIUM','BELARUS') "GREATEST" FROM DUAL
LEAST From a list of expressions, Oracle determines which has the lowest numeric precedence.	The following query results in BELARUS: SELECT LEAST('BELIZE','BELGIUM','BELARUS') "LEAST" FROM DUAL

Numeric Functions

Function	Description	Examples
CEIL	Returns the smallest integer value that is greater than or equal to the value (n).	13551.8 becomes 13552
FLOOR	Returns the largest integer value that is less than or equal to the value (n).	13551.8 becomes 13551
ROUND	Rounds a number to the specified number of decimal places: Negative arguments indicate rounding to the left of the decimal point. Positive arguments round to the right of the decimal point.	25.193 with an argument of 1 will round to 25.2
SIGN	Returns a number that represents the sign of a value (n).	-1 if n < 0; 0 if n = 0; 1 if n >0
TRUNC	Truncates a value to the specified number of decimal places.	25.193 with an argument of 1 will truncate to 25.1

Character Functions

Function	Description	Examples
aptStringConcat	Creates a comma-separated list of strings. The Oracle function, aptStringConcat with DISTINCT or UNIQUE, cannot be used to concatenate values in a method, even though the method will validate and save. When that method is used in a report template, it will fail. Use collectString in a method to get this functionality.	aptStringConcat (array_name) can be used to create a comma-separated list of arrays within an array family.

Function	Description	Examples
collectString	Use this function in the Method Designer to concatenate distinct values.	`SELECT rtd.collectString(SET(CAST(COLLECT (cast(client_id as varchar2(10))) AS stringListType)),', ') name from apt_v_job;`
CONCAT	Concatenates multiple character strings.	'Storage' + 'HQ' = 'StorageHQ'
INITCAP	Displays a string with the first letter of each word in uppercase.	'daily backup schedule' becomes 'Daily Backup Schedule'
LOWER	Displays a string with all letters of each word in lowercase.	'DAILY BACKUP SCHEDULE' becomes 'daily backup schedule'
LPAD	Pads the left of a string with the specified characters, up to the total string length.	Preface an alert message with a string of asterisks; for example, ****Warning
LTRIM	Trims the specified character set from the left of a string. This is useful for removing redundant words, characters, or labels. When no string is supplied, it defaults to a single blank.	'RAID5' and 'RAID6' would be trimmed to simply '5' and '6'
REPLACE	Substitutes one character string for another character string; in addition, you can remove character strings.	Substituting 'HDS' for 'Hitachi' changes 'Hitachi array' to 'HDS array'
RPAD	Pads the right of a string with the specified characters, up to the total string length.	Postfix text with a string of asterisks; for example, Error***
RTRIM	Trims the specified character set from the right of a string. This is useful for removing redundant words, characters, or labels. When no string is supplied, it defaults to a single blank.	'RAID 5' and 'RAID 6' could be trimmed to simply 'RAID'
SUBSTR	Extracts a portion of a character string.	Use this function to list the first three characters of policy names.

Function	Description	Examples
TO_DATE	Converts a character string to a date.	to_date('10/09/13', 'DD/mm/YY') to_date('10-sep-13', 'DD-MON-YYYY')
TRANSLATE	Makes several single-character substitutions; one-to-one substitution in one operation.	Use this function to replace all spaces with an underscore character. 'System Reference Guide' would become 'System_Reference_Guide'
TRIM	Removes leading or trailing characters (or both) from a character string	Use this function to remove leading zeroes from object identifiers (00049 --> 49)
UPPER	Changes all characters in a string to uppercase	'veritas' becomes 'VERITAS'; 'Aptare' becomes 'APTARE'

Date Functions

Function	Description	Examples
ADD_MONTHS	Takes the date and adds n months. The result has the same day as input date. If the input date is the last day of the month or has fewer days than the resulting month, then the returning date will be the last day of the month.	June 22nd + 3 = September 22nd
ROUND	Rounds the date to the unit format that you specify: Minute, Hour, Day, Month, Quarter, Year.	If you specify Month, the function rounds up on the 16th day of the month
SYSDATE	Returns the current date and time for the system on which the database resides.	06-13-2012 09:34:41 (displayed in the format: MM-DD-YYYY HH:MM:SS)
TO_CHAR	Converts a date or interval value to a character data type in the specified format: Date, Timestamp.	finish_Date, 'DD-MON-YYYY HH24:' startDate, 'YYYY-MM-DD HH24:MI:SS'

Function	Description	Examples
TRUNC	Truncates the time portion of a day to the specified format unit: Minute, Hour, Day, Month, Quarter, Year. In Oracle, date values contain a year, month, and day and also the hour, minute, and second.	10-July-2012 09:34:41 becomes 10-July-2012

Character String Returning Numeric Values

Function	Description	Examples
INSTR	Searches a string for a sub-string and returns an integer that is the position in the string of the first character of the sub-string. The INTSR and LENGTH functions are mutually exclusive.	Search for 'error' in the string: 'A big system error caused the problem.' This results in 14.
LENGTH	Returns the number of characters of a character string. The INTSR and LENGTH functions are mutually exclusive.	Length of 'array type' would be 10

Null Function

Function	Description	Examples
NVL	Checks for a null and substitutes another value.	Use this function to display N/A instead of a blank value

Decode Function

Function	Description	Examples
DECODE	Checks for a value and if there is a match, replaces it with another constant or database field value.	Use this function to display Success if the status is 0, Warning if the status is 1, and Failed if the status is 0. `DECODE(summary_status, 0, 'Success', 1, 'Warning', 2, 'Failed')`

Unique Function

Function	Description	Examples
UNIQUE	Enables aggregation on a unique field. This function can be applied to only one field in a report template.	`COUNT(UNIQUE(client_id)), MAX(UNIQUE(server_id)), job_type FROM apt_job GROUP BY job_type`

Examples of Dynamic Templates Containing Functions

The best way to learn how to define functions in a report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
Reports					
Backup Executive Summary	Table	All: F	Failed	Column	DECODE, SUM
		All: G	Failed Count	Hidden	DECODE, UNIQUE, COUNT
		All: U, All: V	Start Date	Hidden	MAX
		All: C	Total		COUNT
		All: D	Total		UNIQUE, COUNT

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
Job Status Summary	Bar	All: T	Failure	Bar	DECODE, SUM
Job Summary	Table	Veritas NetBackup: P1	Client	Column	NVL
		Veritas NetBackup: O1	Host Id	Hidden	NVL
		CommVault Simpana: X1	Reason	Column	DECODE
		CommVault Simpana: R	Reason	Hidden	NVL
		Tivoli Storage Manager: C	Server	Column	UPPER
		HP Data Protector: V	Status	Column	NVL, DECODE
Job Volume Summary	Bar	All: D	# of Files Backed Up or Restored	Line	SUM
		All: B	# of Jobs	Hidden	COUNT
		All: C	Backup/Restore Volume	Bar	SUM
Job Duration	Bar	All: F	Job Duration	Bar	SUM
Error Log Summary by Server	Table	All: E	Error Occurrences	Column	COUNT
		All: F	Last Error Date	Column	MAX
Error Log Summary by Client	Table	All: F	Error Occurrences	Column	COUNT
		All: G	Last Error Date	Column	MAX
Error Log Summary by Policy	Table	All: E	Error Occurrences	Column	COUNT
		All: F	Last Error Date	Column	MAX

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
Consecutive Errors By Client	Table	All: N	# Consecutive Errors	Column	COUNT
		All: I	First Error Date	Column	MIN
Largest Backup Volume	Pie	All: C	Job Size	Sector	SUM
Data Domain Reports					
Data Domain Snapshot History	Table	EMC Data Domain: K	Days Left To Expire	hidden	TRUNC
		EMC Data Domain: I	Expires After (Days)	Column	SIGN, DECODE, TO_CHAR, NVL
Reports					
Array Executive Summary	Table	All: A1	Thin	Column	SUM, NVL
LUN Utilization Summary	Table	All: F Hitachi Data Systems: C Symmetrix: D NetApp: E IBM: E IBM XIV: E IBM SVC: E HP EVA: C	Device Nbr	hidden	TO_CHAR, LPAD
NetApp Cluster-Mode Volume Summary	Table	NetApp Cluster-Mode: N	Type	Column	UPPER
EMC Isilon Nodes	Table	EMC Isilon: Y	# Int Up Interfaces	Column	DECODE, TO_NUMBER, SUM
		EMC Isilon: T	External IP Address	Column	DECODE, aptStringConcat

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
EMC Isilon File System Performance by Protocol	Line	EMC Isilon: H	Max Bytes Out	Line	MAX
		EMC Isilon: A	# Active Clients	Line	AVG

Create User-Defined Fields with the Field Builder

Often a report template requires fields that are calculated or derived from other fields. You can create user-defined fields using the field builder, as described in the following sections.

Create Fields with the Field Builder

1. In the Dynamic Template Designer, click **Create Field**.

   Create a new field by combining several Alias Names into an expression, enclosed in: ${ }

   Examples:

   ${A + B}

   ${A / (A+B)}

   ${F.time - E.time}

   ${A == 0 ? 'Success' : A == 1 ? 'Warning' : 'Error'}

Convert Dates to Numbers in the Field Builder

To calculate a time duration, such as the duration of a backup job, dates need to be converted to a numeric data type.

- Append a **.time** suffix to a field to convert the date to a number.
  With this suffix applied, the result is the number of milliseconds that have elapsed since the epoch time (midnight of January 1, 1970).
  Example of a Duration Calculation:

  ```
  ${D.time - C.time}
  ```

  Where D is the Job Finish Time field and C is the Job Start Time field. Both C and D are date fields.

Syntax for Calculated Fields

Description	Syntax
Adding	+
Subtracting	-
Multiplying	*
Dividing	/
If-then-else	(condition) ? (value if true) : (value if not true)
Equality check	==
Not equal check	!=
Greater than	>
Less than	<
Greater than or equal to	>=
Less than or equal to	<=
And (joining 2 conditions)	&&
OR (joining 2 conditions)	||

Examples of Calculated Fields in a Dynamic Template

Note that all expressions need to be enclosed in the following format: **${ }**

Description	Expression
Display status ID having alias A and values 0, 1, 2 to display Success, Warning and Error.	${A == 0 ? 'Success' : A == 1 ? 'Warning' : 'Error'}
Add two numeric fields A and B.	${A + B}
Calculate percentage A from A and B.	${A / (A+B)}
Convert from A which is in KB to GB.	${A / (1024 * 1024)}
Calculate duration by subtracting values of date fields	${G.time - F.time}

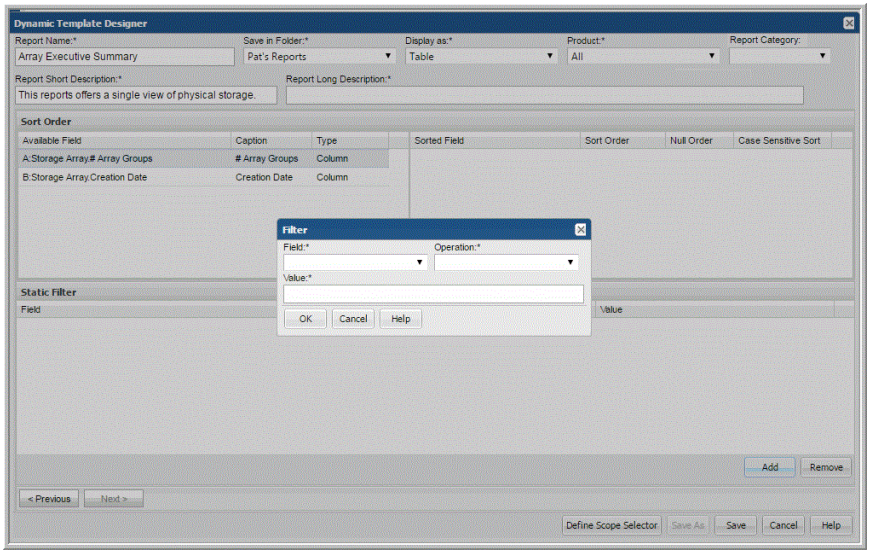
Configure a Static Filter

In addition to the scope selector, additional filtering can be achieved via the Dynamic Template Designer. This filter is a static filter--that is, unlike the filter in the Report Scope Selector, this filter cannot be changed at report run time.

Filters can be defined in the Dynamic Template Designer and also in the report scope selector. At run time, the system uses the filter that is the most restrictive. For example, if you defined a date filter of the Last 12 Hours, but your scope selector is configured for Last 4 Hours, the report results will display the events for the Last 4 Hours.

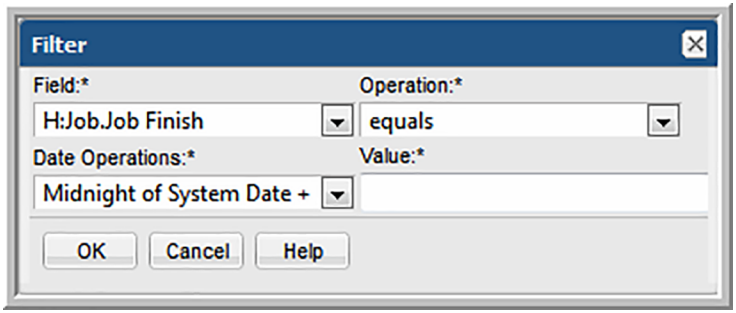
Note: When applying more than one static filter, a Boolean AND operator is used to combine the filters. A Boolean OR currently is not supported.

Click **Add** to configure the Field, Operation, and Value of a filter.



Filter on Date Fields

When the time range is not enough, additional granular filtering can be achieved. When you enter 1, the unit of measure is 1 day; 1/24 represents 1 hour.



Static Filter vs. Tabular Report Filter

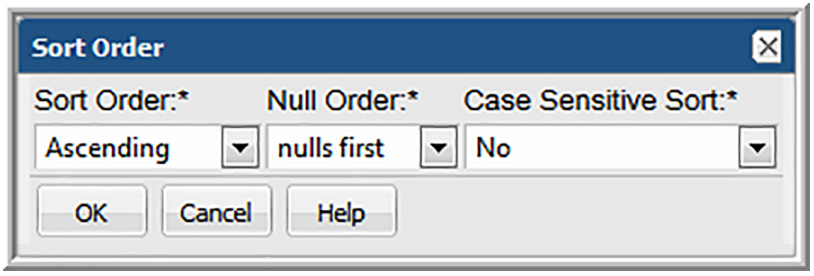
A static filter has several advantages over the filter that can be applied to a tabular report.

Filter in a Tabular Report: The tabular report filter (right-click and filter within a rendered tabular report) is a client-side filter that filters the data after it has been retrieved from the database.

Static Filter in a Dynamic Template: A static filter is a server-side filter that extracts only the filtered subset of data from the database. The static filter is useful when doing totals because the counts will reflect the sum of only the filtered data returned from the database. Also, static filters are recommended for reports where there is so much data that pagination is required when the report is rendered.

Configure the Field Sort Order

To define the **Sort Order**, drag fields into the right pane. Then, double-click the field to configure Sort Order, Null Order, and Case Sensitive Sort. The default sort order is case-insensitive.

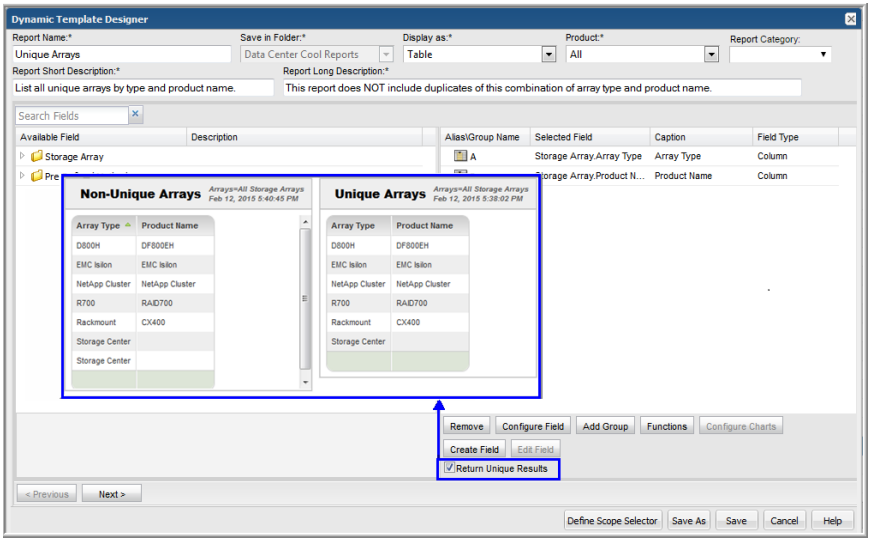


Return Unique Results

In certain situations, duplicate sets of values may be returned from the database query, causing the output to be cluttered with unnecessary data. In this case, distinct results are desired.

After defining the fields in the Dynamic Template, click the **Return Unique Results** checkbox.

A simple use case might be a list of array types by product. The following example illustrates the results, before and after the Return Unique Results checkbox is checked.



Dynamic Template Scope Selector Components

While creating a report template, you define enterprise objects such as storage arrays, host groups or a list of hosts to include in the report scope. Report scope simply refers to the criteria that you specify to filter the data that's included in a report. You can select from the these filters when building a report template for an enterprise object:

- See "Data Domain Enterprise Object Scope Selector Components" on page 284.
- See "Host Enterprise Object Scope Selector Components" on page 285.
- See "Job Enterprise Object Scope Selector Components" on page 287.
- See "Storage Array Enterprise Object Scope Selector Components" on page 289.

Note that every enterprise object will have some mandatory components that are checked and greyed out. These are automatically set by the designer to ensure optimal performance and to enable components that are required for successful rendering of the data.

Also, the list of available components is specific to the enterprise object. In addition to the enterprise object components that enable report filtering, some components enable you to control graphical representation of the data; for example, Maximum Legends for Pie Chart. And, some components enabled advanced filtering, as described in the following section.

See "Scope Selector Component - Custom Filter" on page 280.

To configure the scope selector of a report template:

1. Click **Define Scope Selector** at the bottom of the Dynamic Template Designer window.
2. In the Scope Components window, check the components that you want to appear in the Scope Selector window when generating a report from a report template. These components restrict the scope of the report.
   - **Show**: Check the box to include this component in the Scope Selector that is displayed when generating a report. Some scope components already are selected, as they are mandatory for the report template.
   - **Scope Component**: Only elements that are relevant for the report template will be listed here. Some of these components provide more robust functionality.
   - **Default Value**: The default value for the scope component, if relevant.
   - **Product**: This column indicates the product or subsystem to which these components are applicable. The notation, Base, means it applies to all products for that enterprise object.
   - **Description**: Describes the component and its purpose.

Scope Selector Component - Custom Filter

The Custom Filter scope selector component provides advanced report filtering capabilities. Use this filter to define free-form fields that enable data filtering at run time.

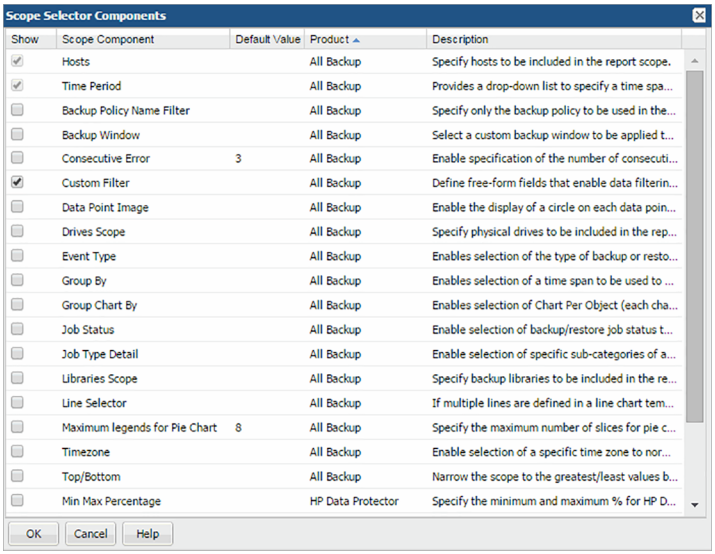
When defining a Custom Filter for a Dynamic Template, the following rules apply:

- The Dynamic Template is restricted to one product only. For example, the Dynamic Template must be defined for a single product, such as Veritas NetBackup; only data for that product is relevant.

- Fields that contain aggregation functions will not be listed for Custom Filter selection at run time.

To configure a Dynamic Template to include the Custom Filter scope selector:

1. Once you have defined the fields to be included in the template, click **Define Scope Selector** to view the Scope Selector Components. At run time, the user can supply values to achieve more granular filtering.



2. Check the **Custom Filter** component and click **OK**.

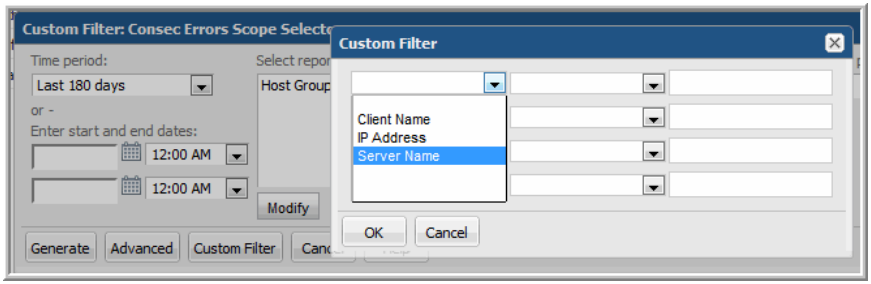
   When the report is run, the Scope Selector displays the fields that can be selected to specify values to filter the report results.

Define a Custom Filter at Run Time

This Custom Filter functionality in a Dynamic Template scope selector is particularly useful for bar and line charts, where it is desirable to tailor the amount of data reflected in the chart. This filter is similar to the advanced filtering that is available in tabular reports, although the operators are slightly different.

When you run a report, the system determines which fields can be filtered, based on what has been defined in the report template and the following restrictions:

- Numeric and String fields can be filtered.
- Date fields cannot be filtered.



Four filters can be defined by selecting a field name and an operator, then typing the value on which to filter. These filters are ANDed together to determine the results.

The following operations are permitted, by field type.

- See "Custom Filter Operators for Numeric Fields" on page 282.
- See "Custom Filter Operators for String Fields" on page 283.

Custom Filter Operators for Numeric Fields

Numeric-filtering operations have the following requirements and restrictions:

- For decimal and % values, the formatter rounds and truncates decimal places. Therefore, use the **greater than** and **less than** operators to find matches for these values.

Operator	Description
equals	Filters data where the value of the associated field is equal to the value entered.
not equal	Filters data where the value of the associated field is not equal to the value entered.
greater than	Filters data where the value of the associated field is greater to the value entered.
less than	Filters data where the value of the associated field is smaller than the value entered.
greater than or equal to	Filters data where the value of the associated field is greater than or equal to the value entered.
less than or equal to	Filters data where the value of the associated field is less than or equal to the value entered.

Operator	Description
in	Looks for a match in a comma-separated list. Wildcards are not supported. **Note:** Do not include spaces after the commas in the comma-separated list.
not in	Returns data for values that are not in the comma-separated list. Wildcards are not supported for values in this list. **Note:** Do not include spaces after the commas in the comma-separated list.
is null	Checks for null values. This is a special operation that looks for the absence of values.
is not null	Checks for non-null values. This is a special operation that looks for values that are not null.

Custom Filter Operators for String Fields

String-filtering operations have the following requirements and restrictions:

- String comparisons are case-insensitive
- Wildcards are supported for substring matching. This is applicable for the **like** and not **like operators**. Use the * to filter on substrings, as shown in the following examples.
  - Example of substring searches for arrays with the same prefix (for example, hqfin01). In this example, to look for all arrays in the hqfin group, the following wildcards could be used: **like hqfin***
- When using the **in** operator with a comma-separated list, do not include spaces after the commas in the list.

Operator	Description
equals	Filters data where the value of the associated field is equal to the value entered.
not equal	Filters rows where the value of the associated field is not equal to the value entered.

Operator	Description
in	Looks for a match in a comma-separated list. **Note:** Do not include spaces after the commas in the comma-separated list.
not in	Returns data for values that are not in the comma-separated list. **Note:** Do not include spaces after the commas in the comma-separated list.
like	Matches patterns of characters. Wildcards are supported for substring matches: *, %.
not like	Returns data for strings that are not in the pattern of characters. Wildcards are supported for substring matches: *, %.
is null	Checks for null values. This is a special operation that looks for the absence of values.
is not null	Checks for non-null values. This is a special operation that looks for values that are not null.

Data Domain Enterprise Object Scope Selector Components

The following scope selector components are specific to the Data Domain enterprise object.

Custom Filter	Define free-form fields that enable data filtering at run time. When the report is run, the Scope Selector displays the fields that can be selected to specify values to filter the report results. When the Custom Filter is selected, the report template is restricted to one product only.
Data Point Image	Enable the display of a circle of each data point in a line chart. This allows for quick identification of roll-over information in the chart.
Group By	Enables selection of a time span to be used to group the data by: Hours, Days, Weeks, Months, Quarters, or Years. When using Group By, in general, an aggregation function, such as sum, min, or max, should be included in the report template definition.

Group Chart By	Enables selection of Chart Per Object (each chart is for a single object, with one or more metric lines) or Chart Per Metric (each chart is for one metric, with one or more objects represented as lines). This is a required scope selector component if the report template includes a field of type, Group by Object.
Hosts	Specify hosts to be included in the report scope.
Line Selector	If multiple lines are defined in a line chart template, this component enables the user to select the line to be shown in the chart.
Maximum legends for Pie Chart	Specify the maximum number of sectors for pie chart rendering.
Time Period	Provides a drop-down list to specify a time span, such as last 90 days or previous month.
Top/Bottom	Narrow the scope to the greatest/least values based on a given metric.

Host Enterprise Object Scope Selector Components

The following scope selector components are specific to the Host enterprise object.

Backup Server Type	The type of host in a backup environment, such as Client or Media Server.
Byte Size	Select the units for displayed capacity values; for example, TB or PB.
Custom Filter	Define free-form fields that enable data filtering at run time. When the report is run, the Scope Selector displays the fields that can be selected to specify values to filter the report results. When the Custom Filter is selected, the report template is restricted to one product only.
Data Point Image	Enable the display of a circle of each data point in a line chart. This allows for quick identification of roll-over information in the chart.
Domain	The domain to which the enterprise object belongs.

Group By	Enables selection of a time span to be used to group the data by: Hours, Days, Weeks, Months, Quarters, or Years. When using Group By, in general, an aggregation function, such as sum, min, or max, should be included in the report template definition.
Group Chart By	Enables selection of Chart Per Object (each chart is for a single object, with one or more metric lines) or Chart Per Metric (each chart is for one metric, with one or more objects represented as lines). This is a required scope selector component if the report template includes a field of type, Group by Object.
Host	Name of the host, as defined in the collected product.
Host Creation Date	The date and time the host was added to the reporting database.
Host Type	Indicates how a host has been commissioned in an enterprise, such as VM Server, VM Guest, or Other.
Host Updated Date	The date and time that the host data was last updated in the reporting database.
Hosts	Specify hosts to be included in the report scope.
Line Selector	If multiple lines are defined in a line chart template, this component enables the user to select the line to be shown in the chart.
List of Client IDs	List of Client IDs for drilldown.
Make	Make of the host, such as Dell.
Maximum legends for Pie Chart	Specify the maximum number of sectors for pie chart rendering.
Model	Model of the host, such as Windows-x86.
OS Platform	Host's OS type, such as Linux or Windows.
OS Version	Version of the host's operating system.
Product Collected	Find hosts that have been collected from this product, such as Veritas NetBackup.
Product Group Collected	Find hosts that have been collected from this product group, such as capacity collection.
Product Group Not Collected	Find hosts that have not been collected from this product group, such as capacity collection.

Product Not Collected	Find hosts that have not been collected from this product, such as Veritas NetBackup.
Time Period	Provides a drop-down list to specify a time span, such as last 90 days or previous month.
Timezone	Enable selection of a specific time zone to normalize the report by any time zone equivalent in the world. The default setting for this parameter is the time zone setting of the location of the Data Collector. The Data Collector must be installed on a server that is in the same time zone as the subsystem from which data is collected.
Top/Bottom	Narrow the scope to the greatest/least values based on a given metric.

Job Enterprise Object Scope Selector Components

The following scope selector components are specific to the Job enterprise object. Note that some of these components are specific to a backup vendor/product.

Backup Policy Name Filter	Specify only the backup policy to be used in the report's scope.
Backup Window	Select a custom backup window to be applied to the report. Typically, backups begin at the end of the business day, but they do not finish before the end of the day-thereby skewing the success statistics for the day. To more accurately reflect backup SLA metrics, you can re-define a day with a custom backup window. These custom backup windows are defined by the Portal administrator.
By Primary Server Only	This selection filters data for the NetBackup Primary Server only.
Consecutive Error	Enable specification of the number of consecutive errors logged for the object, within the report's timeframe.
Custom Filter	Define free-form fields that enable data filtering at run time. When the report is run, the Scope Selector displays the fields that can be selected to specify values to filter the report results. When the Custom Filter is selected, the report template is restricted to one product only.

Data Point Image	Enable the display of a circle of each data point in a line chart. This allows for quick identification of roll-over information in the chart.
Drives Scope	Specify physical drives to be included in the report scope.
Event Type	Enables selection of the type of backup or restore, such as Restores or Incremental Backups.
Group By	Enables selection of a time span to be used to group the data by: Hours, Days, Weeks, Months, Quarters, or Years. When using Group By, in general, an aggregation function, such as sum, min, or max, should be included in the report template definition.
Group Chart By	Enables selection of Chart Per Object (each chart is for a single object, with one or more metric lines) or Chart Per Metric (each chart is for one metric, with one or more objects represented as lines). This is a required scope selector component if the report template includes a field of type, Group by Object.
Hosts	Specify hosts to be included in the report scope.
Job Status	Enable selection of backup/restore job status to be included in a report. For example, Warning or Failed jobs.
Job Type Detail	Enable selection of specific sub-categories of a backup vendor's jobs; for example, duplication, catalog, or vault.
Libraries Scope	Specify backup libraries to be included in the report scope.
Line Selector	If multiple lines are defined in a line chart template, this component enables the user to select the line to be shown in the chart.
List of Client IDs	List of Client IDs for drilldown.
Maximum legends for Pie	Specify the maximum number of sectors for pie chart rendering.
Min Max Percentage	Specify the minimum and maximum % for HP Data Protector sessions.
NetBackup Policy Types	List of NetBackup Policy Types.
Parent Job Only	Enable reporting on the status of the backup parent job only. For NetBackup, the Portal groups the jobs with the same parent-child relationship that NetBackup uses.

Session Status	For HP Data Protector, enable selection of a session status: Unknown, Completed, Completed/Failures, Failed, Completed/Errors, Aborted.
Time Period	Provides a drop-down list to specify a time span, such as last 90 days or previous month.
Timezone	Enable selection of a specific time zone to normalize the report by any time zone equivalent in the world. The default setting for this parameter is the time zone setting of the location of the Data Collector. The Data Collector must be installed on a server that is in the same time zone as the subsystem from which data is collected.
Top/Bottom	Narrow the scope to the greatest/least values based on a given metric.
TSM Scope	Include TSM backup/restore data in the report scope.

Storage Array Enterprise Object Scope Selector Components

The following scope selector components are specific to the Storage Array enterprise object. Note that some of these components are specific to a storage array vendor/product.

Array Disk Class	Specify array disk classes to be included in the report scope.
Array RAID Types	Enables selection of specific, relevant RAID types, such as RAID_5_6+P.
Array States	Specify storage array states to be included in the report scope.
Arrays Scope	Specify storage arrays to be included in the report scope.
Byte Size	Select the units for capacity values; for example, TB or PB.
Celerra VNX Volume Types	Include specific Celerra volume types, such as Slice, Stripe, Meta, Disk, and Pool, in a report.
Custom Filter	Define free-form fields that enable data filtering at run time. When the report is run, the Scope Selector displays the fields that can be selected to specify values to filter the report results. When the Custom Filter is selected, the report template is restricted to one product only.
Data Point Image	Enable the display of a circle of each data point in a line chart. This allows for quick identification of roll-over information in the chart.

Disk Classes	Report on specific disk classes, such as enterprise fibre channel drives (ENT) or Nearline ATA drives (NL).
Disk Stages	Enable selection of specific disk states, such as Installing, Formatting, or Rebuilding.
Disk Type	Enable selection of specific disk types, such as FC, SCSI, or SATA.
Disk Usage	Enable selection of specific disk usage types, such as Unconfigured or Spare.
Extent Pool Storage Types	Enables selection of specific, relevant Extent Pool storage types, such as fixed block.
File System Types	Include specific file system types, such as Unix, Raw, and Mirror, in a report.
Group By	Enables selection of a time span to be used to group the data by: Hours, Days, Weeks, Months, Quarters, or Years. When using Group By, in general, an aggregation function, such as sum, min, or max, should be included in the report template definition.
Group Chart By	Enables selection of Chart Per Object (each chart is for a single object, with one or more metric lines) or Chart Per Metric (each chart is for one metric, with one or more objects represented as lines). This is a required scope selector component if the report template includes a field of type, Group by Object.
Hard Empty Percentage	For IBM XIV storage pools, enable the selection of a utilization % that provides an indicator of empty physical pool capacity.
Hard Near Capacity Percentage	For IBM XIV storage pools, enable selection of a utilization % that provides an indicator of reaching physical pool capacity.
Hard Over-Provisioned High Percentage	For IBM XIV storage pools, enable selection of a utilization % that provides the high-end indicator for the range of over-provisioned physical pool capacity.
Hard Over-Provisioned Low Percentage	For IBM XIV storage pools, enable selection of a utilization % that provides the low-end indicator for the range of over-provisioned physical pool capacity.
HDS DP Pool IDs	HDS dynamic provisioning pool IDs.
HDS DP Type	HDS dynamic provisioning type.
HDS Open Reserved	HDS Open Reserved.
Hosts	Specify hosts to be included in the report scope.

IBM Array Site Disk Classes	Enable specification of IBM array site disk classes in a report. For example, ENT or NL.
IBM Array Site States	Enable specification of IBM array site states in a report. For example, Assigned, Unassigned, or Unavailable.
IBM Extent Pool IDs	Unique Identifier for the IBM extent pool.
IBM Rank IDs	List of IBM rank IDs.
In Use	For EMC VNX (Celerra) SnapSure data, provide a selection to include data for checkpoints that are in use. A checkpoint is in use if a filesystem is registered in the mount table of a Data Mover.
Line Selector	If multiple lines are defined in a line chart template, this component enables the user to select the line to be shown in the chart.
LUN Status	Enable selection of the LUN status to be included in the report. For example, Allocated but Unused or Allocated but undiscovered.
Maximum Legends for Pie Chart	Specify the maximum number of sectors for pie chart rendering.
Qtree Status	Enable selection of the NetApp Qtree status to be included in the report. For example, Snap Vaulted or Read Only.
RAID State	Enable selection of one or more array states, such as partner, zeroing, or reconstructing.
RAID Status	Enable selection of one or more raid status, such as normal, copying, reconstruct.
RAID Type	RAID type.
Rank Array States	Enable specification of IBM Rank states in a report. For example, Below, Exceeded, or Full.
Rank RAID Types	Include specific IBM Rank RAID types, such as fixed block, in a report.
Rank Status	Select an IBM rank status for report filter; for example, normal, configuring, unassigned, or reserved.
Security Style	Enable selection of the NetApp security style to be included in the report. For example, Unix, NTFS, or Mixed.
Session ID	Unique Identifier for the HPDP Session
Snapshot Busy	Choose to report data for NetApp Snapshot busy state: Yes or No.
Snapshot Only	Snapshot Only

Soft Near Capacity Percentage	For IBM XIV storage pools, enable selection of a utilization % that provides an indicator of reaching virtual (thin-provisioned) pool capacity.
Space Guarantee	Select a space guarantee for NetApp volumes; for example, volume, file, or none.
Thin-Provisioned	Included thin-provisioned array storage in the report scope.
Time Period	Provides a drop-down list to specify a time span, such as last 90 days or previous month.
Top/Bottom	Narrow the scope to the greatest/least values based on a given metric.
Volume Styles	Include specific volume style, such as flex, striped, infinitevol.
Volume Type	Select a NetApp volume type, such as flex or trad.

Customize and Export Dynamic Templates

Dynamic Templates that are shipped with the product or that are downloaded from the Cloud cannot be overwritten and will only allow save as to a different name.

To customize a Dynamic Template, search or navigate the report folders to find the Dynamic Template to be modified. When you select the template, an action bar provides a guide for actions that can be taken.

The **Customize** button launches the Dynamic Template Designer tool. All Dynamic Templates can be customized. If the template is an out-of-the box report, it can be customized and then saved in a user-defined report folder.

See “Modify an Out-of-the-Box Dynamic Template” on page 222.

See “Save a Dynamic Template After Edits” on page 294.

The **Export** button enables the export of a file that will have a file extension of: .rtd. This report template can then be imported into a user-define report folder on another Portal. Only Dynamic Templates in a user-defined report folder can be exported.

Note that IT Analytics Portal exports only digitally signed templates.

Export/Import Dynamic Templates Containing Custom Attributes

When you export a report template that was created with the Dynamic Template Designer, if the template relies on custom, user-defined attributes, those attributes are included in the exported report template definition (.rtd file). When the report template is imported into another Portal, those custom attributes will appear

automatically in that portal's reporting database. Custom attribute types supported in the Dynamic Template Designer include: Array, Drive, Host, and Library.

Note that IT Analytics Portal exports digitally signed templates and allows import of only of digitally signed templates.

Note: When the template is imported, if the import fails for any reason (such as version incompatibility), the report will not be created in the destination portal; however, the attributes will have already been created and will be available for use in the portal.

Export a dynamic template that contains attributes

When you export a report template that was created with the Dynamic Template Designer, if the template relies on custom attributes, those attributes are included in the exported report template definition (.rtd file). Attribute name validation is performed during an export to ensure that attribute names comply with the rules.

Note that all exported templates are digitally signed for security reasons.

See "Attribute naming rules" on page 550.

See "Import a dynamic template that contains attributes" on page 293.

Import a dynamic template that contains attributes

When a dynamic report template is imported into a Portal:

- Only digitally signed templates can be imported by default. The steps to import a non-digitally signed template are described below.
- Custom attributes are saved automatically in the Portal's database.
- If the system detects that the imported attributes are duplicates of existing attributes, a list of the impacted attributes and their domains will be displayed so that you can rename or edit the attributes.
  See " Edit or rename attributes " on page 544.
- If the import fails for any reason (such as version incompatibility), the report template will not be created in the destination Portal, however, the attributes that are not duplicates will have already been created and will be available for use in the Portal.
- See "Export a dynamic template that contains attributes" on page 293.

Import a non-digitally signed template

When you try to import a non-digitally signed report template, the portal displays an error: **The template <template_file_name> is not signed.** To enable importing a non-digitally signed template, you need to set the `portal.allowUnsignedReportImport` parameter to **True** as described in the procedure below.

To set the portal.allowUnsignedReportImport parameter:

1. On the IT Analytics Portal, go to **Reports** > **My Reports** and select the folder to which you want to import the non-digitally signed template.
2. Click **Import** > **Choose File** and select the template you want to import.

Save a Dynamic Template After Edits

When you edit an existing report template in the Dynamic Template Designer, you can save it with a different report name and/or a different menu group. Note that report templates that are shipped with the product or that are downloaded from the Cloud cannot be overwritten and will only allow save as to a different name.

See "Customize and Export Dynamic Templates" on page 292.

See "Modify an Out-of-the-Box Dynamic Template" on page 222.

Format the Dynamic Template Output

Often data can be best represented by a specific chart rendering. Several charting types are supported for the Dynamic Template Designer.

- See "Configure a Bar Chart Dynamic Template" on page 295.
- See "Configure an Area/Stacked Area Chart Dynamic Template" on page 303.
- See "Configure a Donut Chart Dynamic Template" on page 309.
- See "Configure a Horizontal Bar Chart Dynamic Template" on page 310.
- See "Configure a Horizontal Stacked Bar Chart Dynamic Template" on page 311.
- See "Configure a Line Chart Dynamic Template" on page 312.
- See "Configure a Pie Chart Dynamic Template" on page 324.
- See "Configure a Stacked Bar Chart in a Dynamic Template" on page 326.
- See "Configure a Table Dynamic Template" on page 328.

Configure a Bar Chart Dynamic Template

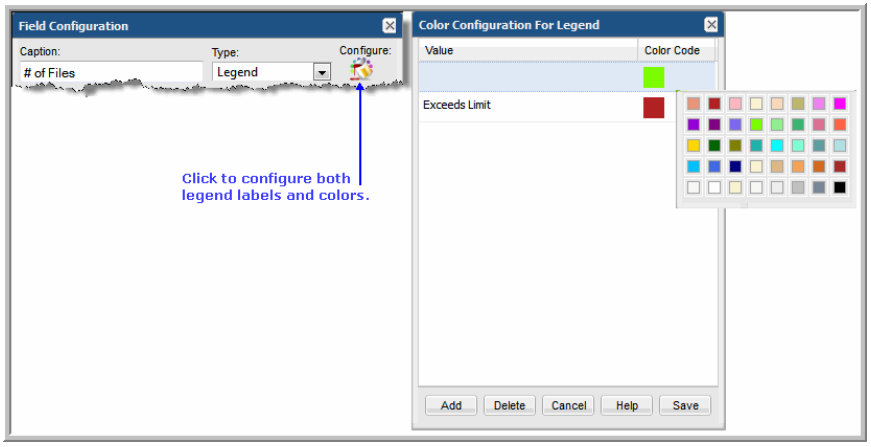
Bar charts are available as two types:

- **Dynamic**: With this data-driven type of chart, the system represents the data that is available and you don't need to know how many distinct values are available. Typically, this means that you do not select colors for the bars, but let the system select them. In some cases, colors are available for selection for a limited set of known values. A dynamic bar chart requires 1 Caption, 1 Bar, and 1 Legend. In a Dynamic Bar Chart, you cannot define the tooltips for the bars.
- **Static**: With this type of chart, the author of the template pre-defines only the values and colors that will be represented in the chart--typically, a subset of the full set of values. A static bar chart **requires** 1 Caption and at least 1 Bar. Since each bar is pre-defined, you can include a percentage line in a Static Bar Chart.

Best Practices

Use the **Static** option if you require one of the following:

- tabular report created from a bar chart
- specific set of colors in reports
- specific captions for reports; for example, job status values of Successful, Warning, and Failed.



Steps to Create a Bar Chart Dynamic Template

The easiest way to learn how to create a bar chart is by example, as described in the following sections:

Bar Chart Dynamic Template: Use Case

Use Case: As a manager, I'd like to see a list of daily backup jobs so that I can easily determine the success and failure rate. In addition, I'd like to be able to drill down to the details, especially for failed backups.

When creating a report, always start with a problem statement so that you can identify the characteristics that are essential for your report. For this example, our problem statement requires:

- Time component, such as a backup job start or finish date
- Status of the backup job, such as successful, warning, or failed

Bar charts are available as two types:

- Dynamic: With this type of chart, you don't need to know how many distinct values are available. Typically, this means that you do not select colors for the bars, but let the system select them. In some cases, colors are available for selection for a limited set of known values. A dynamic bar chart requires 1 Caption, at least 1 Bar, and 1 Legend.
- Static: With this type of chart, you pre-define only the values and colors that you want represented in the chart--typically, a subset of the full set of values. A static bar chart requires 1 Caption and at least 1 Bar.

Bar Chart Dynamic Template Step 1: Create a Dynamic Template

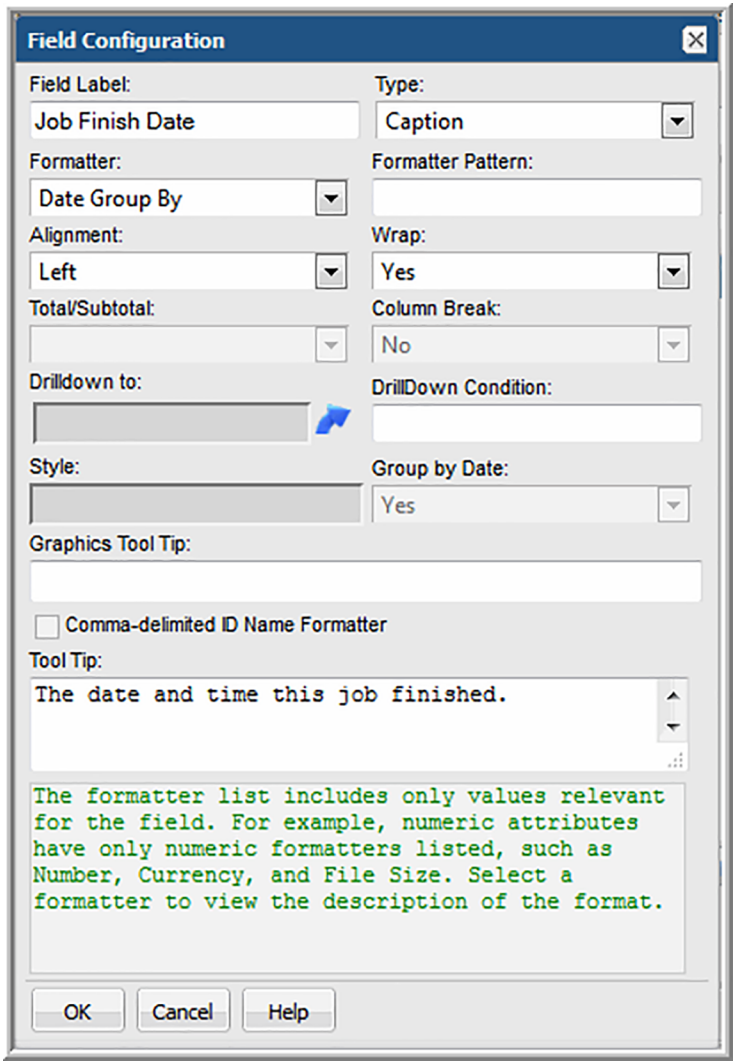
For this example--creating a report template for daily backup job status--you start with a Job enterprise object.

1. Select a Job enterprise object.
2. In the Dynamic Template Designer, enter text or select options at the top of the window. Be sure to select **Display as Bar Chart**.
3. In the Dynamic Template Designer, expand the folders in the left panel and drag the following fields into the Selected Fields panel at the right. Note that as you drag fields into the selected panel, an **Alias Name** gets assigned to the field. As you remove and add fields, these alias names do not get reused.
    - **Job Finish Date**- This provides the time component of the chart. The dates are used for the x-axis labels.

- **Summary Status** - In the description for this field, you'll see that the summary status lists the explicit codes that denote the success or failure of a backup job. For the purpose of this example, we'll be using three Summary Status fields to display Success, Warning, and Failed status.

Bar Chart Dynamic Template Step 2: Configure Basic Field Functionality

1. "/>Double-click the **Job Finish Date** field in the Selected Fields panel and configure as shown:
   - Type = **Caption**
   - Formatter = **Date Group By** (This formatting is required in conjunction with the **Group By** Scope Selector setting to group values by the selected date.)



2. Double-click the **Summary Status** field and configure as follows:
   - Field Label = **Success**
   - Type = **Bar**
   - Color = **Green**
3. Drag two additional Summary Status fields into the panel and configure each to represent the Warning and Failed backup job status bars.

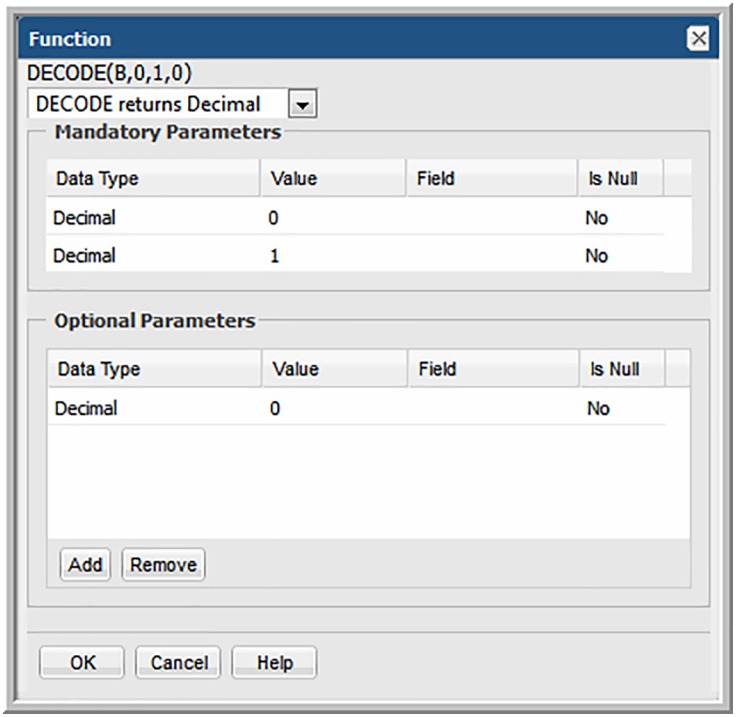
Bar Chart Dynamic Template Step 3: Configure Functions for a Field

Each of the Summary Status fields require two functions:

- DECODE, which detect and differentiate the status values, create counters for the status so that the number of backup jobs for the particular status can be tallied.
- SUM, which sums the number of backup jobs per status.

1. In the selected fields panel, select the **Summary Status** field that you've defined to represent the success status.
2. Click the **Functions** button at the bottom of the panel.
3. In the Function Builder window, click **Add**.
4. Select **DECODE returns Decimal** from the drop-down list of functions.

   Note that only functions relevant to the selected field will be available. The majority of these functions are Oracle functions that enable you to manipulate values.
5. In the Function window, two mandatory parameters and one optional parameters must be configured.



- For the first Decimal, click in the Value cell and enter a **0**. (Recall that in the description of this field, 0 = Success.)
- For the second Decimal, click in the Value cell and enter a **1**.
- Click **Add** in the Optional Parameters section and enter a value of **0**.
- Click **OK** to save the DECODE function's configuration.

This configuration tells the system that whenever a zero is encountered for a job summary status, make it a 1 so that it can be added to the count of successful jobs; then, any other status will be set to 0 so that it will not get counted in this status.

6. Configure similar settings for the Warning Summary Status field:
   - For the first Decimal, click in the Value cell and enter a **1**. (Recall that in the description of this field, 1= Warning.)
   - For the second Decimal, click in the Value cell and enter a **1**.
   - Click **Add** in the Optional Parameters section and enter a value of **0**.
   - Click **OK** to save the DECODE function's configuration.
7. Configure similar settings for the Failed Summary Status field, where the DECODE parameters will be **2**, **1**, **0**.
8. For each Summary Status field, click **Functions** and select **SUM returns Decimal**.

Bar Chart Dynamic Template Step 4: Configure Drilldowns for a Field

While a bar chart provides an at-a-glance, visual representation of backup job success, it's often useful to be able to drill down to details. You can achieve this by configuring a drilldown for each Summary Status field.

1. In the selected fields panel, double-click the **Summary Status** field that represents the success status.
2. Click the **Drilldown to** arrow.
3. In the Drilldown Configuration window, select a report that will be displayed when the user clicks on a section of a bar in the chart. For this example, select **Job Summary**.
4. Now you must tell the template what values should be passed to the drilldown report. For example, you want the Success portion of a bar to drill down to a Job Summary report that only lists successful backup jobs.
   - In the Drilldown Configuration window, in the Parameter section at the bottom, click **Add**.

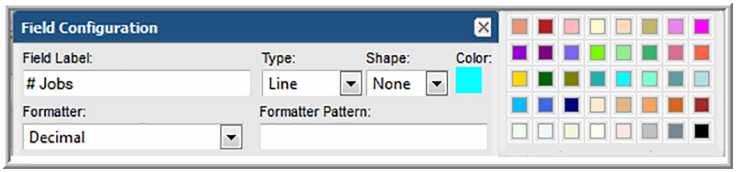
- In the Drilldown Parameter window, from the Parameter drop-down list, select the **Summary Status** field that represents the Success jobs.
- Operation = **equals**, Value = **0**

5. Repeat these steps for each of the Summary Status fields, setting **1** for Warning and **2** for Failed.
6. Next, define the scope selector options.

Bar Chart Dynamic Template Step 5: Configure a Line Field

A Line can be added to a Bar Chart to connect data points.

1. Drag a numeric field into the selected panel.
2. Double-click the field to configure the shape of the data points on the line (None, Rectangle, Triangle, Circle, or Diamond) and the color of the line.



Bar Chart Dynamic Template Step 6: Define the Scope Selector

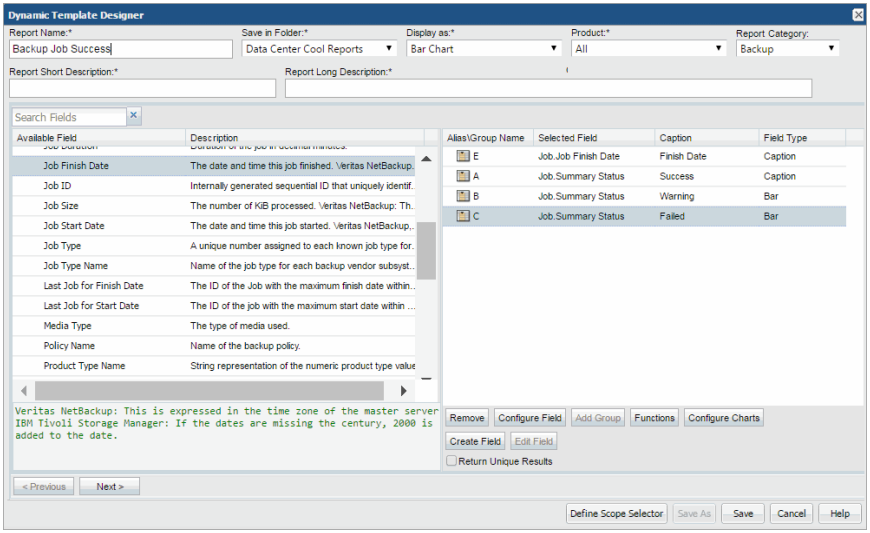
1. At the bottom of the Dynamic Template Designer window, click **Define Scope Selector**.
2. In the Scope Selector Components window, check **Group By**.

   This setting must be selected so that the data can be grouped rather than enumerated for every time value. When the user runs this report template, a Group By selector will be available to choose the time grouping: Hours, Days, Weeks, Months, Quarters, or Years.
3. Click **OK** to return to the main Dynamic Template Designer window.

Bar Chart Dynamic Template Step 7: Save and Run the Report Template

1. In the main Dynamic Template Designer window, click **OK**to save your work.

   Your final report template window should look something like this:



2. Find your report template in the **Reports** window and run it to check your results.

Examples of Bar Chart Dynamic Templates

The best way to learn how to define a bar chart report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
Reports					
Job Status Summary	Bar	All: T	Failure	Bar	DECODE, SUM

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
Job Volume Summary	Bar	All: D	# of Files Backed Up or Restored	Line	SUM
		All: B	# of Jobs	Hidden	COUNT
		All: C	Backup/Restore Volume	Bar	SUM
Job Duration	Bar	All: F	Job Duration	Bar	SUM

Configure an Area/Stacked Area Chart Dynamic Template

- See "Area/Stacked Area Chart Dynamic Template Step 1: Create a Dynamic Template" on page 303.
- See "Bar Chart Dynamic Template Step 2: Configure Basic Field Functionality" on page 297.
- See "Bar Chart Dynamic Template Step 3: Configure Functions for a Field" on page 299.
- See "Bar Chart Dynamic Template Step 5: Configure a Line Field" on page 301.
- See "Bar Chart Dynamic Template Step 7: Save and Run the Report Template" on page 301.

Area/Stacked Area Chart Dynamic Template Step 1: Create a Dynamic Template

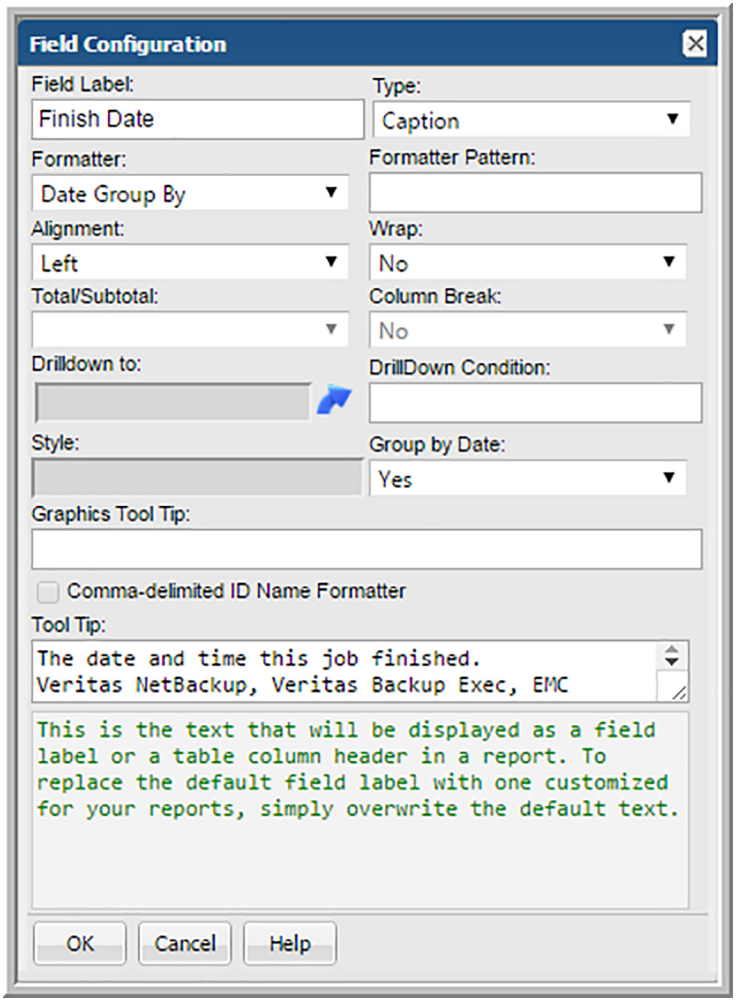
For this example--creating a report template for daily backup job status--you start with a Job enterprise object.

1. Select a Job enterprise object.
2. In the Dynamic Template Designer, enter text or select options at the top of the window. Be sure to select **Display as Area/Stacked Area Chart**.
3. In the Dynamic Template Designer, expand the folders in the left panel and drag the following fields into the Selected Fields panel at the right. Note that as you drag fields into the selected panel, an **Alias Name** gets assigned to the field. As you remove and add fields, these alias names do not get reused.

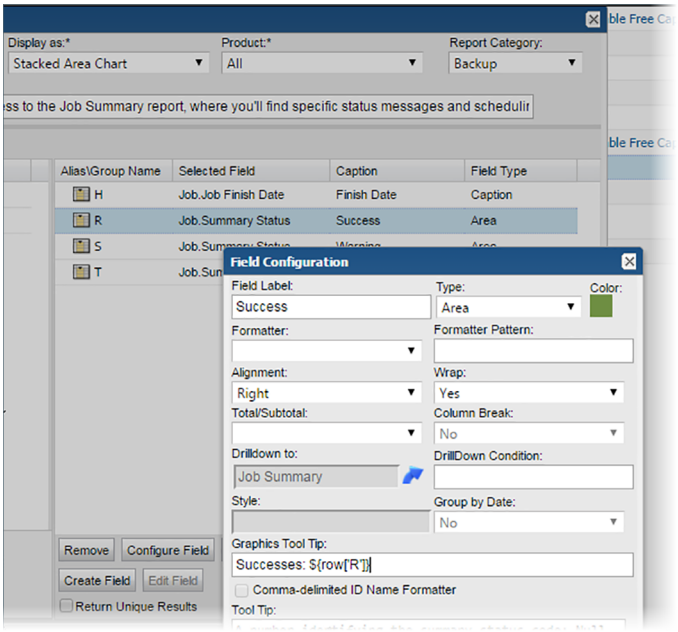
- **Job Finish Date**- This provides the time component of the chart. The dates are used for the x-axis labels.
- **Summary Status** - In the description for this field, you'll see that the summary status lists the explicit codes that denote the success or failure of a backup job. For the purpose of this example, we'll be using three Summary Status fields to display Success, Warning, and Failed status.

Area/Stacked Area Chart Dynamic Template Step 2: Configure Basic Field Functionality

1. "/>Double-click the **Job Finish Date**field in the Selected Fields panel and configure as shown:
   - Type = **Caption**
   - Formatter = **Date Group By** (This formatting is required in conjunction with the **Group By** Scope Selector setting to group values by the selected date.)



2. Double-click the **Summary Status** field and configure as follows:
   - Field Label = **Success**
   - Type = **Area**
   - Color = **Green**



3. Drag two additional Summary Status fields into the panel and configure each to represent the Warning and Failed backup job status bars.

Area/Stacked Area Chart Dynamic Template Step 3: Configure Functions for a Field

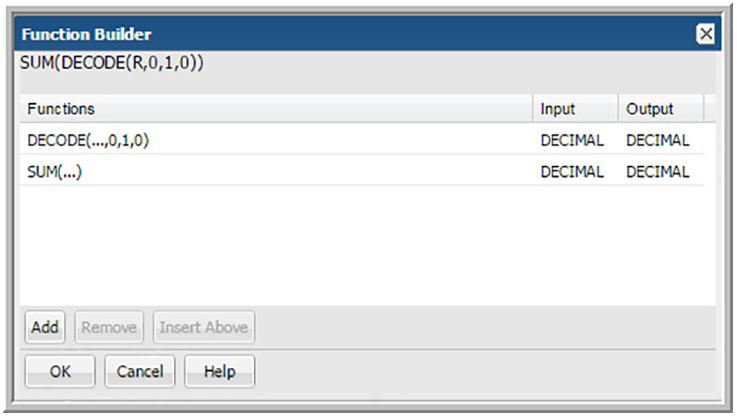
Each of the Summary Status fields require two functions:

- DECODE, which detect and differentiate the status values, create counters for the status so that the number of backup jobs for the particular status can be tallied.
- SUM, which sums the number of backup jobs per status.

1. In the selected fields panel, select the **Summary Status** field that you've defined to represent the success status.
2. Click the **Functions** button at the bottom of the panel.
3. In the Function Builder window, click **Add**.
4. Select **DECODE returns Decimal** from the drop-down list of functions.

   Note that only functions relevant to the selected field will be available. The majority of these functions are Oracle functions that enable you to manipulate values.

5. In the Function window, two mandatory parameters and one optional parameters must be configured.



   - For the first Decimal, click in the Value cell and enter a **0**. (Recall that in the description of this field, 0 = Success.)
   - For the second Decimal, click in the Value cell and enter a **1**.
   - Click **Add** in the Optional Parameters section and enter a value of **0**.
   - Click **OK** to save the DECODE function's configuration.

   This configuration tells the system that whenever a zero is encountered for a job summary status, make it a 1 so that it can be added to the count of successful jobs; then, any other status will be set to 0 so that it will not get counted in this status.

6. Configure similar settings for the Warning Summary Status field:
   - For the first Decimal, click in the Value cell and enter a **1**. (Recall that in the description of this field, 1= Warning.)
   - For the second Decimal, click in the Value cell and enter a **1**.
   - Click **Add** in the Optional Parameters section and enter a value of **0**.
   - Click **OK** to save the DECODE function's configuration.
7. Configure similar settings for the Failed Summary Status field, where the DECODE parameters will be **2**, **1**, **0**.
8. For each Summary Status field, click **Functions** and select **SUM returns Decimal**.

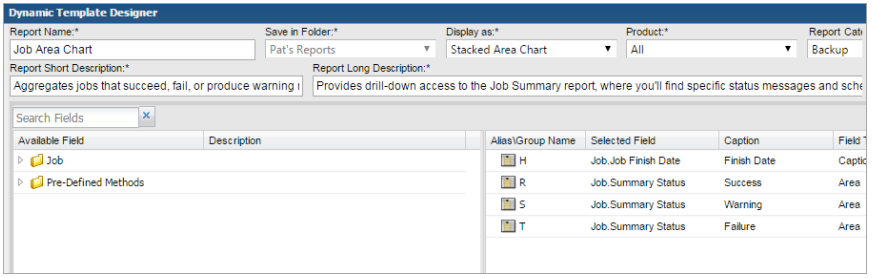
Area/Stacked Area Chart Dynamic Template Step 4: Define the Scope Selector

1. At the bottom of the Dynamic Template Designer window, click **Define Scope Selector**.
2. In the Scope Selector Components window, check **Group By**.

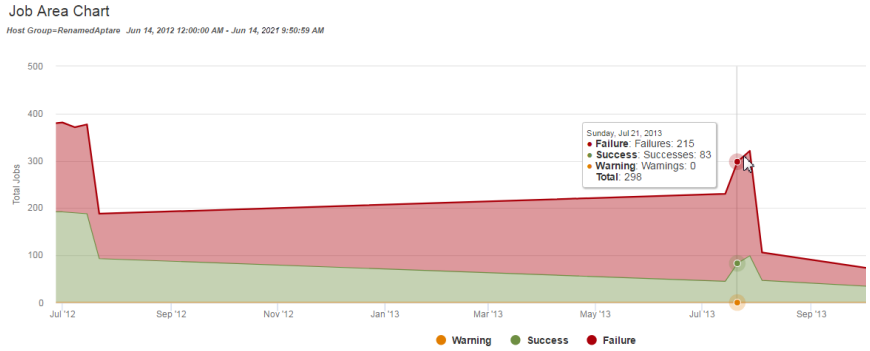
   This setting must be selected so that the data can be grouped rather than enumerated for every time value. When the user runs this report template, a Group By selector will be available to choose the time grouping: Hours, Days, Weeks, Months, Quarters, or Years.
3. Click **OK** to return to the main Dynamic Template Designer window.

Area/Stacked Area Chart Dynamic Template Step 5: Save and Run the Report Template

1. In the main Dynamic Template Designer window, click **OK**to save your work. Your final report template window should look something like this:



2. Find your report template in the Reports window and run it to check your results. The end result should look similar to the following report.



Configure a Donut Chart Dynamic Template

A Donut Chart is similar to a Pie Chart.

- See "Configure a Pie Chart Dynamic Template" on page 324.
- See "Example of a Donut Chart Dynamic Template" on page 309.

Example of a Donut Chart Dynamic Template

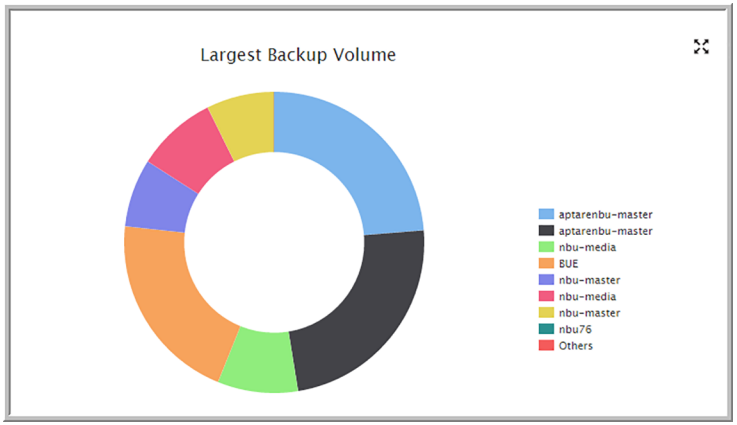
Donut charts require both a caption field and a numeric sector field. To create sectors for a non-numeric field, such as OS name, you simply can configure a Count function to render the count of the instances of the value.

1. In the Dynamic Template Designer window, select **Donut Chart** from the **Display as** list.
2. Drag or double-click fields to add them to the pane at the right of the window. Typically, you will need only a caption field and a sector field.
3. **Sector Field**: Double-click the field to configure it as a Donut sector.
4. **Caption Field**: Double-click the field to configure it as a caption.
5. Select the **Sector** field and click **Functions**.
6. In the Function Builder window, click **Add** and select a **COUNT** function to count the instances of the value to return a decimal that can be graphed as a donut sector.
7. For this example, the Donut Chart should include the following fields and field types.

Alias\Group Name	Selected Field	Caption	Field Type
E	Job.Client ID	Client Id	Hidden
D	Host (Client ID).Host Disp...	Display Name	Caption
C	Job.Job Size	Job size	Pie
F	Job.Summary Status	Summary Status	Exclude
G	Host (Client ID).Host ID	Host Id	Hidden

8. Be sure to complete the other mandatory fields for the template: Report Name, Short Description, and Long Description.
9. Click **Define Scope Selector** to configure the **Maximum legends for Pie/Donut Chart**. Overwrite the default value with the maximum that you want to see.
10. Save the Report Template.
11. Generate the report from the Report Template.

The report will look something like this:



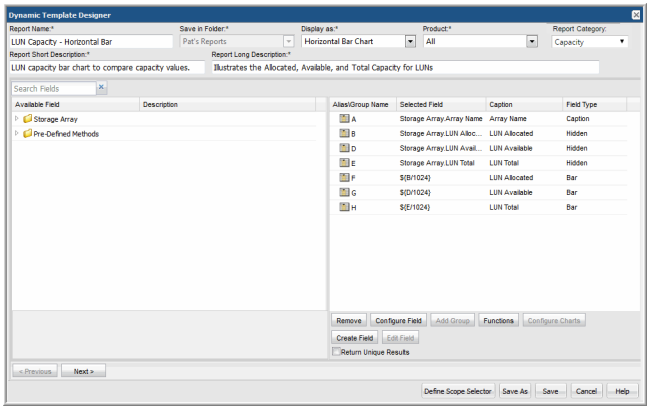
Configure a Horizontal Bar Chart Dynamic Template

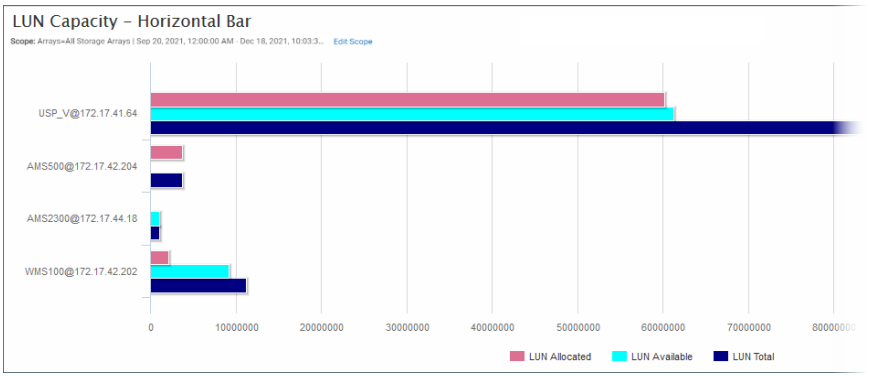
A horizontal bar chart is similar to a bar chart. The only thing that differs is the orientation, horizontal vs. vertical. To build a horizontal bar chart, refer to the steps in the following sections and simply choose a display type of Horizontal Bar Chart, instead of Bar Chart.

- See "Configure a Bar Chart Dynamic Template" on page 295.
- See "Example of a Horizontal Bar Chart Dynamic Template" on page 310.
- See "Configure a Horizontal Stacked Bar Chart Dynamic Template" on page 311.
- See "Configure a Line in a Bar Chart" on page 327.

Example of a Horizontal Bar Chart Dynamic Template

This example generates a horizontal stacked bar chart to represent LUN capacity: allocated, available, and total.



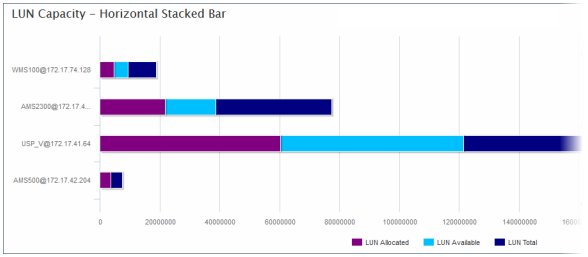


Configure a Horizontal Stacked Bar Chart Dynamic Template

A horizontal stacked bar chart is similar to a stacked bar chart. The only thing that differs is the orientation, horizontal vs. vertical. To build a horizontal stacked bar chart, refer to the steps in the following sections and simply choose a display type of Horizontal Stacked Bar Chart, instead of Bar Chart.

- See “Configure a Bar Chart Dynamic Template” on page 295.
- See “Example of a Horizontal Stacked Bar Chart Dynamic Template” on page 312.
- See “Configure a Horizontal Bar Chart Dynamic Template” on page 310.
- See “Configure a Line in a Bar Chart” on page 327.

Example of a Horizontal Stacked Bar Chart Dynamic Template



Configure a Line Chart Dynamic Template

Line charts provide an effective way to visualize metrics on a time line. These charts are particularly useful for performance metrics, where you want to compare multiple metrics to assess historical performance. There are times when a single chart containing multiple metrics will suffice. Likewise, there are times when a better comparison can be made with multiple charts containing a single metric per chart.

For details on two line chart approaches, see:

- See "One Object Per Line Chart, One or More Metrics Per Chart" on page 318.
  - For example, view the total, allocated, and available capacity history for each array in the report's scope.
- See "Multiple Objects Per Line Chart, One Metric Per Chart" on page 321.
  - For example, compare available capacity for 10 arrays in a single chart.

Note: Tip: As a best practice, choose a line chart that is provided as an out-of-the-box report template, then copy and customize it to see how it is configured.

Line charts can be configured for the following:

- **All subsystems**- For example, all storage array vendors in your environment.
- **Single subsystem** - For example, a single storage array vendor, such as EMC Symmetrix. You cannot multi-select a subset of all of your subsystems for a report template.

See "Steps to Create a Line Chart Dynamic Template" on page 314.

Line Charts for Performance Metrics

Line charts provide a meaningful visualization of performance over time. Use line charts to represent performance data, such as array performance metrics.

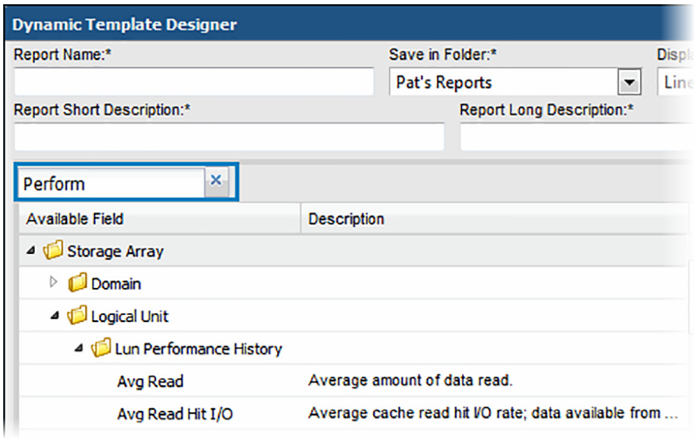
Array performance data is collected from certain storage arrays:

- NetApp ONTAP Cluster-Mode arrays
- Isilon arrays
- Array LUN performance statistics are collected for Dell Compellent, EMC VNX (CLARiiON), EMC Symmetrix, HDS Tuning Manager, HP 3PAR, IBM SVC, IBM XIV, NetApp ONTAP (Block only), NetApp Cluster-Mode and Pure Storage.

For collected array LUN performance raw data, the default retention period (504 hours or 21 days) controls the retention of the data. When the daily performance data ages out, it gets deleted.

Identify Array Performance Data Fields for a Template

Use the **Search** feature in the Dynamic Template Designer to find performance statistics in the database fields.



Historical performance metrics are collected for the following enterprise objects:

- **Host** - Chargeback, CPU, Memory, and Network history is available. Search on History in the Dynamic Template.
- **Job** - The backup/restore job itself has historical information associated with it. Each job record is the history.

- **Storage Array** - Search on Log in the Dynamic Template. Search on History in the Dynamic Template.

Rules and Guidelines for Performance Line Charts

Several rules govern the creation of a line chart for performance metrics, as listed in the following table.

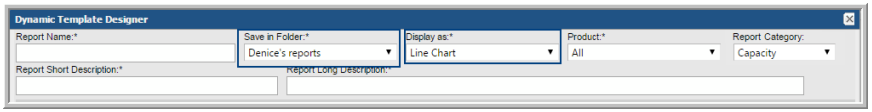
Rule for Performance Metrics Line Charts	Description/Example
1. If aggregation is defined for a line, all lines must have an aggregation specified.	See "Aggregation Functions" on page 267. See "Line Chart Field Requirements" on page 315.
2. The Group By scope selector component is required only if the lines have aggregation functions applied.	See "Dynamic Template Scope Selector Components" on page 279.
3. If no aggregation has been specified for lines, the raw data will be returned in the report.	Array performance data is captured and stored as raw, hourly, and daily records. When raw data is used, the dates reflect exactly when the data was captured, including the seconds.
4. A date field should be defined as the caption.	This provides the x-axis time period markers. Note that time-series charts automatically calculate the x-axis labels based on the data. Therefore, your Group By selection in the scope selector may not always render with the set of dates that you expect for the time period.
5. The Time Period scope selector is always required for performance line charts.	Performance metrics are best represented as time-series charts. See "Dynamic Template Scope Selector Components" on page 279.
6. Use the **Add Group**option to create folders when you need to group data fields that have disparate units of measure.	For example, you might have a capacity value in KiB, but a performance value in Kbps. This requires two charts with different metrics scales to plot the values. See "Add a Group to Create Separate Line Charts in a Report" on page 263.

Steps to Create a Line Chart Dynamic Template

To create a line chart report template, start with the following steps.

1. Select **Reports > My Reports> Your Custom Name Folder**.
2. Click the **New Dynamic Template** button.
3. Choose an **Enterprise Object** (Data Domain, Host, Job, or Storage Array) and click **OK**.

4. In the Dynamic Template Designer, select a folder to save the template into and the display type. Be sure to select Line Chart, as it enables options specific to this feature.



5. Enter a report name, short description, and long description.
6. Optionally, assign a category to set how report templates are categorized in the Inventory. When displayed in the Inventory, templates are sorted into information categories such as performance, storage or forecasting. If a category has not been selected, these templates are displayed under the heading Uncategorized. You can always assign a category, by customizing an existing template.
7. Drag fields into the report scope and configure the fields.
8. At the bottom of the Dynamic Template Designer, click **Configure Charts** to define the label and metric unit to be displayed in the y-axis.
9. At the bottom of the Dynamic Template Designer, click **Define Scope Selector**.

Line Chart Field Requirements

In its simplest form, a line chart requires a **Caption** field and any number of **Line** fields. For more complex scenarios,

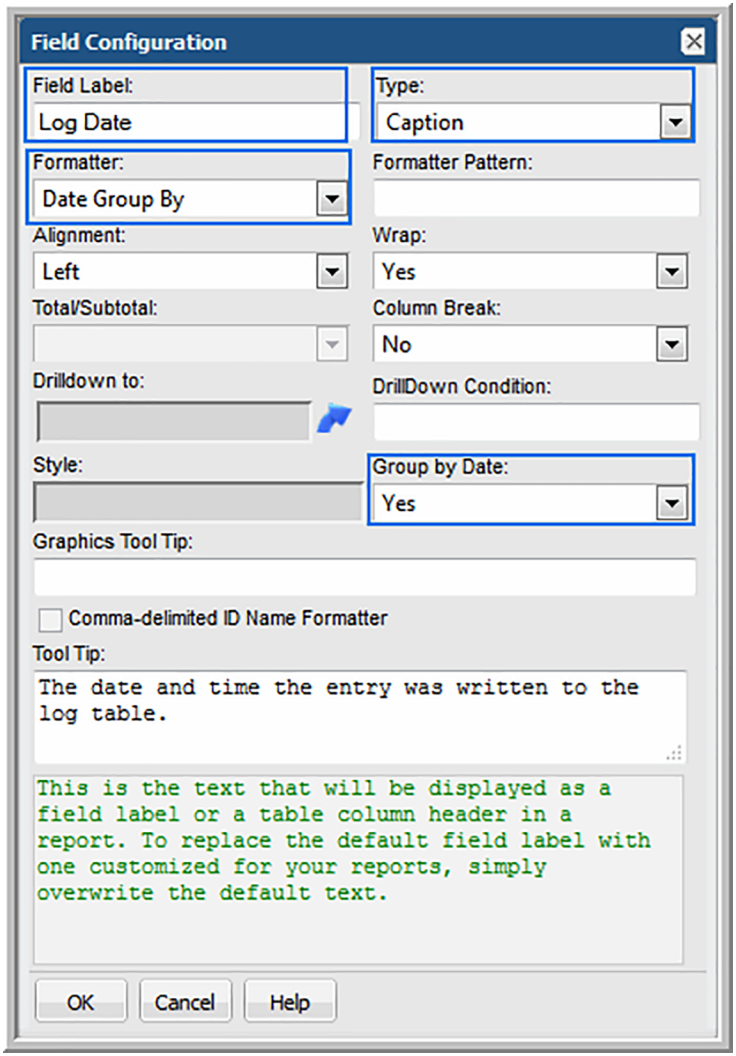
See “One Object Per Line Chart, One or More Metrics Per Chart” on page 318.

See “Multiple Objects Per Line Chart, One Metric Per Chart” on page 321.

Line Chart: Caption Field

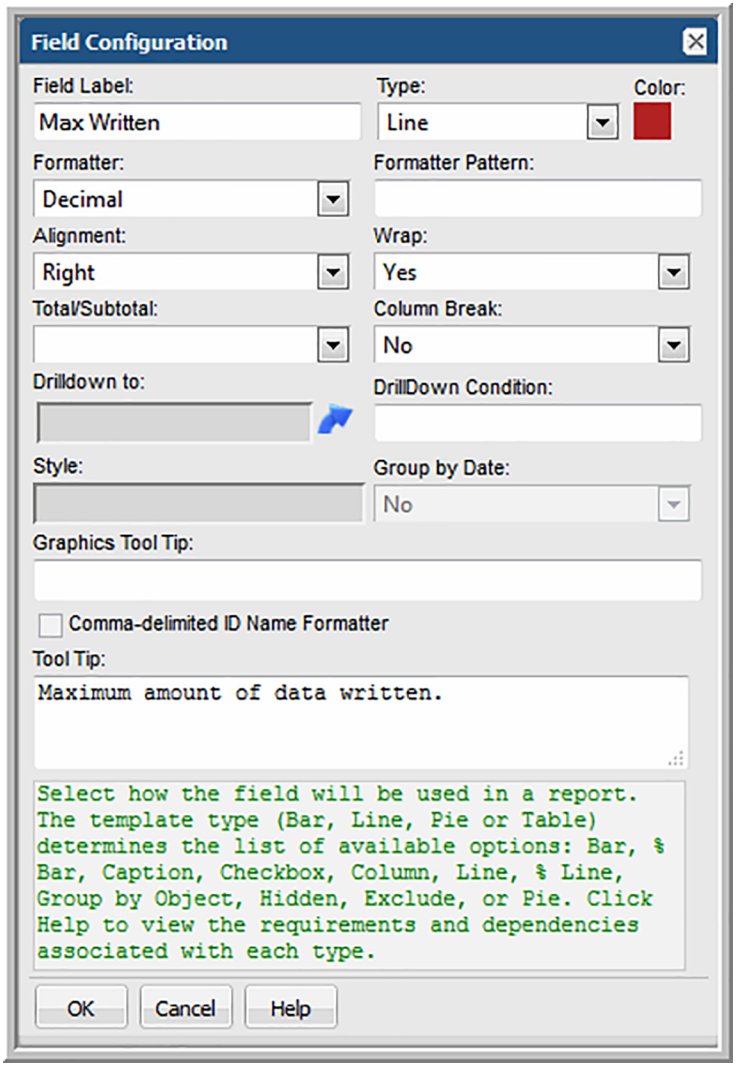
This required field type supplies the category caption for the x-axis. Typically, a date field becomes the caption because line chart data is best represented over a time line. When you configure a date field as a caption, the field configuration automatically sets **Group By** to **Yes**, so that in the scope selector, a list of time frames is provided.

See “Format Line Chart Fields” on page 322.



Line Chart: Line Field

At least one line field must be defined in a line chart template. A numeric field can be configured as a line. One or more line fields can be included in a line chart report template. Select **Line** for the field type and then select a color Also, use an aggregation function, such as SUM to provide meaningful values in the chart.



Note that line colors do not apply to the “multiple objects per chart” scenario. In this scenario, line colors will be determined dynamically by the system, since the number of objects represented by the lines is unknown until run time.

Line Chart: Group by Object Field

When you configure a Group by Object field, essentially you are defining, “for each object, display metric lines.” Group by Object has the following requirements:

- Only one **Group by Object** field can be defined for a line chart.
- When a **Group by Object** field is configured, one field in the template must have **Column Break** set to **Yes**. This Column Break field should have unique values. For example, LUN names typically are duplicated in a storage environment, but the LUN ID uniquely identifies a LUN. In this example, a LUN ID should be included in a template as a Hidden field with a Column Break. Only one field in a template can have the Column Break set to Yes.
- If the line chart report template includes a field of type, **Group by Object**, the scope selector must have the **Top/Bottom** and the **Group Chart By**selections checked.

Line Chart: Hidden Field

A Hidden field comes in handy when you want to ensure that an object is unique, but you don't want to display that field in the report. For a scenario where a Hidden field is useful, refer to the following.

See "Line Chart: Group by Object Field" on page 317.

Line Chart: Exclude Field

There will be cases when a field will be needed only for filtering purposes and you do not want this field to be included in the selected fields for the report template. For example, you may want to filter on backup jobs that are only failed and partial. In this case, you would drag in the job status field, make it an Exclude field type, and configure a static filter.

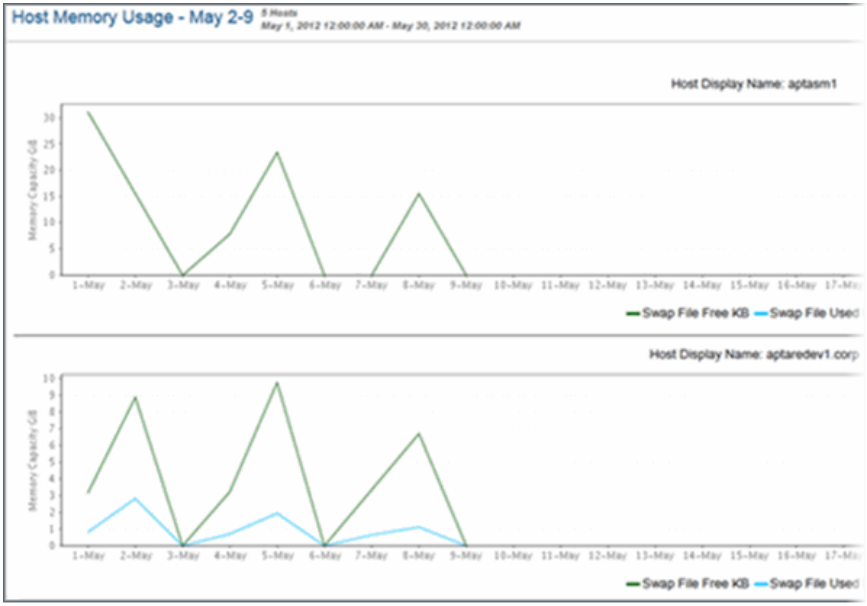
See "Configure a Static Filter" on page 277.

One Object Per Line Chart, One or More Metrics Per Chart

A single line chart can be used to visualize multiple metrics for a single enterprise object. For example, you might have a host with degraded performance and you want to view all the metrics for that host within a given time period. You want to examine swap file activity, free and used, over time. Each chart would have two metrics lines. At run time, the user might select 10 hosts and the resulting report would show 10 charts, one per host enterprise object.

At run time in the Scope Selector:

- Choose one or more objects to be included in the chart.
- Choose from the **Group Chart By** pick list: Chart Per Object.



Template Fields to Render Two Metrics in a Chart

Alias\Group Name	Selected Field	Caption	Field Type
E	Host Memory History.Log...	Log Date	Caption
F	Host Memory History.Sw...	Swap File Free KB	Line
G	Host Memory History.Sw...	Swap File Used KB	Line
H	Host.Host ID	Host ID	Group By Object

Configure One Object Per Line Chart

Configure the following settings to create a report template for a chart that contains several metric lines per object chart. Each object in the report's scope will have its own chart in the report.

Field Configuration for One Object Per Line Chart

- **Caption Field**- This required field type supplies the category caption for the x-axis. Typically, a date field becomes the caption because line chart data is best represented over a time line. When you configure a date field as a caption, the field configuration automatically sets **Group By Date**to **Yes**.
- **Line Field** - A numeric field can be configured as a line. One or more line fields can be included in a line chart report template. Select **Line** for the field type and then select a color.

- **Add Group** - (Optional) Use this option to create folders when you need to group data fields that have disparate units of measure. For example, you might have a capacity value in KiB, but a performance value in Kbps. This requires two charts with different metrics scales to plot the values. In the Dynamic Template Designer, click **Add Group** to create a folder for each group so that each group will have its own chart. Once you create a group, every field within the template must reside within a group, so you will likely need to create additional groups that will include the remaining fields. Note that the Group functionality is used in tabular reports to create spanning headers, but in the case of line charts, it enables separate charts.
- **Column Break** - (Optional) Enable this setting on either an object name or object ID; for example, array name or array ID. This setting creates a chart for each object, such as each array in the report scope. For Line Chart report templates, if the template includes a field of type, Group By Object, one field must have Column Break set to Yes. Typically, a field of type, Group By Object or Hidden, is used for a Column Break. This field must have unique values, such as an Array ID. Only one field can have the Column Break set to Yes.
- **Group by Date** - If a date field is configured as a caption, the Group by Date setting is automatically set. However, if more than one date field is included in the template, you must select which date field should have the **Group by Date** setting set to **Yes**.

Scope Selector Configuration for One Object Per Line Chart

- **Group By -** If the scope selector has the Time Period component checked, the Group By component must also be checked. This enables the user to select the time units at run time: Hours, Days, Weeks, Months, Quarters, or Years.
- **Group Chart By**- This component lets the user select: Chart Per Object or Chart Per Metric. If the line chart report template includes a field of type, Group By Object, the scope selector must have the Group Chart By selection checked.
- **Time Period**- Use this component to specify the time span for the generated report.
- **Top/Bottom** - This enables selection of the Top n or Bottom n values, based on a metric/field. If the line chart report template includes a field of type, Group By Object, the scope selector must have the Top Bottom selection checked. When Top/Bottom is selected, an additional “Based on” selector will appear in the scope selector to choose an associated metric.

Multiple Objects Per Line Chart, One Metric Per Chart

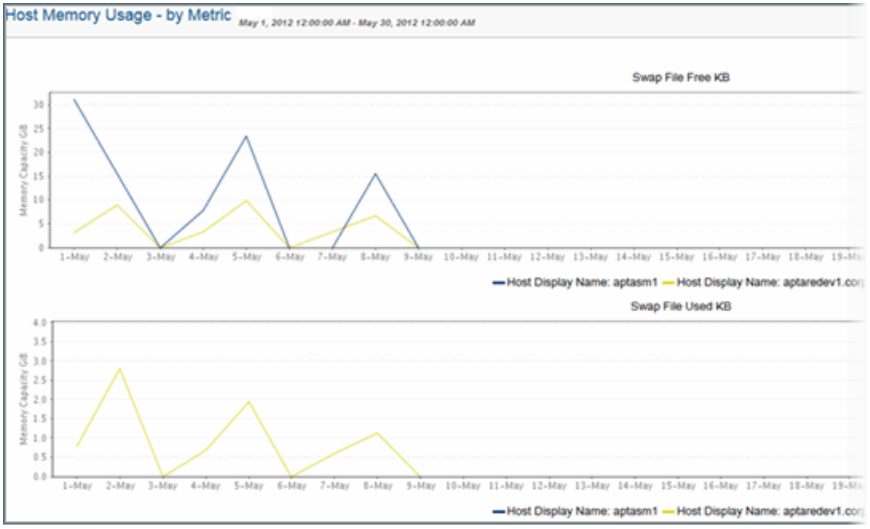
A single line chart can be used to compare a single metric for a group of enterprise objects. As shown in the following example.

See "One Object Per Line Chart, One or More Metrics Per Chart" on page 318.

The same template definition can be used. At run time, the user selects Chart Per Metric and chooses 10 hosts. The resulting report would show two charts, one for swap file free and second for swap file used, with each chart displaying 10 lines for the 10 hosts in the scope.

At run time in the Scope Selector:

- Choose one or more objects to be included in the chart.
- Choose from the **Group Chart By** pick list: Chart Per Metric.



"Group By Metric" Selection in Scope Selector Displays Multiple Hosts Per Metric

Alias\Group Name	Selected Field	Caption	Field Type
E	Host Memory History.Log...	Log Date	Caption
F	Host Memory History.Sw...	Swap File Free KB	Line
G	Host Memory History.Sw...	Swap File Used KB	Line
H	Host.Host ID	Host ID	Group By Object

Field Configuration for Multiple Objects Per Line Chart, One Metric Per Chart

- **Caption Field**- This required field type supplies the category caption for the x-axis. Typically, a date field becomes the caption because line chart data is best represented over a time line. When you configure a date field as a caption, the field configuration automatically sets **Group By** to **Yes**.
- **Line Field** - A numeric field can be configured as a line. One or more line fields can be included in a line chart report template. Select **Line** for the field type. In this multiple objects per chart scenario, a line represents a specific object. Note that line colors do not apply to this “multiple objects per chart” scenario. Line colors will be determined dynamically by the system, since the number of objects represented by the lines is unknown until run time.
- **Column Break** - Use this setting to define how many points are represented by the line. For Line Chart report templates, if the template includes a field of type, **Group By Object**, one field must have **Column Break** set to **Yes**. This field should have unique values. For example, LUN names typically are duplicated in a storage environment, but the LUN ID uniquely identifies a LUN. In this example, a LUN ID should be included in a template as a Hidden field with a Column Break. Only one field in a template can have the Column Break set to Yes.
- **Group by Date** - If a date field is configured as a caption, the Group by Date setting is automatically set. However, if more than one date field is included in the template, you must select which date field should have the **Group by Date** setting set to **Yes**.

Format Line Chart Fields

When formatting a line field, note that the Unit Converter only formats the mouse-over text, not the label for the Y-axis.

See “Valid Formatter Patterns” on page 248.

Line Chart Scope Selections

- If a **Time Period** is selected for the scope selector, **Group By** must also be checked if the lines contain an aggregation function.
  See “Aggregation Functions” on page 267.
- If the line chart report template includes a field of type, Group by Object, the scope selector must have the **Top/Bottom** selection checked.

- If the line chart report template includes a field of type, Group by Object, the scope selector must have the **Group Chart By** selection checked.

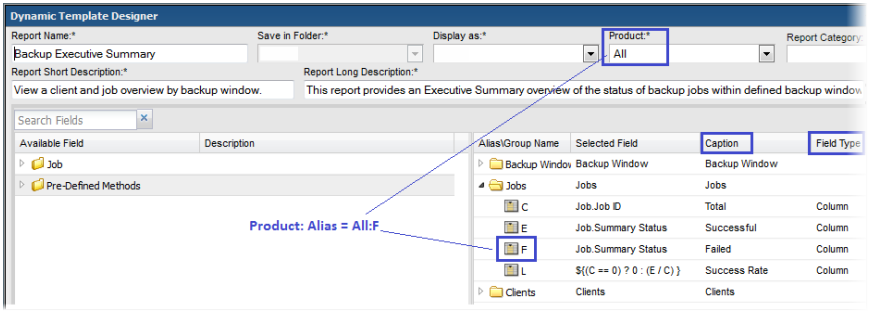
Examples of Line Chart Dynamic Templates

The best way to learn how to define a line chart report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A



Report Template	Product: Alias	Caption	Field Type	Functions	Group
Reports					
EMC Isilon File System Performance by Protocol	EMC Isilon: H	Max Bytes Out	Line	MAX	Yes
	EMC Isilon: A	# Active Clients	Line	AVG	Yes

Report Template	Product: Alias	Caption	Field Type	Functions	Group
NetApp Cluster-Mode NFS Performance by Vserver	NetApp Cluster-Mode: B	Cluster Name	Group By Field		Yes
	NetApp Cluster-Mode: D	Avg Write Latency	Line	AVG	Yes
NetApp Cluster-Mode Performance by Node	NetApp Cluster-Mode: H	Log Date	Caption		Yes
	NetApp Cluster-Mode: N	Disk Read	Line	AVG	Yes

Configure a Pie Chart Dynamic Template

Pie charts require both a caption field and a numeric sector field. To create pie sectors for a non-numeric field, such as OS name, you simply can configure a Count function to render the count of the instances of the value.

Note: Tip: As a best practice, choose a pie chart that is provided as an out-of-the-box report template, then copy and customize it to see how it is configured.

---

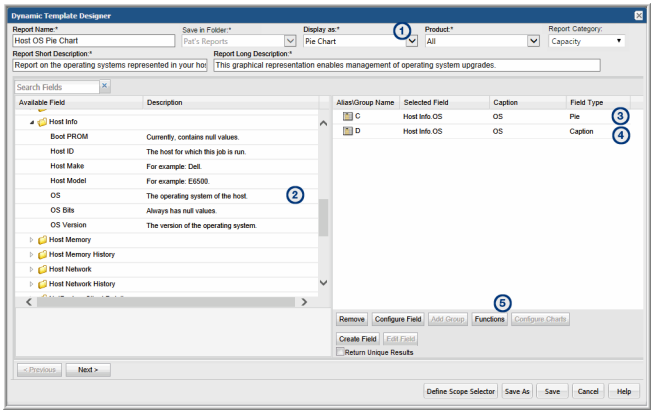
1. In the Dynamic Template Designer window, select **Pie Chart** from the **Display as** list.
2. Drag or double-click fields to add them to the pane at the right of the window. Typically, you will need only a caption field and a sector field.
3. **Sector Field**: Double-click the field to configure it as a pie sector.
4. **Caption Field**: Double-click the field to configure it as a caption.
5. Select the **Pie** field and click **Functions**.
6. In the Function Builder window, click **Add** and select a **COUNT** function to count the instances of the value to return a decimal that can be graphed as a pie sector.

7. Be sure to complete the other mandatory fields for the template: Report Name, Short Description, and Long Description.
8. Click **Define Scope Selector** to configure the **Maximum legends for Pie Chart**. Overwrite the default value with the maximum that you want to see.
9. Save the Report Template.

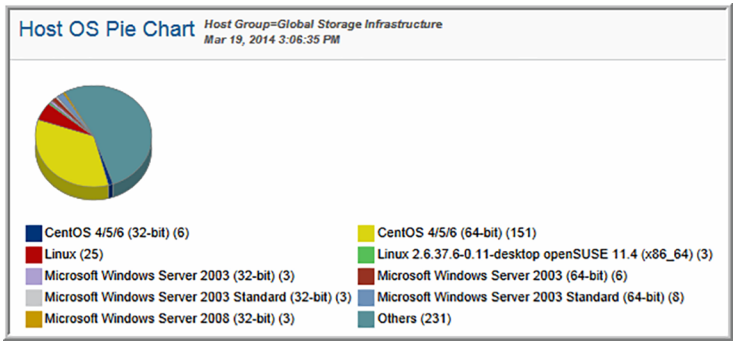
The following example shows the steps to create a pie chart for the operating systems represented by the hosts in your environment. The steps noted in the following illustration reflect the main steps described in the following.

See "Configure a Pie Chart Dynamic Template" on page 324.

Note that the HostInfo.OS field is used twice in this template, one for the Pie Sector and one for the Caption. Also, for the Sector field type, a COUNT function must be applied.



The Report Template defined in this example results in the following pie chart.



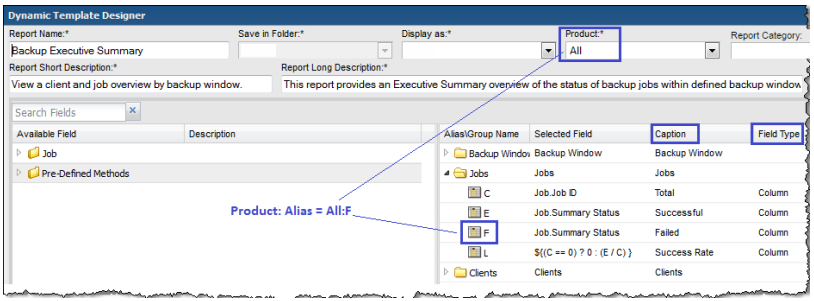
Examples of Pie Chart Dynamic Templates

The best way to learn how to define a pie chart report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A



Report Template	Report Type	Product: Alias	Caption	Field Type
Report				
Largest Backup Volume	Pie	All: C	Job size	Sector

Configure a Stacked Bar Chart in a Dynamic Template

A stacked bar chart is similar to a bar chart. The only thing that differs is the values are stacked in one bar, rather than represented as separate bars. To build a stacked bar chart, refer to the steps in the following sections and simply choose a display type of Stacked Bar Chart, instead of Bar Chart.

- See “Configure a Bar Chart Dynamic Template” on page 295.

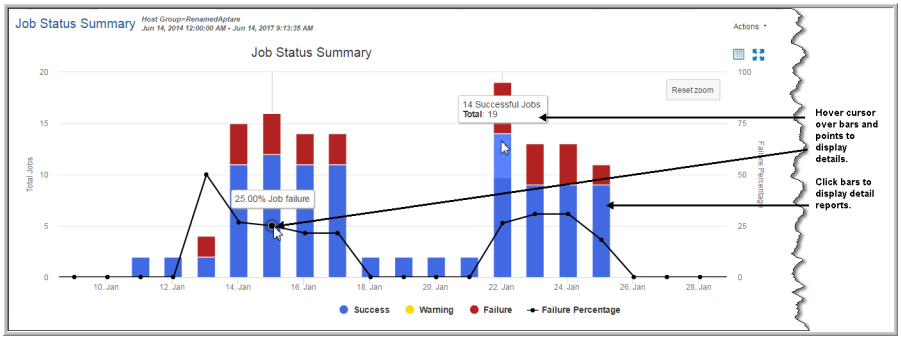
- See “Example of a Horizontal Stacked Bar Chart Dynamic Template” on page 312.
- See “Configure a Horizontal Bar Chart Dynamic Template” on page 310.

Example of a Stacked Bar Chart Dynamic Template

The stacked bar chart is best represented in an out-of-the-box report, Job Status Summary.

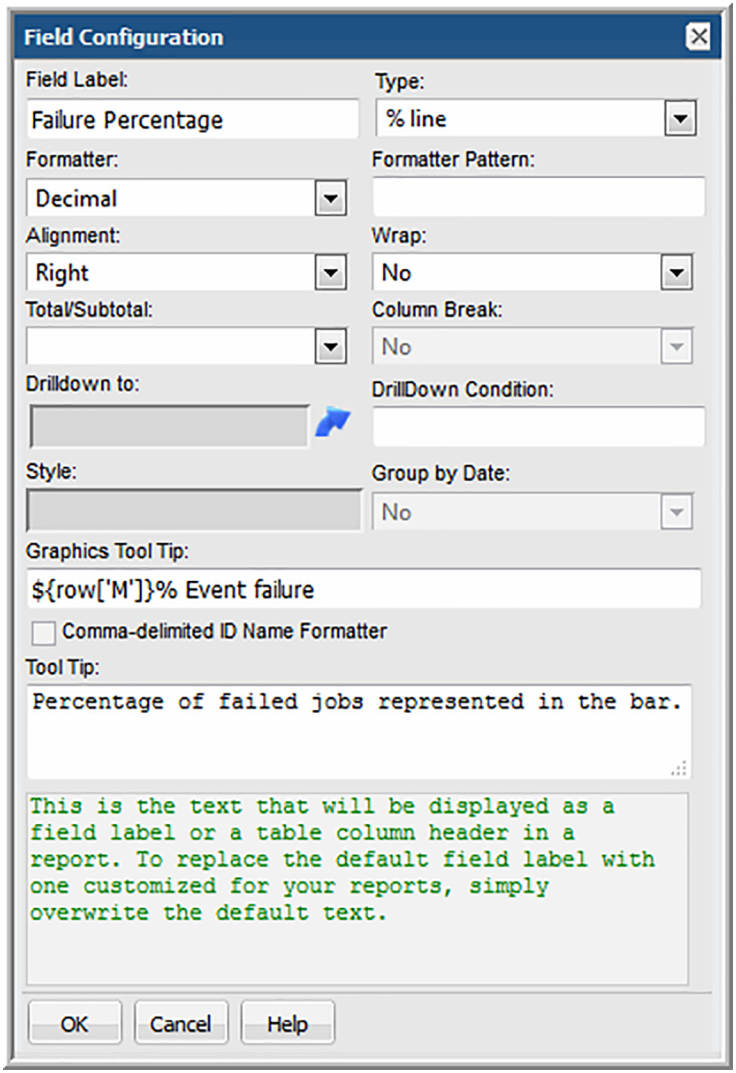
Refer to examples of other bar charts.

- See “Examples of Bar Chart Dynamic Templates” on page 302.
- See “Example of a Horizontal Bar Chart Dynamic Template” on page 310.
- See “Example of a Horizontal Stacked Bar Chart Dynamic Template” on page 312.



Configure a Line in a Bar Chart

Use the following example of how the line in the Job Status Summary bar chart was configured.



Configure a Table Dynamic Template

To get started with examples of out-of-the-box templates refer to the following.

See “Examples of Tabular Dynamic Templates” on page 329.

Examples of Tabular Dynamic Templates

The best way to learn how to define a bar chart report template is to learn by example.

The following list includes a sampling of the out-of-the-box reports that are shipped with the product, along with a few examples of the relevant information to help you isolate an example of the functionality that you are trying to implement.

To help you identify a field in the template that contains this functionality, the following convention is used in the list of report templates.

- **Product:Alias** combines the values listed in the template. In the following example, the selected **Product** is **All**, while the selected **Alias** is **F**.
- Some templates are vendor product-specific, in which case you may have a Product:Alias such as EMC Avamar:A

Report Template	Report Type	Product: Alias	Caption	Field Type	Functions
Reports					
Backup Executive Summary	Table	All: E	Successful	Column	DECODE, SUM
Job Summary	Table	All: L	Start Date	Column	
Monthly Backup Summary	Table	All: E	Last Full Backup	Column	
Reports					
Array Executive Summary	Table	All: C	Usable Internal Capacity	Hidden	SUM, NVL
		All: K	Replica	Column	SUM, NVL
Array Port Utilization	Table	All: L	Port Nbr	Hidden	TO_CHAR
Array Utilization Summary	Table	All: H	Used Capacity	Column	NVL
LUN Utilization Summary	Table	All: F	Device Nbr	Hidden	TO_CHAR, LPAD

Create a Sparkline Chart in a Tabular Dynamic Template

A sparkline presentation of the data can be added as a column to a tabular report to plot a series of values. Sparklines enable an at-a-glance view of data spikes that may indicate an issue that requires attention.

The basic requirements for a sparkline chart in a Dynamic Template include:

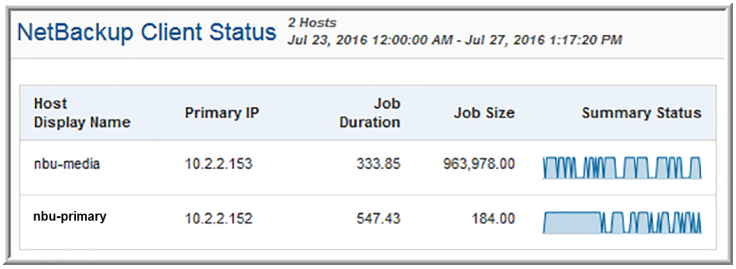
- **Table template type**. Sparklines are embedded in a table cell.
- **Data over time**. Scope Selector should have a Time Period component if you want to see values over time.
- **Comma-separated list of string values**. Use the aptStringConcat function to achieve this.
- **Field Configuration Formatter**: Select either Area Sparkline, Column Sparkline or Line Sparkline.
- **Method**. A Method may provide an efficient way of including a sub-query (WITH clause) to supply a series of values to be rendered by the sparkline.

Example of a Client Backup Sparkline Chart in a Dynamic Template

The following example provides a sparkline chart to represent backup job failures over time.

To achieve this report's backup client status results, the following settings were configured:

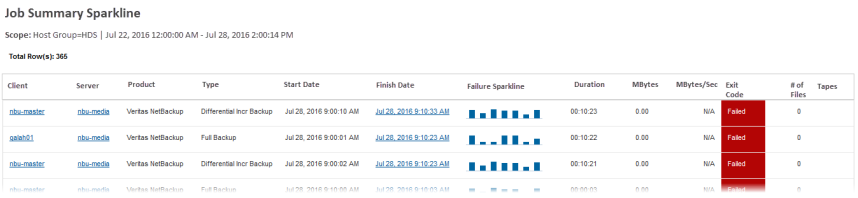
- The client's host name and IP address are included to retrieve data from the database.
- Job Duration contains a **SUM** function.
- Job Size contains a **SUM** function.
- Summary Status contains an aptStringConcat function.
- Summary Status has the Area Sparkline formatter selected in the Field Configuration.
- Time Period is selected in the Scope Selector.


NetBackup Client Status
2 Hosts
Jul 23, 2016 12:00:00 AM - Jul 27, 2016 1:17:20 PM

Host Display Name	Primary IP	Job Duration	Job Size	Summary Status
nbu-media	10.2.2.153	333.85	963,978.00	
nbu-primary	10.2.2.152	547.43	184.00	

Example of Using a Method to Create a Sparkline Chart in a Dynamic Template

The following example illustrates how to create a method to supply re-usable code that can be inserted into a column of a tabular report.


Job Summary Sparkline

Scope: Host Group=HDS | Jul 22, 2016 12:00:00 AM - Jul 28, 2016 2:00:14 PM

Total Row(s): 365

Client	Server	Product	Type	Start Date	Finish Date	Failure Sparkline	Duration	MBytes	MBytes/Sec	Exit Code	# of Files	Tapes
nbu-master	nbu-media	Veritas NetBackup	Differential Incr Backup	Jul 28, 2016 9:00:10 AM	Jul 28, 2016 9:10:33 AM		00:10:23	0.00	N/A	Failed	0	
galah01	nbu-media	Veritas NetBackup	Full Backup	Jul 28, 2016 9:00:01 AM	Jul 28, 2016 9:10:23 AM		00:10:22	0.00	N/A	Failed	0	
nbu-master	nbu-media	Veritas NetBackup	Differential Incr Backup	Jul 28, 2016 9:00:02 AM	Jul 28, 2016 9:10:23 AM		00:10:21	0.00	N/A	Failed	0	

To add a sparkline chart to the Job Summary Dynamic Template, take the following steps.

1. Create a client failure sparkline method.
2. Customize the Job Summary Dynamic Template.
3. From the Pre-Defined Methods folder, open the Client Failure Sparkline folder.
4. Drag the Failure Sparkline field into the report template's list of fields.
5. Configure the Failure Sparkline field and select a Column Sparkline formatter.
6. Click **Save As** and supply your own report name to be saved into a specified folder.

Client Failure Sparkline Method

To enable a sparkline to be dropped into a backup job report, a user-defined method for a Job enterprise object needs to be created with the following elements.

See "Overview of Method Creation" on page 334.

- With Clause

```
with spark as (
select trunc(start_date), client_id, client_name, count(job_id)
failed_count
  FROM apt_v_job
 WHERE client_id in(${hosts})
   AND start_date between ${startDate} AND ${endDate}
   AND summary_status = 2
  group by client_id, client_name, trunc(start_date)
)
```

- Query

```
select client_id, client_name, aptStringConcat(failed_count)
failed_count
  from spark s
 group by client_name, client_id
```

- Join

```
apt_job.client_id = ClientFailureSparkline.client_id (+)
```

- Fields

```
CLIENT_NAME - Check this field so that it can be included in a
Dynamic Template.
FAILED_COUNT - Check this field so that it can be included in a
Dynamic Template.
CLIENT_ID - Do not check this field as it is only needed for the
method's code.
```

Configure Chart Axes

For field types that have a field type of Bar or Line, additional elements need to be configured.

See "Steps to Create a Bar Chart Dynamic Template" on page 296.

Note: Tip: As a best practice, choose a bar chart that is provided as an out-of-the-box report template, then copy and customize it to see how it is configured.

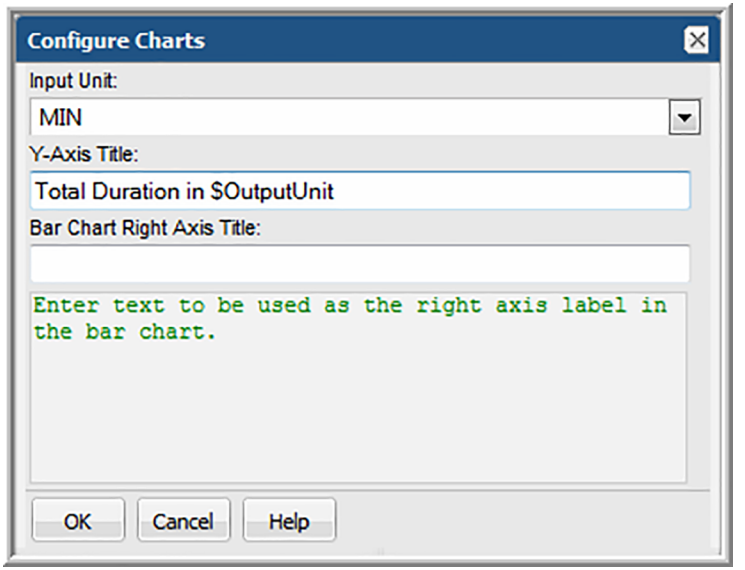
---

1. In the Dynamic Template Designer window, once you have configured Selected Fields that have a Field Type of Bar, click **Configure Charts** at the bottom right of the window.

   The Configure Charts option is not available under the following conditions:

   - If the report type is a Table.
   - If the report type is a Line Chart and has one or more groups defined in the template.

2. In the Configure Charts window, select an Input Unit that is relevant for the field values: KB, MB, GB, TB, Sec, Min, Hr, Day.



3. Enter text to be used as the Y-axis label in the bar chart. To append the output unit of measure, use **$OutputUnit**.
4. Enter text to be used as the right axis label in the bar chart.
5. To configure a bar chart legend refer to the following.

   See "Configure a Bar Chart Dynamic Template" on page 295.

Using Methods in Dynamic Templates

Methods enable special processing to be incorporated into a Dynamic Template. This mechanism enables modular programming, where processing, such as a sub-query, can be defined once and then used in many Dynamic Templates. A Method can take parameters, perform an action, such as a complex calculation,

and then return a value. For example, for the job enterprise object, a method can be used to derive the details about the last full backup: client ID, finish date, and job ID.

Methods have the following advantages:

- Modular: Create small, modular pieces of code to query specific data. Combine several methods in one Dynamic Template.
- Re-usable: Create SQL code that can be used in many Dynamic Templates. Simply drag-and-drop methods into a Dynamic Template.
- Easier to Maintain: Use similar logic in multiple templates. Modifications to a method then get applied to all relevant templates.
- More Filter Parameters: More tokens, selected at run time, can be passed into methods (compared to SQL templates). Example: client list.

This advanced feature requires experience in SQL (Structured Query Language) query development, similar to queries written in the SQL Template Designer. Methods can be used only in report templates that have been created using the Dynamic Template Designer.

A set of pre-defined system methods are provided to be used to incorporate complex processing into report templates. Administrators with SQL knowledge can create additional methods. When an administrator creates a method, it is available only to users within the same Domain.

For a complete, current list of pre-defined methods available in the Portal, go to

Admin>Reports>Method Designer.

Overview of Method Creation

Action	Description
See "Viewing Methods" on page 341.	View the list of system and user-defined methods that have already been created. Only user-defined methods can be edited.
See "Modify Reports" on page 148.	Several components within a method require configuration: See "Accessing Methods" on page 343.

System Methods

The following methods are packaged with the system, for use in the Dynamic Template Designer.

System Method	Vendor Product	Description
Methods: Job Enterprise Object		
Backup Window	All Supported Backup Products	Returns the start and finish dates in a backup window.
Client Last Backup Job Info	All Supported Backup Products	Returns the finish date and the corresponding job IDs of the most recent full or incremental backup jobs that are in a successful or partial state. The data returned is per client. This is a stand-alone function and cannot be used in combination with other objects or functions.
CommVault Backup Copy Info	CommVault Simpana	Returns the minimum Backup Copy Identifier for a given CommVault Simpana job.
Commvault Job Attempts	CommVault Simpana	Returns Simpana backup jobs and try count for a client list for CommVault Simpana jobs.
CommVault Running and Queued Job Count	CommVault Simpana	Returns the count of CommVault Simpana backup jobs that are in the running state in the last 90 days. The counts returned are per server.
Data Domain Last Backup Job Info	Data Domain	Returns the details for Data Domain backup jobs. The data returned is per client and logical storage unit.
Disk and Tape Usage Summary	All Supported Backup Products	For the most recent successful or partial job, this method returns the finish date, corresponding job ID, number of kilobytes saved in tape media, number of kilobytes saved in disk media, number of files saved in tape media and number of files saved in disk media. The data returned is per server, client, product type and job type.
Failed Backup Job Detail	All Supported Backup Products	For the failed backup jobs within the selected, required time period, this method returns the following data on a per-client basis: client ID, along with its corresponding job ID, the count of failed backups, and the date of the last failed backup job.

System Method	Vendor Product	Description
Failed Full Backup Job Detail	All Supported Backup Products	For the failed full backup jobs within the selected, required time period, this method returns the following data on a per-client basis: client ID, along with its corresponding job ID, the count of failed full backups and the last failed full backup date.
HP DP Session Tape Media List	HP Data Protector	For any HP Data Protector (HP DP) session, this method returns a comma-separated list of the tape media associated with the job. The data will be returned in the following format: id1,name1|id2|name2. The IDs can be used for report drilldowns.
Job Tape Media List	All Supported Backup Products	For any successful or partial job, this method returns a comma-separated list of the tape media associated with the job. This method returns the tape media list if the job's status is success or warning.
NetBackup Job Error Msg	Veritas NetBackup	Relevant only for NetBackup. It returns the error message if either the job or its parent job is in an error state.
NetBackup Job Tape Media List	Veritas NetBackup	Relevant only for NetBackup. For any successful or partial NetBackup job, this method returns a comma-separated list of the tape media associated with the job. This method returns the tape media list if the job or its parent job's status is success or warning.
NetBackup Running and Queued Job Count	Veritas NetBackup	Relevant only for NetBackup. The method returns the number of jobs that are in the running state and the number of jobs in the queued state. The counts returned are per backup server.
NetBackup Running Job Info	Veritas NetBackup	Relevant only for NetBackup. The method returns a value of 1 if a NetBackup job is long-running, slow-running, or stalled, based on the policies that are configured for a host group.

System Method	Vendor Product	Description
NetWorker Running Job Count	EMC NetWorker	Relevant only for EMC NetWorker. The method returns the number of backup jobs in a running state. The data returned is per NetWorker server.
NetWorker Running Job Info	EMC NetWorker	Relevant only for EMC NetWorker. The method returns a value of 1 if a NetWorker job is long-running, slow-running, or stalled, based on the policies that are configured for a host group.
Restore Successful Job Detail	All Supported Backup Products	For successful restore jobs within the selected, required time period, this method returns the following data on a per-client basis: client ID, along with its corresponding job ID, the count of successful restore jobs, the number of files restored, and the last successful restore date.
Successful Backup Job Detail	All Supported Backup Products	For the successful full and partial backup jobs within the selected, required time period, this method returns the following data on a per-client basis: client ID along with its corresponding job ID, the count of the number of successful backups, the last successful backup date, amount of data (in MB) saved to tape media, and the amount of data (in MB) saved to disk media.
Successful Full Backup Job Detail	All Supported Backup Products	For the successful full backup jobs within the selected, required time period, this method returns the following data on a per-client basis: client ID along with its corresponding job ID, the count of the number of successful full backups, the last successful full backup date, amount of data (in MB) saved to tape media, and the amount of data (in MB) saved to disk media.
Methods: Data Domain Enterprise Object		
Data Domain Filesystem Capacity	Data Domain	This method returns total and used filesystem capacities at the Data Domain system level. Requirement: When using this method in a template, the Scope Selector must have the Filesystem Names component selected.

System Method	Vendor Product	Description
Data Domain Raw Capacity	Data Domain	Returns raw capacities at the Data Domain system level. Use this method to include the total raw capacity in a report template.
Data Domain Tape Capacity	Data Domain	Returns tape capacities at the Data Domain system level. Use this method in a report template to display the total tape capacity, as well as the total used tape capacity.
Data Domain VTL Count	Data Domain	Returns the number of virtual tape libraries (VTLs) at the Data Domain system level. Use this method to include a VTL count in a Data Domain report template.
Methods: Array Enterprise Object		
Array Site Disk Info	IBM	Relevant only for IBM Arrays. For each Array Site, the method returns aggregated information from the physical disk.
Data Mover CPU and Memory Usage	EMC VNX (Celerra)	Returns data only for EMC VNX (Celerra) arrays: Peak CPU Usage Percentage and Peak Memory Usage Percentage for the Movers for the specified arrays.
LUN Array Group List	All Arrays	Returns a comma-separated Array Group list for each logical unit.
LUN Array Port WWN List	All Arrays	Returns a comma-separated Array Port WWN list of all logical units.
Top 10 Tiers Chargeback Capacity	All Arrays	Lists the chargeback capacity per array in kilobytes and the policy name for the top 10 chargeback policies defined in the system.
Methods: Host Enterprise Object		
Analytics Host Summary	File Analytics	Lists a capacity summary for hosts from which file data has been collected. Use this method to get an overview of how devices and volumes are being used, including an indicator of how much is duplicated.
Methods: Host Enterprise Object		

System Method	Vendor Product	Description
VM CPU Performance Detail		This method returns details about VM CPU Performance, based on the time period or dates selected in the report's scope selector. The VM CPU Performance Summary enables a Storage Administrator to have an at-a-glance look at CPU I/O performance statistics.
VM Disk Performance Detail		This method returns details about VM Disk Performance, based on the time period or dates selected in the report's scope selector. The VM Disk Performance details enable a Storage Administrator to have an at-a-glance look at I/O performance statistics for disks.
VM Guest Detail		This method returns details about VM Guests. Use this report to find the largest Virtual Machines. This report lists virtual guests/hosts with usage and status information. NOTE: VMTools must be installed to enable collection of the IP address, Host name, mount points, and guest operating system of the VM.
VM Memory Performance Detail		This method returns details about VM Memory Performance, based on the time period or dates selected in the report's scope selector. The VM Memory Performance Summary enables a Storage Administrator to have an at-a-glance look at memory I/O performance statistics,
VM Server Detail		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.
Methods: Host Enterprise Object		
Host Zones and Hops Details		Returns the total number of zones and hops for each host within a fabric.

System Method	Vendor Product	Description
Switch and Fabric Name List Details		Returns a comma-separated list of the switch and fabric names associated with a host.
Methods: Host Enterprise Object		
Filesystem Summary		Use this method to list file system capacity and the most/least used file systems.
Hosts At Risk Details		Identifies hosts that require your immediate attention. Hosts at risk are determined by evaluating current and projected usage.
Host Utilization Details		Use this method to get an overview of host utilization to determine if there is contention for storage.
Over Provisioned Host Details		This method returns details about over-provisioned hosts, based on the time period or dates selected in the report's scope selector.

Stand-Alone Methods

- Some pre-defined methods are stand-alone methods--that is, they cannot be combined with fields or methods that are not in the stand-alone method. These stand-alone methods typically contain complex processing that falls outside of the scope of basic Dynamic Template Designer functionality. If you use fields from a stand-alone method in a Dynamic Template, no other methods or fields can be added to the template.
- A stand-alone method enables you to build a SQL query that can be used in a Dynamic Template; therefore, SQL knowledge is required to author a stand-alone method.
- Stand-alone methods are not associated with an enterprise object, but with specific products, such as backup products.
- To determine if a pre-defined method is a stand-alone method, view/edit the method and note that the JOIN clause is null.
- For Line Charts, if the template uses a stand-alone method, the **Top/Bottom** scope selector component is not supported.

Viewing Methods

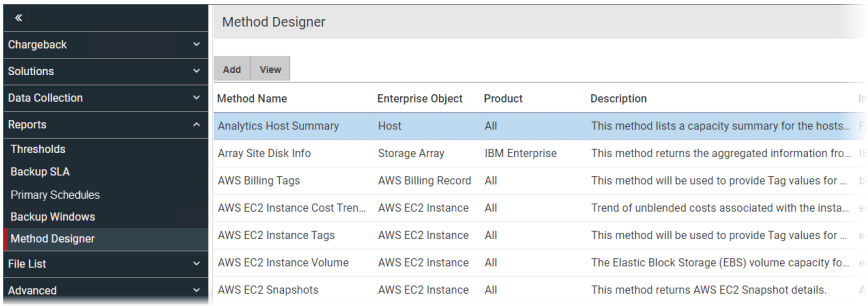
Methods, for use in the Dynamic Template Designer, are listed in the Methods View panel.

See “Using Methods in Dynamic Templates” on page 333.

To access methods, take the following steps:

1. In the Portal, select **Admin>Reports> Method Designer**

   The methods view panel lists existing methods, with details about each.



Method Name	This name is provided by the author of the method.
Product	The IT Analytics product, such as Capacity Manager or Backup Manager.
Subsystem	Some methods are specific to the type of data collected from a particular subsystem, such as a backup product; others are available for all systems for the product.
Description	Description of the value and use case for the method.
Internal Method Name	This unique, internal name is created by the Dynamic Template Designer. Use this name if you are creating or editing a method that has a query that needs a WITH or JOIN clause. During query validation, you may see this name listed in error messages.
Last Modified	The date the method was created.

2. Buttons at the top of the Methods View panel enable the following actions, depending on which method you select: **Add**, **Delete**, **Details**, and **Edit**. These same options are available when you right-click in the browser window. System Methods, provided as part of the software, cannot be modified or deleted.

Adding or Editing Methods

Methods that can be used in report templates are created with the Dynamic Template Designer. These methods are listed in the Methods View panel. This list includes both system methods and user-defined methods.

User-Defined Methods

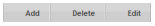
For an example of a user-defined method refer to the following.

See "Example of Using a Method to Create a Sparkline Chart in a Dynamic Template" on page 331.

Only user-defined methods can be edited or deleted. When you select a user-defined method in the Methods View panel, at the top of the Methods View pane, you'll see the following buttons that enable updates.

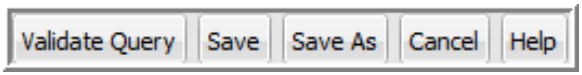
A method cannot be edited under the following circumstances:

- The fields in the method are in use in a report template.
- A table used in the Join clause is not relevant for the enterprise object.



Within the Method Designer window, the following buttons enable functionality that is available only for user-defined methods. Note that System Methods will have certain buttons disabled, as described in the following section.

See "System Methods" on page 343.

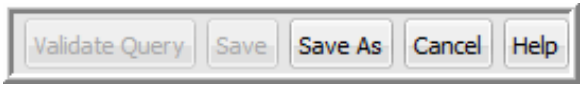


- **Validate Query** - Validates valid syntax only for the query and the With clause, if used.
- **Save** - Before saving, the query is validated.
  See "Validate and Save a Method" on page 348.
- **Save As** - Enables a copy and save function so that you can clone and modify methods.
- **Cancel** - Ignores any changes that have been made since the last save.
- **Help** - Accesses the context-sensitive help for methods, with details that can assist in method creation.

System Methods

System methods shipped with the product cannot be edited or deleted. When you select a system method in the Methods view panel, at the top of the Methods view pane, you'll see the following available buttons.

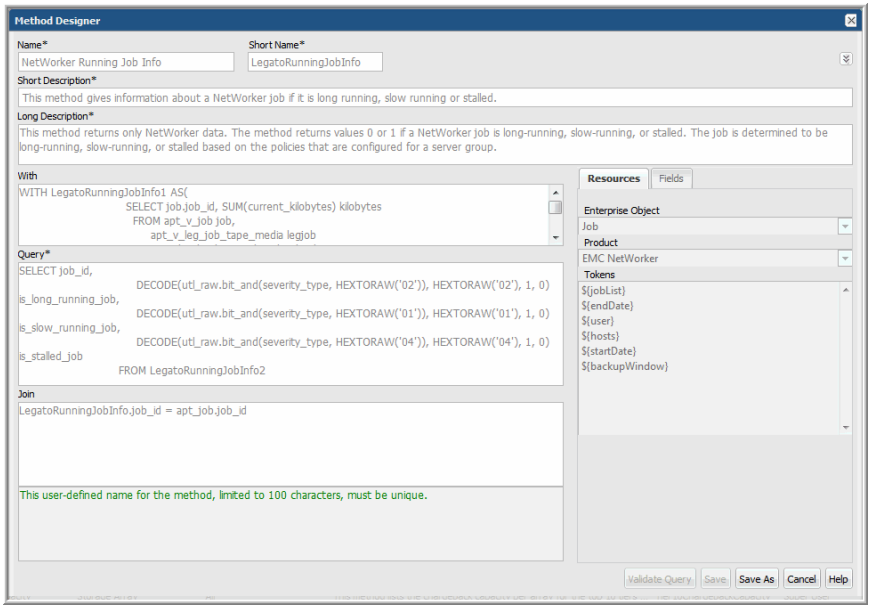
Within the Method Designer, the following buttons are available only for user-defined methods.



- **Validate Query** and **Save**are disabled for System Methods.
- **Save As** - Clone and modify system methods to take advantage of factory-shipped functionality.
- **Cancel** - Dismisses the dialog window.
- **Help** - Accesses the context-sensitive help for methods, with details that can assist in method creation.

Accessing Methods

1. In the Portal, select **Admin>Reports>Method Designer**.
2. To Add a Method, click **Add** at the top of the view panel to launch the Method Designer window.
3. To Edit a Method, double-click a method or select a method and click **Edit**.



4. In the Method Designer, define and edit the following method components.

Name	This user-defined name for the method, limited to 100 characters, must be unique.
Short Name	This user-defined, case-insensitive name must be unique, with the following constraints: ■ Limited to up to 23 characters. ■ Cannot begin with a number. ■ Cannot contain special characters or spaces. ■ Cannot contain Oracle reserved words. For a list of reserved words, see the Oracle documentation. Use this name when defining a WITH or JOIN clause. During query validation, you may see this name listed in error messages.
Short Description	This description, limited to 200 characters, is displayed in the methods view panel.
Long Description	The long description is limited to 4,000 characters. Use this description to provide essential details, such as prerequisites, which will aid anyone wanting to customize the method.

With Clause

Use the SQL WITH clause to assign a name to a subquery block so that it can be easily referenced in the query. The WITH clause must follow these rules:

- "/>WITH clause aliases must be unique across all the functions used in a report template; therefore, create an alias in the following format:
  <function_name><number>
  Example: BackupWindow1
- "/>-->Cannot contain Oracle reserved words that can alter data in the database; for example, DELETE, UPDATE, or INSERT.

Query

Define the SQL query with SELECT statements to return data from the database.

- Cannot contain Oracle reserved words that can alter data in the database; for example, DELETE, UPDATE, or INSERT.
- All columns in the SELECT clause must be aliased; for example,

```
SELECT decode(client_id,1,'121212',2,'2323',5) A
FROM apt_v_job
```

- The function,collectString, can be used in the Method Designer to concatenate distinct values. The Oracle function, aptStringConcat with DISTINCT or UNIQUE, cannot be used to concatenate values in a method, even though the method will validate and save. When that method is used in a report template, it will fail.

See "collectString" on page 425.

Join Clause

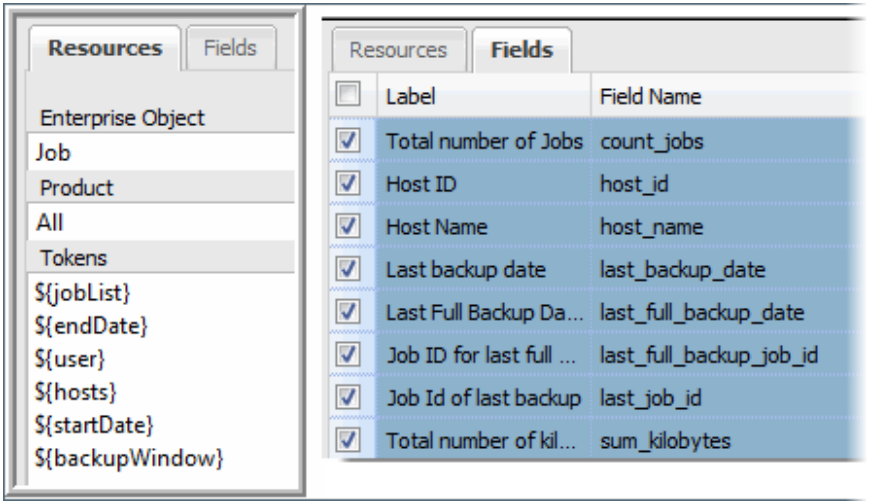
Use the SQL JOIN clause to combine results from the Method with the results of the template. The JOIN clause must follow these rules:

- The JOIN conditions should use relevant enterprise object tables only. See the Dynamic Template Designer to view the table associated with a field.
- Only the equal sign (=) is supported.
- The method's field should always be defined on the left side of the equal sign (=) .
- SQL functions, such as UPPER, DECODE, or LOWER cannot be used in a JOIN clause.
- Certain database tables cannot be used in JOIN clauses.
  See "Exception Tables for a Method's Join Clause" on page 347.

5. Click **Validate Query** at the bottom of the Method Designer window to validate the query.

   Validate Query only validates the With and Query components. If an error exists, both components will be bordered in red, indicating that the validation has failed. In addition, after validation, the Fields pane will be populated with the fields (database columns) that are used in the query.

6. At the right of the Method Designer window, the **Fields** and **Resources** buttons enable the following configurations and capabilities.



Fields

When adding a Method, the Fields pane will be populated only after the query has been successfully validated.

Check the boxes to select the fields that will be available in the Dynamic Template Designer so that the method can be included in a report template.

Several entries are required for each field before the method can be successfully saved.

- Label: User-supplied. This label is what will appear in the Report Template.
- Field Name: Pre-defined by the Method Designer.
- Short Description: User-supplied.
- Long Description: User-supplied.

Resources

- Enterprise Object: Select an enterprise object. This selection makes a method available to only templates created for that object. It also drives the list of tokens that are available for use in a query to enable variable substitution at report run time.
- Product: Select a specific vendor product (subsystem) for the enterprise object; for example, for a storage array enterprise object, EMC VNX (Celerra).
- Tokens: These tokens can be used in a query to enable variable substitution at run time. Place your cursor at the insertion point in the query and double-click the token.

  See "Tokens for Methods" on page 347.

  The list of relevant tokens is driven by the selected Product Group.

7. Click **Save**.

Tokens for Methods

To access the pre-defined tokens provided for certain product groups, take the following steps:

1. In the Portal, select **Admin>Reports>Method Designer**.
2. To Add a Method, click **Add** at the top of the view panel to launch the Method Designer window.
3. To Edit a Method, double-click a method or select a method and click **Edit** at the top of the view panel to launch the Method Designer window.
4. In the Method Designer, click **Resource** at the right of the window.
5. Select a **Product Group** to view the relevant tokens.

Product Group	Token Name
Backup Manager	${backupWindow}
Backup Manager	${endDate}
Backup Manager	${startDate}
Backup Manager	${jobList}
Backup Manager	${hosts}
Backup Manager	${user}
Capacity Manager	${domainList}
Capacity Manager	${endDate}
Capacity Manager	${startDate}
Capacity Manager	${arrayList}
Capacity Manager	${user}
Capacity Manager	${spArrayList}

Exception Tables for a Method's Join Clause

The following tables cannot be used in a JOIN clause of a method.

Enterprise Object	Table Name
Job	apt_server_instance
Job	apt_host
Job	apt_leg_volume_pool
Job	ptl_volume_pool
Job	apt_leg_cr_group
Job	ptl_nbu_storage_unit
Job	ptl_nbu_retention_level
Job	apt_tsm_stgpool_contents
Job	apt_attr_host
Job	apt_host_detail
Job	apt_host_info
Array	aps_emc_sym_storage_pool
Array	aps_vnx_storage_pool

Validate and Save a Method

Once you have configured all the mandatory elements of a method, you can save the method.

- **Save** - When saving, the method is validated. This button is disabled for System Methods.
- **Save As** - “Save As” also validates the method. This option enables a copy-and-save capability so that you can clone and modify methods. This is especially useful for tailoring System Methods to take advantage of factory-shipped, advanced functionality.

Method Validation

When a method is saved, the following checks are made.

- All mandatory elements have values: Name, Short Name, Short Description, Long Description, and Query
- Short Name is unique

- Query validates
- At least one field is checked
- All checked fields have been filled in: Label, Short Description, Long Description

If any of the above validations have issues, the relevant method pane(s) will be flagged with a red border. When possible, a red information icon will display, which you can click for additional information.

Delete a Method

The product ships with system methods that cannot be deleted. Only user-defined methods can be deleted. And, a method can be deleted only if it is not being used by a report template.

To delete a method, take the following steps:

1. In the Portal, select **Admin>Reports> Method Designer**.

   The methods view panel lists existing methods.

2. Click on a method to select it.
3. Click the **Delete** button at the top of the methods view pane.

   If the method is currently being used by a report template, a pop-up window will notify you and prevent you from deleting the method.

Troubleshooting Dynamic Templates

If an error occurs while generating a Dynamic Template, use the **Ctrl + Alt + D** key sequence within the browser and rerun the template. This sequence lets you turn on database debugging, which will write error messages to scon.log, enabling you to email those messages to Cohesity Support.

Chapter 9

Work with the SQL template designer

This chapter includes the following topics:

- SQL template designer overview
- Database Published Views
- Steps to Create SQL Templates
- Create a SQL Template
- Use Functions in Queries
- Format the SQL Template Output
- Configure a Bar Chart SQL Template
- Example of a Bar Chart SQL Template: Host CPU Performance
- Configure a Line Chart SQL Template
- Configure a Donut Chart SQL Template
- Example of a Donut Chart SQL Template: NetBackup Job Size by Job Type
- Configure a Gauge Chart SQL Template
- Configure a Horizontal Bar Chart SQL Template
- Example of a Horizontal Bar Chart SQL Template: Host CPU Performance
- Configure a Horizontal Stacked Bar Chart SQL Template
- Example of a Horizontal Stacked Bar Chart SQL Template: Available/Allocated Capacity

- Configure an Area Chart SQL Template
- Example of Area Chart SQL Template: Allocated Available
- Configure a Pie Chart SQL Template
- Example of a Pie Chart SQL Template: Overall Job Status Summary
- Configure a Pivot Table SQL Template
- Example of a Pivot Table SQL Template: NetBackup Job Size
- Configure a Stacked Bar Chart SQL Template
- Example of a Stacked Bar Chart SQL Template: Host CPU Performance
- Configure a Table in the SQL Template Designer
- Save and Share Report Templates
- Advanced SQL Report Template Options
- Export/Import SQL Templates

SQL template designer overview

Note: The SQL Template Designer is considered to be an advanced feature, requiring experience in SQL (Structured Query Language) query development.

The following reporting capabilities are provided:

- Out-of-the-box Report Templates: Report templates packaged with your installation.
- Dynamic Template Designer: A tool to create custom reports by dragging and dropping database components, which are then used to dynamically generate the SQL query “behind the scenes” of the Portal.
  See “Dynamic Template Designer Overview” on page 220.
- SQL Template Designer: A tool to enable custom report creation using SQL skills. Query a large set of database published views to develop reports.
- Customer Report Library: A collection of custom reports that you can use as a starting point, importing them into the Portal: http://reportlibrary.aptare.com/

The SQL Template Designer offers a way for you to augment your report inventory. Once you have become familiar with the report templates that are shipped with the Portal, you may want to design reports that draw on the rich set of data gathered

by the Data Collectors, to serve your unique reporting needs. Using the SQL Template Designer and your knowledge of SQL, you can build advanced reporting solutions to support efficient storage resource management.

See “Steps to Create SQL Templates” on page 352.

Database Published Views

The product provides a set of published read-only database views into the reporting database and related object model. The database views provide a native SQL interface into the reporting database. IT Analytics publishes a collection of views that overlay the underlying database tables. Using these read-only views, you can write your own reports or stored procedure handlers to query the reporting database. These views provide a simple and fast mechanism to access a read-only view into the raw data maintained within the database.

Access the Database Published Views from the CLI

The underlying Oracle database can be accessed via the command-line interface (CLI). The product includes a separate Oracle user account that is specifically designed and granted SELECT access to the Published Database Views.

The Oracle user account for the Published Database Views is called **aptare_ro**. The default password for this account is set to **aptaresoftware123**. This user account can ONLY access the published database views and does not have any update or insert privileges.

Access the Database Published Views from the SQL Template Designer

You can quickly access the available database views, fields and corresponding descriptions for query creation directly from the SQL Template Designer.

See “Help for Database Views” on page 368.

Steps to Create SQL Templates

Use the SQL Template Designer to create custom reports by developing SQL (Structured Query Language) queries to tap into the database.

See “SQL template designer overview” on page 351.

Also, visit the Report Library for additional report template examples:

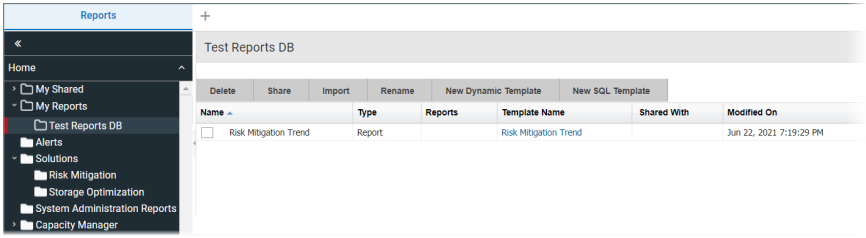
https://reportlibrary.veritas.com

This section includes details for the following SQL Template Designer options and features:

- See “Create a SQL Template” on page 353.
- See “Configure SQL Template Scope Selector Components” on page 356.
- See “Construct the SQL Query” on page 367.
- See “Help for Database Views” on page 368.
- See “Accessing the Published Database Views from the CLI” on page 369.
- See “Special Characters in SQL Queries” on page 369.
- See “Sample SQL Queries” on page 370.
- See “Use Functions in Queries” on page 375.
- See “Format the SQL Template Output” on page 376.
- See “Use Functions in Queries” on page 375.
- See “Save and Share Report Templates” on page 407.
- See “Advanced SQL Report Template Options” on page 408.
- See “Create mouse-over hovers” on page 411.
- See “Export/Import SQL Templates” on page 418.

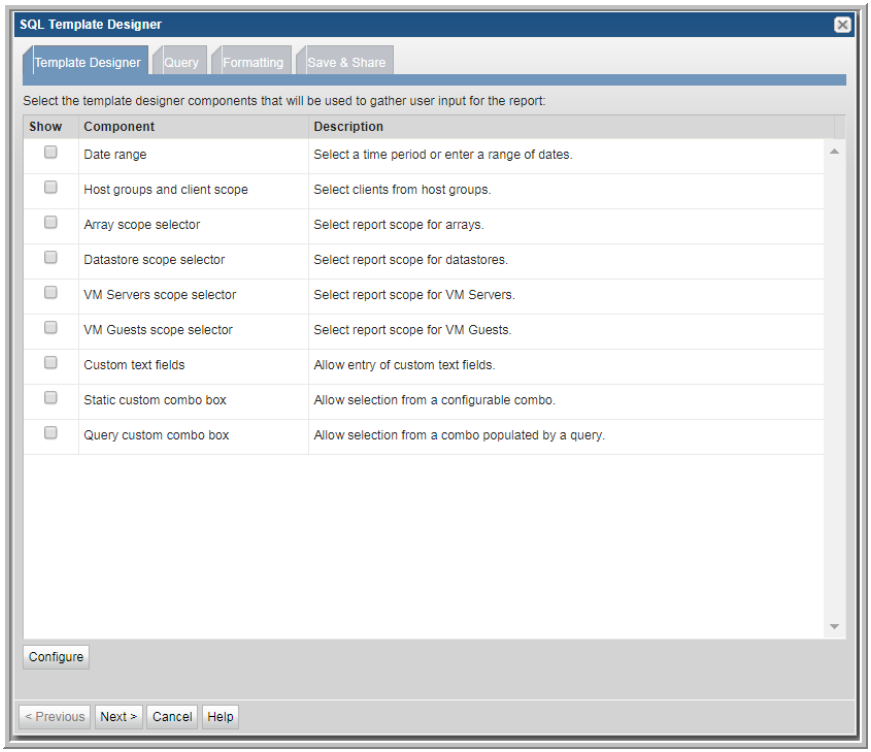
Create a SQL Template

1. Select **Reports > My Reports> Your Custom Name Folder**.
2. Click the **New SQL** Template button.



The initial display of the SQL Template Designer window displays:

- Tabs to navigate the configuration steps. Use either the tabs at the top of the window, or click **Next** at the bottom of the window.
- Designer components that can be configured for users to select at run-time. See “Configure SQL Template Scope Selector Components” on page 356.


SQL Template Designer
Template Designer
Query
Formatting
Save & Share
Select the template designer components that will be used to gather user input for the report:
Show
Component
Description
Date range
Select a time period or enter a range of dates.
Host groups and client scope
Select clients from host groups.
Array scope selector
Select report scope for arrays.
Datastore scope selector
Select report scope for datastores.
VM Servers scope selector
Select report scope for VM Servers.
VM Guests scope selector
Select report scope for VM Guests.
Custom text fields
Allow entry of custom text fields.
Static custom combo box
Allow selection from a configurable combo.
Query custom combo box
Allow selection from a combo populated by a query.
Configure
< Previous
Next >
Cancel
Help

Template Designer	Using the checkboxes, select which parameters you want the user to specify in the Scope Selector window when the report is generated. ■ See "Date Range" on page 356. ■ See "Host Groups and Client Scope" on page 358. ■ See "Array Scope Selector" on page 359. ■ See "Datastore Scope Selector" on page 359. ■ See "VM Servers Scope Selector" on page 359. ■ See "VM Guests Scope Selector" on page 359. ■ See "Custom Text Fields" on page 359. ■ See "Static Custom Combo Box" on page 362. ■ See "Query Custom Combo Box" on page 363. To customize any of the above options refer to the following. See "Configure SQL Template Scope Selector Components" on page 356. After you complete the selections within the tabbed Report Designer window, click **Next** to go to the window where you will construct the SQL query. See "Construct the SQL Query" on page 367.
Query	Type your SQL statement using the list of valid database views and fields. Click **Validate Query**to check your SQL statement for valid syntax before moving on to formatting and saving. This **Validate Query** task actually executes the query in order to validate the syntax. Seemingly simple queries could return a large amount of data and therefore may take some time to validate. See "Construct the SQL Query" on page 367.

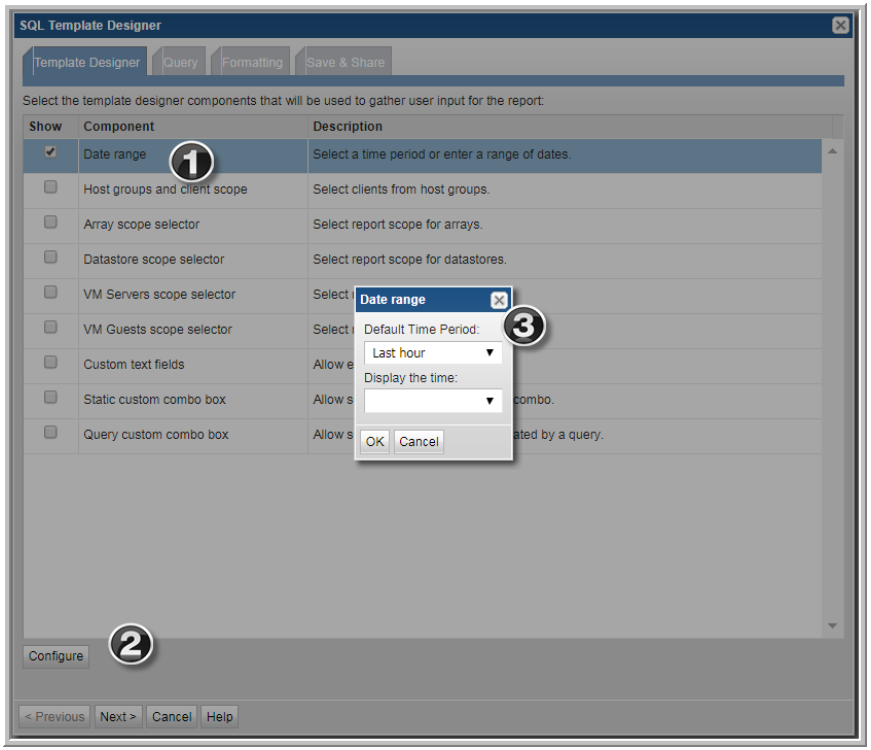
Formatting	To format the output for a report template, select the information that you want included in the report, as well as the report output style--Table, Bar Chart, Donut Chart, Gauge Chart, Area Chart, Stacked Area Chart, Horizontal Bar Chart, Pie Chart, Pivot Table, Stacked Bar Chart, Horizontal Stacked Bar Chart. See "Format the SQL Template Output" on page 376.
Save & Share	Save and share report templates with others. Assign the category for placement and set up the Inventory Report Configuration. See "Inventory Report Configuration" on page 416. You can also select individual users or groups to share with.

Configure SQL Template Scope Selector Components

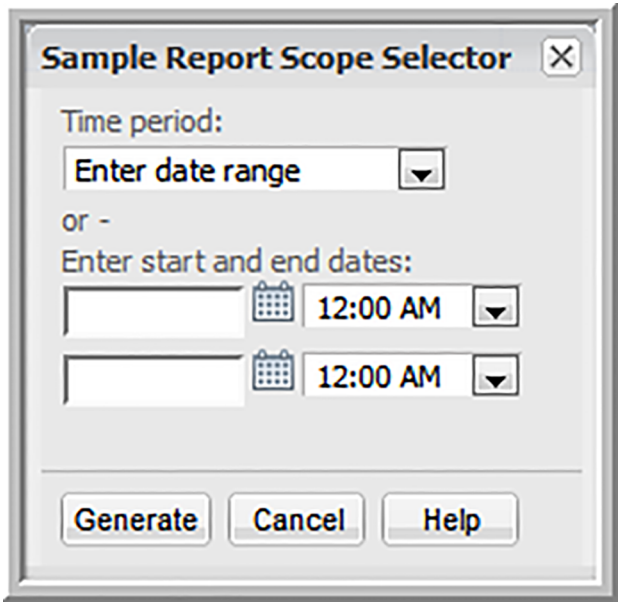
When you create a report template, you must also determine what elements should be provided in the report's Scope Selector--the interface that enables the user to specify the scope of the data to be displayed in the report. Each time a report is generated from the report template, the Scope Selector enables the selection of parameters to be used for report generation.

Date Range

Check this component to include a Time Period and Data Range selector in the report's Scope Selector that is displayed when generating the report.

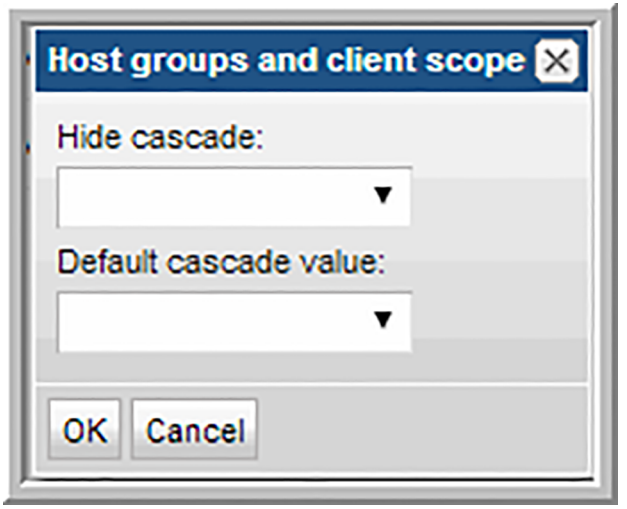


1. Check **Date range**.
2. Select the **Date range** row and click **Configure** to display the Date range window. The Date range window allows you to set a Default Time Period. This becomes the default selection in the scope selector.
3. Configure the Date component to include the time with the date range. Choose **Yes** or **No** to specify if the time will be shown along with the date range in the Scope Selector. This configuration results in a report Scope Selector that includes the following:



Host Groups and Client Scope

This SQL Template Designer component enables you to customize how the report will display the host groups' data.



Use this configuration to specify if the cascade to sub-groups option is available and, if available, the default setting--Checked or Unchecked.

Array Scope Selector

This SQL Template Designer component enables you to add a scope selector for reports with Arrays and groups of arrays.

Datastore Scope Selector

This SQL Template Designer component enables you to add a scope selector for reports with Datastores.

VM Servers Scope Selector

This SQL Template Designer component enables you to add a scope selector for reports with VM Servers.

VM Guests Scope Selector

This SQL Template Designer component enables you to add a scope selector for reports with VM Guests.

Custom Text Fields

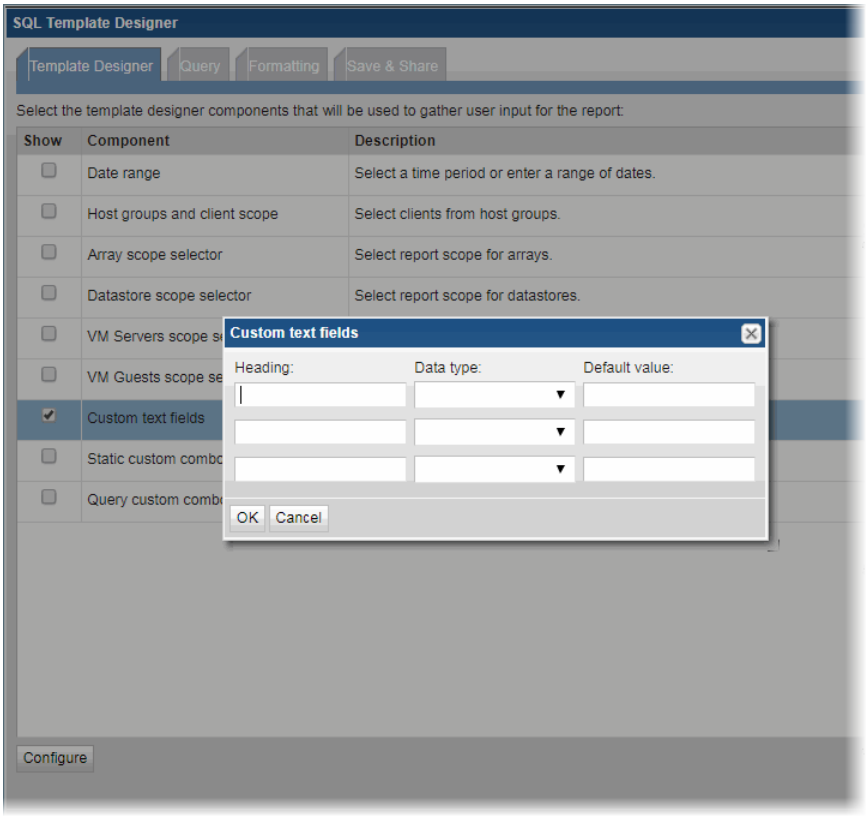
In the Template Designer pane, a user can define fields that can have values substituted into a query. When the report is generated, the Scope Selector window presents the fields that can be selected, thus providing dynamic input for the report.

This option enables you to create up to three free-form text fields, where you can specify a field that you want to use in your query, plus set a default value. For example, you might enter Host Name, so that when you form the query, you can specify a host name to be inserted into the query.

Example of Custom Text Field Configuration in a SQL Template

The following steps illustrate the advantage of Custom Text Fields.

1. In the SQL Template Designer, select the **Custom Text Fields** component and click **Configure**.



2. In the **Custom Text Fields** window, define a field named **Host ID** with a data type of **Number** and click **OK**.
3. To use this newly configured Custom Text Field, be sure to select the checkbox.
4. In the **Query** window, enter the following SQL query by typing a partial statement and then double-clicking to select fields:

   ```
   select * from aps_v_host_volume where host_id > ${freeText1}
   ```

   This query can be constructed by combining typing with double-clicking selections in the window, as shown in the following example. For a complete list of database views and columns, you can access the Database View Help.

   See "Access Help for Database Views" on page 361.

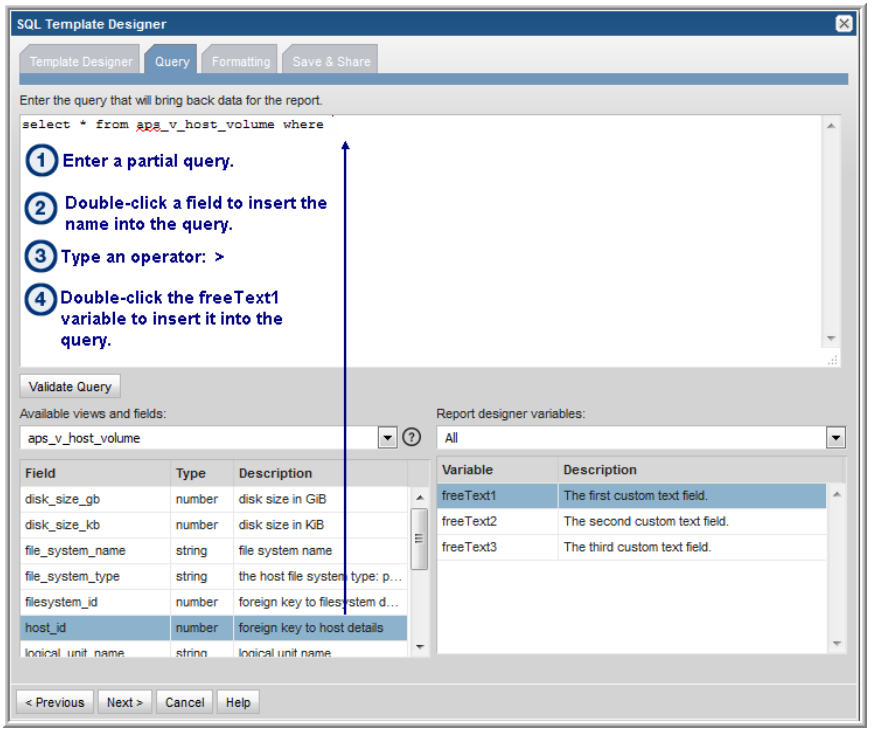
   See "Help for Database Views" on page 368.

   - Enter a partial query: **select * from aps_v_host_volume where**
   - Double-click a field to insert the name into the query.

- Type an operator: **>**
- Double-click the freeText1 variable to insert it into the query.

The resulting query will be:

```
select * from aps_v_host_volume where aps_v_host_volume.host_id
> ${freeText1}
```



5. Click **Validate Query**.
6. In the Formatting window, select all the fields for a table.
7. Save the report template with a name and a menu group.
8. Generate the report from this report template and in its Scope Selector window, provide a value for the Host ID field. This value will be passed to the query.

Access Help for Database Views

1. Click the icon beside the **Available views and fields** drop down to access column descriptions for Portal database views.

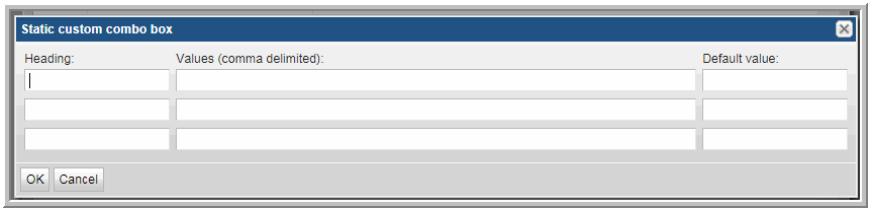
2. Navigate through descriptions for Base Portal Views, base views for licensed modules, and third-party vendor-specific views for your building your query.
3. If required, download a pdf version of the Views for offline use.

See "Help for Database Views" on page 368.

Static Custom Combo Box

The Report Designer Static Combo Box component enables the flexibility of offering the selection of various characteristics when the report is generated--similar to the way the out-of-the-box reports handle options, such as Event Type or Job Type. A Combo Box becomes available in the Scope Selector, enabling a user to select items from a drop-down list. This is particularly useful in environments where custom attributes have been defined for objects, enabling the user to select specific attributes at runtime.

Specify a heading, along with a list of values that will be displayed as a drop-down selection. You can also set the default value to be selected.



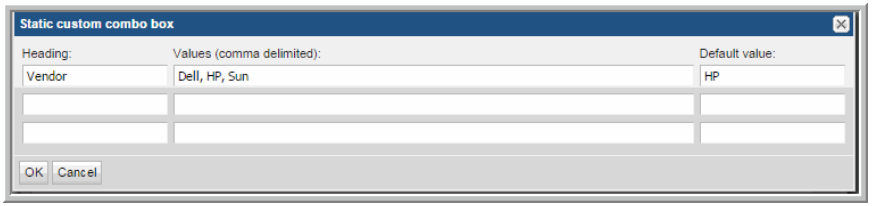
To include a blank or no choice option, specify the list of values in the form:

```
,option 1,option 2
```

If this no choice option is selected when the report is generated, an empty string will be passed to the SQL expression.

Example of a Combo Box Configuration in a SQL Template

1. Create a Custom Combo Box to enable the user to select all hosts for a particular Make:



2. This Combo Box can be used in a report template query, such as:

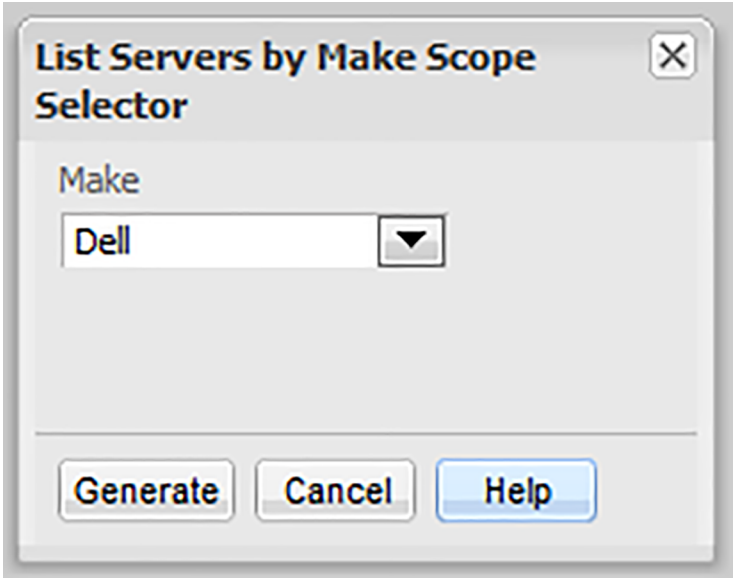
```
select * from apt_v_server where apt_v_server.make =
'${freeCombo1}'
```

Note: Report designer variables are listed at the bottom right of the dialog window. Double-click a freeCombo variable to insert it into the query. If the **${freeCombo1}** value is a string, it must be enclosed in single quotes to be evaluated as a text field, as shown in the above example. To determine if a database view field is a string or numeric, refer to the Database Views Help.

See "Help for Database Views" on page 368.

3. Format this report as a table.
4. Save it as **List Hosts by Make** in a Menu Group.
5. Generate a report from this newly saved report template.

   In this example, a Combo Box heading was specified with a list of values that will be presented in a drop-down list in the scope selector, when the user generates the report:

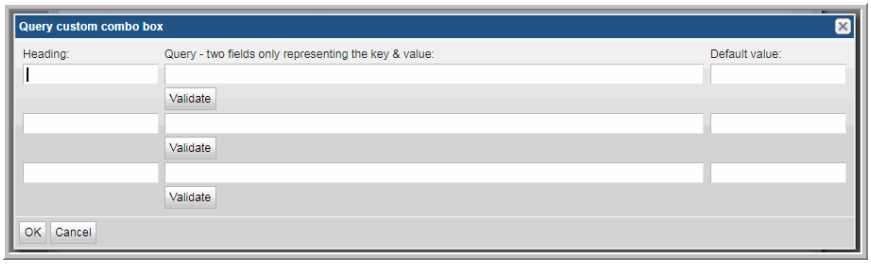


6. Select a **Make** from the drop-down list and click **Generate**.

Query Custom Combo Box

The SQL Template Designer offers a feature to design a combo box that is populated with the results of the query. In addition to configuring the Query Custom Combo

Box, a relevant report template query must be constructed to enable accurate report filtering using the combo box.



1. In the Template Designer component, double-click **Query custom combo box** to access the configuration window.
2. In the Heading field, enter the heading that will appear in the report template's scope selector with the drop-down list.
3. In the Custom Combo Box query field, supply a query with the following components and then click **Validate**.

SQL command	SELECT DISTINCT
Key field, Field value	The comma-separated **key-value pair** is derived from a published view. For example: Key: **storage_array_id** Value: **array_name** The **Key** in the key-value pair is what will substitute the ${queryCombon} variable and the **Value** is shown in the Combo box. Use the help for descriptions of the published database views, including the available fields. See "Help for Database Views" on page 368. Access these views and fields using the SQL Template Designer Query tab. See "Construct the SQL Query" on page 367.
Published view	End the SELECT statement with FROM and the published view, with optional list criteria. For example: `FROM aps_v_array_group ORDER BY array_name ASC`

Sample Query Custom Combo Box Queries

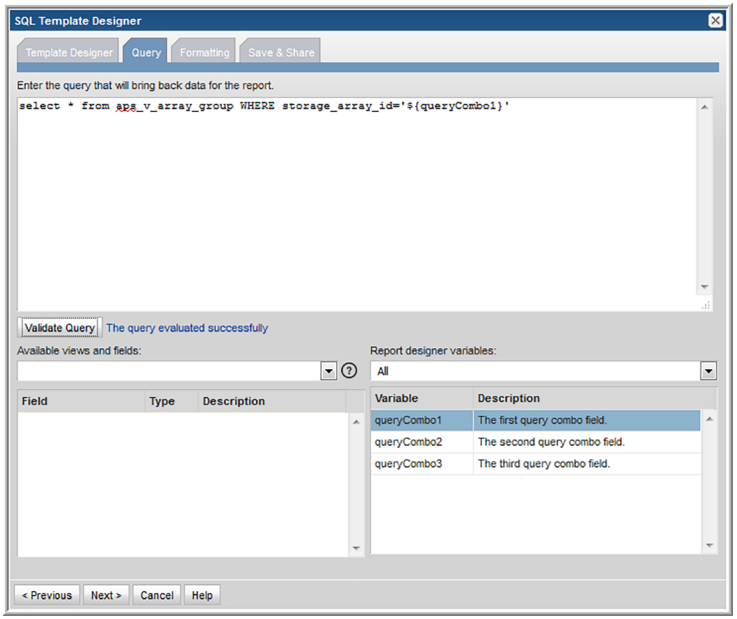
```
SELECT DISTINCT server_id, server_name FROM apt_v_job
SELECT DISTINCT client_id, client_name FROM apt_v_job
SELECT DISTINCT product_type, product_type_name FROM apt_v_job
SELECT DISTINCT job_type, job_type_name FROM apt_v_job
SELECT DISTINCT vendor_status, vendor_status_name FROM apt_v_job
SELECT DISTINCT policy_id, policy_name FROM apt_v_nbu_job_detail
SELECT DISTINCT policy_type, policy_type_name FROM
apt_v_nbu_job_detail
SELECT DISTINCT media_server_id, media_host_name FROM
apt_v_nbu_job_detail
SELECT DISTINCT storage_array_id, array_name FROM
aps_v_array_group
```

4. In the Query tab, enter a SQL query that will use the value selected in the populated combo box. For example:

```
select * from aps_v_array_group WHERE
storage_array_id='${queryCombo1}'
```

Note: Report designer variables are listed at the bottom right of the dialog window. Double-click a queryCombo variable to insert it into the query. If the **${queryCombo1}** value is a string, it must be enclosed in single quotes to be evaluated as a text field. To determine if a database view field is a string or numeric use the Database Views Help. For more information,

See "Help for Database Views" on page 368.



Example of a Query Custom Combo Configuration in a SQL Template

The following query produces the Ultimate NetBackup Job Status Report.

In the SQL Template Designer Query Custom Combo box, enter: **select window_group_id, window_group_name from apt_v_date_window order by window_group_name**

Then, in the Query, use the following query to report Ultimate Job Status:

```
SELECT
        a.nbu_job_id,
        a.client_id,
   b.client_host_name,
        b.policy_name,
        b.job_type_name,
        decode(a.overall_status,0,'Success',
             1,'Partial',2,'Queued',3,'Running','Failed')
overall_status,
        a.start_date,
        a.finish_date,
        a.kilobytes,
```

```
        file_pathlist,
        decode(was_restarted,1,'Y','N') was_restarted
      FROM table (nbu_rtd.listJobSummaryAfterRestart (
      ${startDate},
      ${endDate},
      ${queryCombo1},
      ${spHosts},
      null,
      null,
      null)) a
, apt_v_nbu_job_detail b
      where a.job_id = b.job_id
```

Construct the SQL Query

The SQL Template Designer supports SQL select statements. You cannot specify UPDATE, DROP, ALTER, or CREATE statements.

To view the available database views and fields that can be used when you create a query, click the Help icon.

See "Access Help for Database Views" on page 361.

Refer to examples of queries.

See "Sample SQL Queries" on page 370.

1. In the Query window, construct your select statement using the following tips:

- Use the drop-down list at the bottom left of the window to select views and fields.
- Use the Help to look up the valid tables, fields and variables.
  See "Access Help for Database Views" on page 361.
- To insert a view name into a query, double-click a field and then delete the field portion. For example, in the aps_v_database_datafiles_log view, click on datafile_id. Then, in the query, delete the **.datafile_id** portion.
- Double-click on fields to insert them into the query with the correct syntax.
- Double-click on variables in the list at the right of the window, to insert variables with proper syntax into the query--for example, ${endDate}
- If you checked Host Group and Client Scope in the initial tabbed Template Designer window, you will see that the drop-down list of Template Designer Variables (at the right of the window) includes Report Scope selections. Using these selections, you can enumerate a list of values for Host Groups (hostGroups) and Clients (hosts).

- All evaluated columns--for example, nvl(t2.das_capacity,0)--must have an alias name.
- Certain special characters, when used in SQL queries, must be escaped with a backslash so that they can be evaluated as literals. For example, to treat a **$** character as a literal value, use: **\$**.
  See "Special Characters in SQL Queries" on page 369.
  See "Example of Using a Special Character as a Literal in a SQL Query" on page 375.

1. Click **Validate Query**.

- Correct the statement before you proceed. Use the available lists to view valid tables, fields, and variables.
- Once the query validates, click **Next** to proceed to format the SQL template output.
  See "Format the SQL Template Output" on page 376.

Help for Database Views

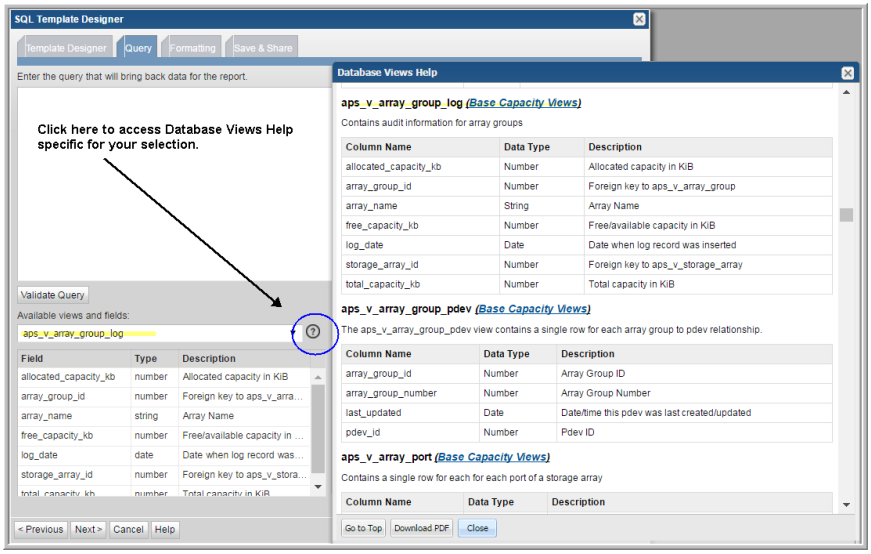
IT Analytics provides a set of published read-only database views into the reporting database and related object model. The purpose of these views is to provide a read-only reporting or data extraction mechanism for advanced users of the product.

Database views provide a native SQL interface into the reporting database. IT Analytics publishes a collection of views that overlay the underlying database tables. Using these read-only views, you can write your own reports or stored procedure handlers to query the reporting database. These views provide a simple and fast mechanism to access a read-only view into the raw data maintained within the reporting database.

There are multiple types of database views:

- **Base Portal Views**. Provides read access to the properties for every inventory item that the product maintains.
- **Base Licensed Module Views**. Provides read access to information about a given underlying IT Analytics licensed module.
- **Vendor Specific Views**. Provides read access to data specific to a third-party vendor, such as NetApp or HDS.

Use the Help specific to the database views to look up the valid tables and fields relevant to your environment. You can use the help as you build the query or you can download it as a PDF to use offline. Access column descriptions for Portal database views by selecting a view and then clicking the icon beside the Available views and fields drop down. You can also just click the icon to see the full listing of all the views.



Navigate through descriptions for Base Portal Views, Base Views for licensed modules, and third-party vendor-specific Views for your building your query.

Accessing the Published Database Views from the CLI

The underlying Oracle database can be accessed via the command-line interface (CLI). The product includes a separate Oracle user account that is specifically designed and granted SELECT access to the Published Database Views.

The oracle user account for the Published Database Views is called **aptare_ro**. The default password for this account is set to **aptaresoftware123**. This user account can ONLY access the published database views and does not have any update or insert privileges.

Special Characters in SQL Queries

Certain characters have a special meaning for Oracle functions. In a SQL Template Designer query, a special character must be escaped in order for the parser to treat it as its intended literal value. The following examples illustrate how to handle commonly used special characters in a query.

See "Example of Using a Special Character as a Literal in a SQL Query" on page 375.

- Backslash (\ ) - Use "\\" instead of "\"
- Dollar sign ($) - The $ character must be prefaced (escaped) with a backslash: '.\$'

Sample SQL Queries

The following queries can serve as a starting point for how to create your own custom queries. In fact, you simply can copy and paste a query from these examples to demonstrate the results in the SQL Template Designer.

Example of a Query for Host Attributes

The following query lists all the attributes associated with a host/server:

```
select *
FROM apt_v_server_attribute
WHERE host_id IN (${hosts})
```

Using the apt_v_server_attribute View in Queries

The database view, apt_v_server_attribute, is dynamically created using the attributes that you create for your environment. When you initially look at this view, you only will see host_id and host_name. This view is recreated during the upgrade process, based on the attributes that you have configured.

If you add or modify server attributes in any way, in order to immediately use this view in queries in the SQL Template Designer, you will need to execute the following steps to manually recreate the apt_v_server_attribute database view:

1. Connect to the database as a Portal user:

   ```
   sqlplus <userID>/<pwd>
   ```

2. Execute the following:

   ```
   EXECUTE dynsql_pkg.recreateDynAttributeView;
   ```

3. Connect to the database as sysdb:

   ```
   sqlplus / as sysdba
   ```

4. Execute the following:

   ```
   CREATE OR REPLACE VIEW aptare_ro.apt_v_server_attribute AS SELECT
    * FROM portal.apt_v_server_attribute WITH READ ONLY;
   ```

Considerations for Attributes Used in SQL Template Designer Queries

The attribute database view, apt_v_server_attribute, is dynamically created from server attributes you enter in the Portal. During the creation of this database view, several rules are applied to facilitate the use of this view in queries in the SQL Template Designer. When you see the attribute names listed in the SQL Template Designer Formatting tabbed window, you'll see that the following conversions have been made to your Attribute Names--that is, the view's column names.

- All characters are lowercase.
- Special characters, such as ":" or "/" or a space, are converted to an underscore.
- If the attribute name begins with a number or a special character, it will be replaced with: c_
- If the conversion process results in duplicate names, the attribute name will have a suffix appended to differentiate the duplicates; for example: _1 or _2
- Names are truncated to 30 characters.

Example of a Query of Failed Backup Jobs

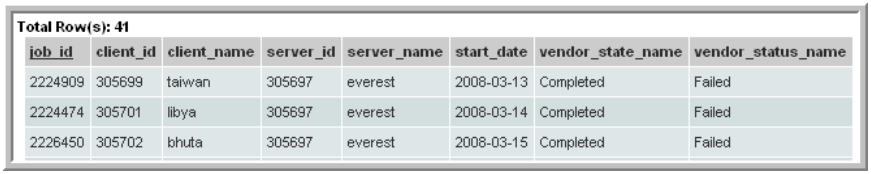
The following example results in a table of failed jobs.

1. In the SQL Template Designer, check both **Date Range**and **Host Groups and Client Scope**.
2. In the Query window, enter the following select statement and click **Validate Query**:

```
SELECT apt_v_job.job_id,apt_v_job.client_id,
apt_v_job.client_name,
apt_v_job.server_id, apt_v_job.server_name,
apt_v_job.start_date,apt_v_job.vendor_state_name,
apt_v_job.vendor_status_name
FROM apt_v_job
WHERE apt_v_job.summary_status = 2 --Failed jobs
  AND apt_v_job.start_date > ${startDate}
  AND apt_v_job.start_date < ${endDate}
  AND apt_v_job.client_id IN (${hosts})
```

3. In the Formatting window, select the fields to be displayed. For this example, it makes sense to **Select All** and display the report as a **Table**.
4. Click **Next**, enter a report name and select a Menu Group. Then, click **Finish**.

5. When you run this report, specify either a time period or start and end dates. You also can modify the scope to generate the report for a specific host group. The output will look something like this:



Total Row(s): 41

job_id	client_id	client_name	server_id	server_name	start_date	vendor_state_name	vendor_status_name
2224909	305699	taiwan	305697	everest	2008-03-13	Completed	Failed
2224474	305701	libya	305697	everest	2008-03-14	Completed	Failed
2226450	305702	bhuta	305697	everest	2008-03-15	Completed	Failed

Example of the SQL Custom Join Feature in a SQL Template

Backup Exec data collection does not populate tables related to tape media--for example, apt_v_tape_media. To include this view in a query so that it will work with Backup Exec data, you'll need to use an "outer join" (as denoted with (+) in the following query).

```
select apt_v_job.server_name, apt_v_job.job_type,
apt_v_job_tape_media.media_name, apt_v_job.client_name,
to_char(apt_v_job.start_date, 'YYYY-MM-DD hh:mm:ss AM') start_date,
 to_char(apt_v_job.finish_date, 'YYYY-MM-DD hh:mm:ss AM') finish_date,
 apt_v_job.summary_status, apt_v_job_message_log.message,
apt_v_job_tape_media.tape_media_id, apt_v_job.kilobytes
from apt_v_job, apt_v_job_tape_media, apt_v_job_message_log
where apt_v_job.job_id = apt_v_job_tape_media.job_id (+)
and apt_v_job.job_id = apt_v_job_message_log.job_id
and apt_v_job.server_id in (${rp.hosts})
and apt_v_job.start_date BETWEEN ${rp.startDate} AND ${rp.endDate}
ORDER BY apt_v_job.server_name, apt_v_job.start_date
```

Example of a Sparkline Query in a SQL Template

A sparkline presentation of the data can be added as a column to a tabular report to plot a series of values. Sparklines enable an at-a-glance view of data spikes that may indicate an issue that requires attention.

The basic requirements for a sparkline chart in a SQL template include:

- **Table template**. Sparklines are embedded in a table cell.
- **Data over time**. The series of values can be derived from: **start_date between ${startDate} AND ${endDate}** in the query.

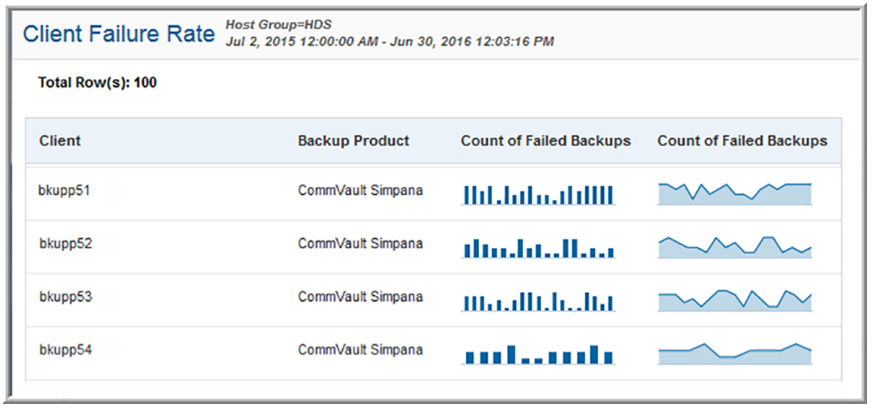
- **Comma-separated list of string values**. Use the pipelined function to achieve this.
  See "collectString" on page 425.

The following example graphs a series of client failure values as a sparkline in a tabular report.

1. In the SQL Template Designer, check **Date Range** and **Host Groups and Client Scope**.
2. In the Query window, enter the following select statement and click **Validate Query**:

```
with spark as (
select trunc(start_date), client_id, client_name,
product_type_name,count(job_id) failed_count
FROM apt_v_job
WHERE client_id in(${hosts})
AND start_date between ${startDate} AND ${endDate}
AND summary_status = 2
group by client_id, client_name, product_type_name,
trunc(start_date)
)
select display_name, product_type_name,
rtd.collectString(cast(collect(TO_CHAR(failed_count)) as
StringListType), ', ') failed_count,
rtd.collectString(cast(collect(TO_CHAR(failed_count)) as
StringListType), ', ') failed_count_area
from apt_v_server h, spark s
where h.server_id = s.client_id
group by display_name, product_type_name
order by 1,2
```

3. In the Formatting window, select all the fields to be displayed and display the report as a **Table**.
4. In the Formatting window, for failed_count, select the **Column Sparkline** formatter and for the failed_count_area, select the **Area Sparkline** formatter.
5. Click **Next**, enter a report name and click **Finish**.The output will look something like this:



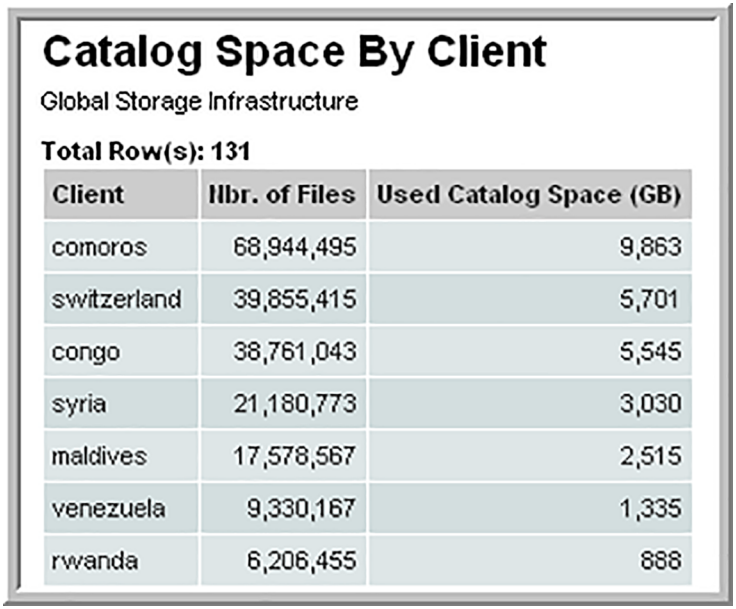
Example of Sums in a SQL Template

The following steps can be used to create the NetBackup Catalog Space by Client List. This example demonstrates how to include sums of a field.

1. In the SQL Template Designer, check **Host Groups and Client Scope**.
2. In the Query window, enter the following select statement and click **Validate Query**:

```
SELECT client_host_name,
sum(nbr_of_files),sum(nbr_of_files)*150/1024/1024
FROM apt_v_nbu_job_detail
WHERE client_id in (${hosts})
AND summary_status is not null
AND expiration_date <= sysdate
GROUP BY client_host_name
ORDER BY sum(nbr_of_files) DESC
```

3. In the Formatting window, select all the fields to be displayed and display the report as a **Table**.
4. Click **Next**, enter a report name and click **Finish**.The output will look something like this:



Catalog Space By Client

Global Storage Infrastructure

Total Row(s): 131

Client	Nbr. of Files	Used Catalog Space (GB)
comoros	68,944,495	9,863
switzerland	39,855,415	5,701
congo	38,761,043	5,545
syria	21,180,773	3,030
maldives	17,578,567	2,515
venezuela	9,330,167	1,335
rwanda	6,206,455	888

Example of Using a Special Character as a Literal in a SQL Query

Outside of the Portal's SQL Template Designer environment, certain characters and symbols have a special meaning to Oracle functions when querying an Oracle database directly. For example, the $ can be used as a parameter in an Oracle function. In the context of the SQL Template Designer, this $ character must be prefaced (escaped) with a backslash, as shown in the following example.

```
select
REGEXP_REPLACE(APT_V_TSM_JOB.DOMAIN_NAME|| '.' ||
APT_V_TSM_JOB.SCHEDULE_NAME,'.\$','')       GROUP_NAME
from APT_V_TSM_JOB
```

Use Functions in Queries

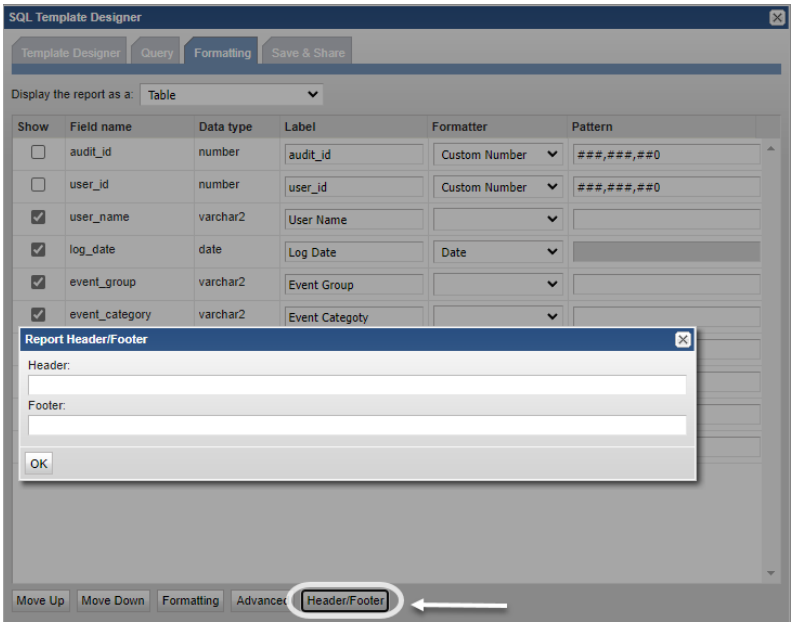
To further expand the functionality of the SQL Template Designer, a number of functions are available. These are sometimes referred to as pipelined functions.

See "About Pipelined Functions" on page 422.

Format the SQL Template Output

1. Format fields in a report within the **Formatting** tab.
   - Check the query fields to select the fields to be displayed in the report.
   - Move fields up and down to position them in the report's output.
   - Modify the **Label**forthe displayed fields to something more relevant to your needs.
   - Modify the number, size, date and time formatting.
     See "Number, Size, Date, and Time Formatting" on page 377.
   - Modify the alignment, aggregation, bar type and bar type color.
     See "Alignment, Aggregation, Bar Type, and Bar Type Color" on page 381.
   - Modify the area, line or column, sparkline format for an SQL tabular report template.
     See "Area, Line or Column Sparkline Format for a SQL Tabular Report Template" on page 383.
2. Choose a Report Type:
   - See "Configure a Bar Chart SQL Template" on page 384.
   - See "Configure a Line Chart SQL Template" on page 386.
   - See "Configure a Donut Chart SQL Template" on page 388.
   - See "Configure a Gauge Chart SQL Template" on page 389.
   - See "Configure an Area Chart SQL Template" on page 395.
   - See "Configure a Horizontal Bar Chart SQL Template" on page 392.
   - See "Configure a Horizontal Stacked Bar Chart SQL Template" on page 393.
   - See "Configure a Pie Chart SQL Template" on page 398.
   - See "Configure a Pivot Table SQL Template" on page 401.
   - See "Configure a Stacked Bar Chart SQL Template" on page 403.
   - See "Configure a Table in the SQL Template Designer" on page 405.
3. Header/Footer

   Click **Header/Footer** button. IT Analytics displays **Report Header/Footer** dialog box is displayed.



In the **Footer** field, specify the current date and time using : `${currentDateTime}` and then click **OK**.

For example: Report generated at ${currentDateTime}

Number, Size, Date, and Time Formatting

The following field formatters are provided:

- Number: This format is based on the Java class NumberFormat. This formatter ignores the Pattern listed in the SQL Template Designer.
  See "Examples of Negative Value Formatting" on page 379.
- Custom Number: This format is based on the Java class NumberFormat. Unlike the Number formatter, this Custom Number formatter takes into account the Pattern when formatting the value.
  See "Examples of Negative Value Formatting" on page 379.
- Decimal: This formatter is based on Java decimal formatters, DecimalFormat.
- Date: This formatter is based on the Java class SimpleDateFormat. Note that the Date Formatter is available for templates developed in earlier releases, however, it is no longer supported in the SQL Template Designer because the format of the date is controlled by the user's profile. For special cases, where custom date formatting is required, use the Oracle TO_CHAR function in the

query. This function converts a date or interval value to a character data type in the specified format: Date, Timestamp.

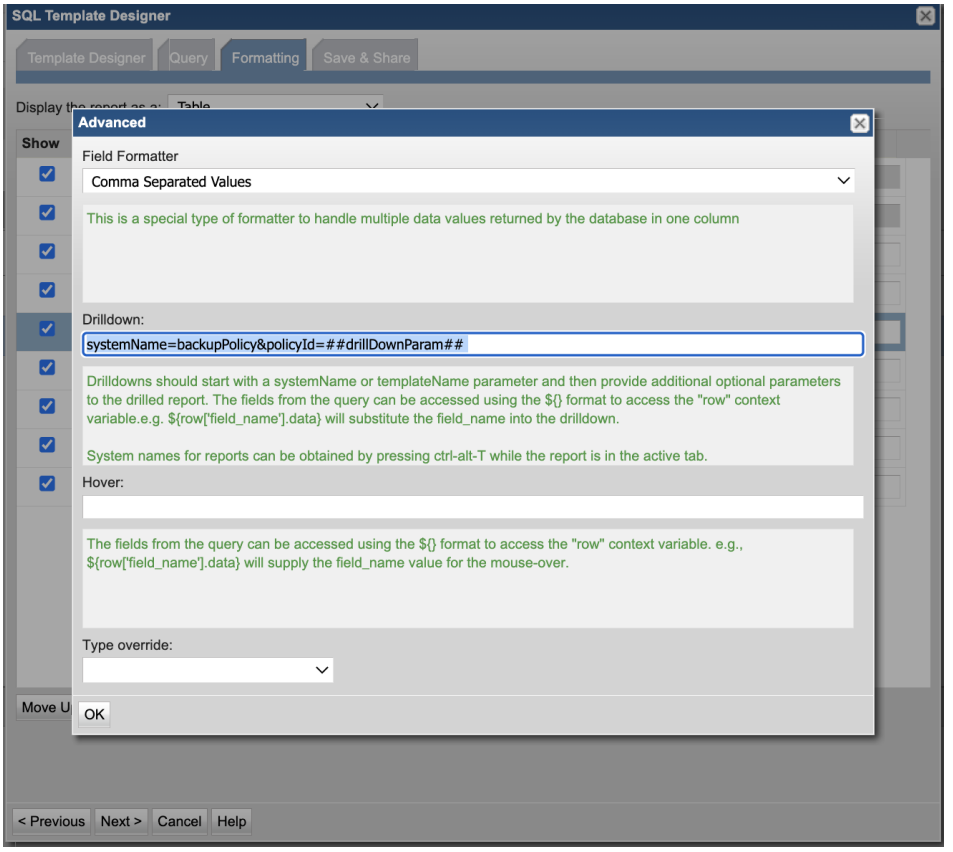
- File Size: This formatter formats the file size with the unit of measure that is most applicable to the field value. For example, 2048 MB would become 2 GB. The pattern field takes the following values: Bytes, KB, MB, GB, and TB (or the Kibibyte multiples: KiB, MiB, GiB, and TiB). This is the unit of the data returned from the Portal database.
- Percentage Bar: The Percentage formatter results in a horizontal bar, representing the value in a tabular report column. For example, **50:75** results in a color designation, where a value < 50 renders green, 50-75 renders yellow, and > 75 renders red.
- Status Icon: Colored icons, relevant to a value, will render in a tabular report column.
  See "Status Icon Values" on page 380.
- Area Sparkline:
  See "Area, Line or Column Sparkline Format for a SQL Tabular Report Template" on page 383.
- Column Sparkline:
  See "Area, Line or Column Sparkline Format for a SQL Tabular Report Template" on page 383.
- Line Sparkline:
  See "Area, Line or Column Sparkline Format for a SQL Tabular Report Template" on page 383.
- Truncate: Truncates a string with ellipsis. The maximum size of the string is 28 characters. Use the pattern to override the maximum size. When using the Truncate formatter, it is recommended to use the Advanced dialog to create a hover tool tip.
  See "Create mouse-over hovers" on page 411.

Field Formatter

A new drop-down **Field Formatter** has been incorporated in SQL Template Designer. The drop down has the following options to select:

- Comma Separated Values

Navigate to **Formatting** tab and select the required field and then click **Advance**. **Advanced** dialog box is displayed with **Field Formatter** drop down.



User can add drill-down to each comma separated values within one column.

To view the drill down report for the SQL output, specify the drill down command in the **Drilldown** text box

For example: `systemName=backupPolicy&policyId=##drillDownPara##`

Examples of Negative Value Formatting

Since the number format is based on Java formatters, you can use a semi-colon (;) sub-pattern boundary to represent negative numbers so that they stand out differently from positive values.

The following examples show how to use the semi-colon to represent the negative value, -1234.56.

- #,##0.00;(#,##0.00) displays the value, -1234.56, as (1,234.56)

- #,##0.00;'<font color=red>'(#,##0.00) displays the value, -1234.56, in red as (1,234.56)

Status Icon Values

Use the following values for a status icon in a tabular report.

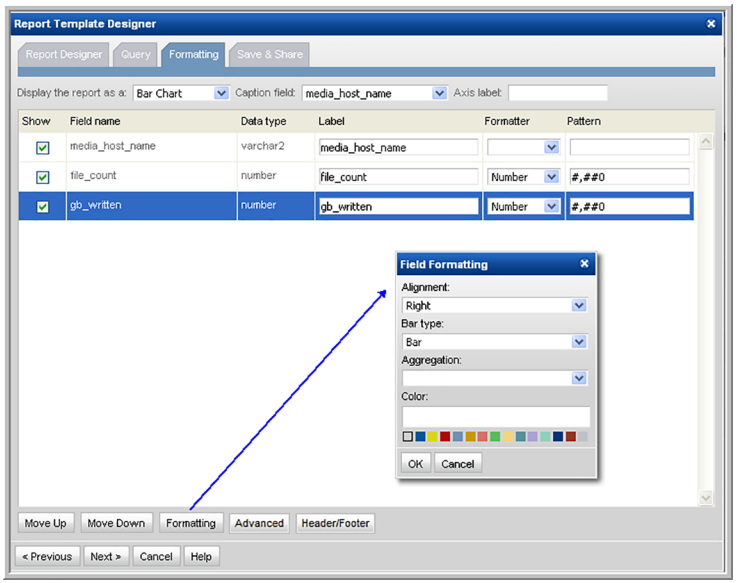
Value	Icon
red	
yellow	
white	
blue	
green	
error	
warning	
success	
fast	
medium	
slow	
status0	
status1	
status2	
status3	

Value	Icon
status4	
status5	
status6	
status7	
status8	
status9	
status10	
status11	
status12	
status13	
status14	
status15	
trend_down	
trend_flat	
trend_up	
unknown	

Alignment, Aggregation, Bar Type, and Bar Type Color

Select the Field and click **Formatting** (at the bottom of the window) to view the formatting options.

- **Data alignment** - Left, Center, or Right.
- **Aggregation** - Sum, Average, Minimum, Maximum
- **Bar type** - Bar or Line (for bar charts only)
- **Color** - Bar or Line (for bar charts only)



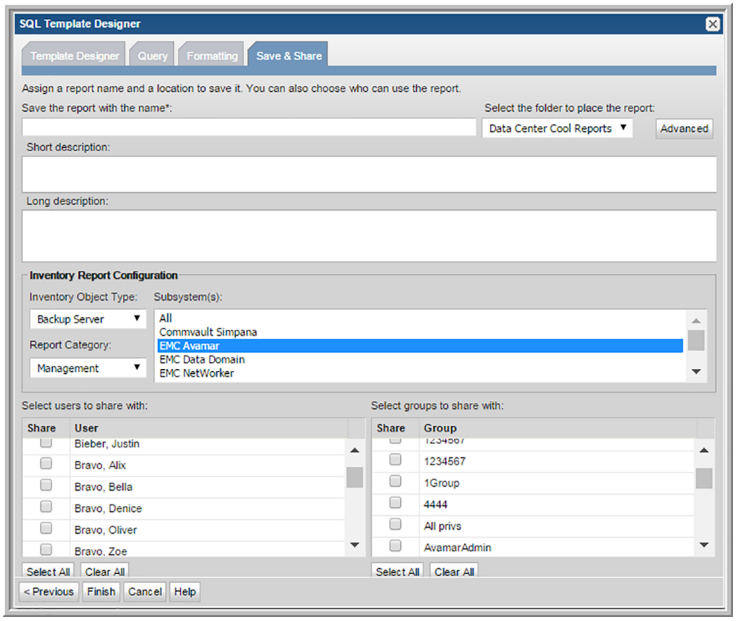
1. At the bottom of the Formatting window, use the **Move Up** and **Move Down** buttons to organize the fields in the order that you want them to appear in a tabular report. Be sure to select a row before shifting the fields.
2. In the **Save & Share** window, complete the appropriate text boxes to save and share the report, then click **OK**. You can choose to share with one or more users and user groups.

   See "Save and Share Report Templates" on page 407.

   You can also assign report configurations for the **Inventory**. This determines how the report templates will be classified and displayed in the Inventory.

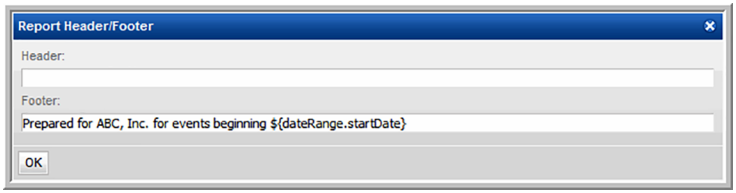
   See "Inventory Report Configuration" on page 416.

   The saved report template will be listed in the report group that you select.



Configuring a Header and Footer in a Custom Report Template

At the bottom of the Formatting window in the SQL Template Designer, click the **Header/Footer** button to launch the window where you can define the text and variables that will be displayed when the report is generated.



Any variables used in the query can be included in a header or footer, as shown in the above example.

Area, Line or Column Sparkline Format for a SQL Tabular Report Template

Sparklines render as small charts embedded in each row of a tabular report, presenting an at-a-glance visualization of data trends. You can display a number series as an Area Sparkline, Line Sparkline, or a Column Sparkline in a column of

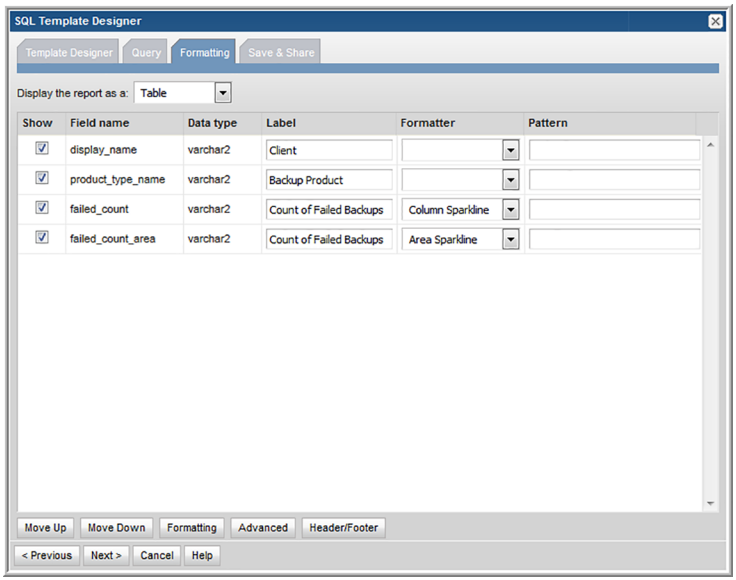
a tabular report to illustrate spikes in data that may indicate issues that require attention. For details on how to configure a sparkline in a SQL Template,

See "Example of a Sparkline Query in a SQL Template" on page 372.

This section covers only the SQL Template formatting selections required to render the data as a sparkline.

To format a table column as a sparkline chart, take the following steps.

1. In the SQL Template Designer Formatting tabbed section, select the field to be rendered as a sparkline chart.
2. From the Formatter list, select either Column Sparkline, Line Sparkline, or Area Sparkline.



Configure a Bar Chart SQL Template

A bar chart represents data with vertical bars. This type of chart can display bars only or it can be configured to include a line that will be charted on the right y-axis.

- See "Example of a Bar Chart SQL Template: Host CPU Performance" on page 385.
- See "Configure a Horizontal Bar Chart SQL Template" on page 392.
- See "Configure a Horizontal Stacked Bar Chart SQL Template" on page 393.

- See "Configure a Stacked Bar Chart SQL Template" on page 403.

Example of a Bar Chart SQL Template: Host CPU Performance

The following example illustrates how to create a Bar Chart and then format it with a dual axis. This example produces a bar chart for Host CPU Performance.

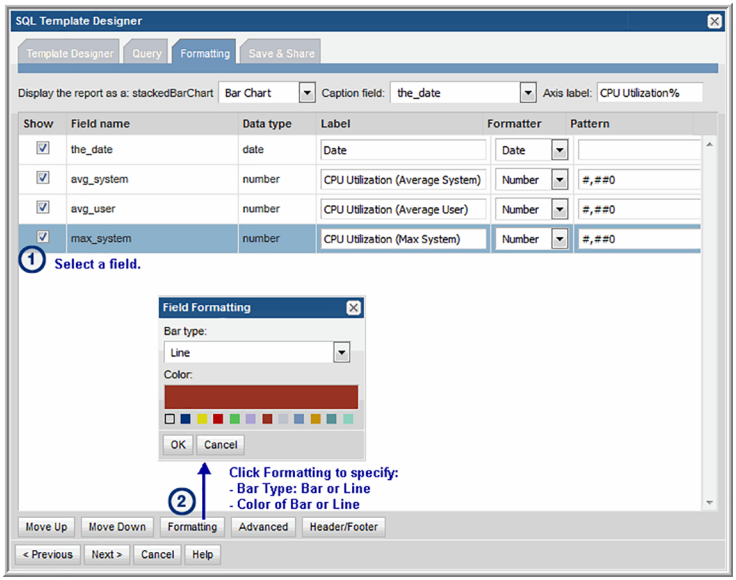
1. From the Template Designer tabbed window, select **Date Range** and **Host Groups and Client Scope** and **Static Custom Combo Box**.
2. Create a query that has a field for the caption, each of the bars, and the line.

```
SELECT
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24','Day',
'DD','Week','WW','Month','MM',
'Quarter','Q','Year')) the_date,
avg(system_processing_time_pct) avg_system,
avg(user_processing_time_pct) avg_user,
max(system_processing_time_pct) max_system
FROM apt_v_host_cpu_log
WHERE log_date between ${startDate} AND ${endDate}
GROUP BY
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24','Day',
'DD','Week','WW','Month','MM',
'Quarter','Q','Year'))
ORDER BY
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24','Day',
'DD','Week','WW','Month','MM',
'Quarter','Q','Year')) ASC
```

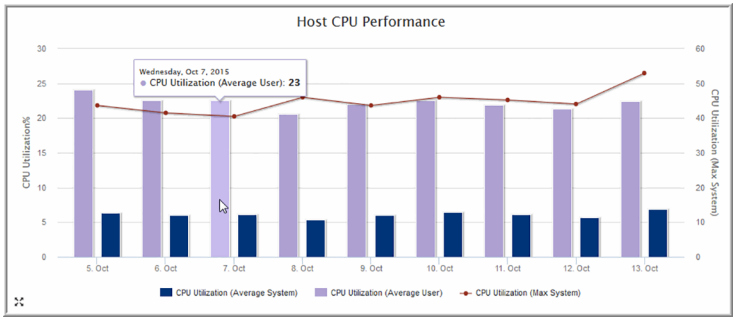
3. In the Formatting tabbed section, select each field (data points) and format each so that you have:
   - 1 caption, n bars, and 1 line

   Note: The line will be charted on the right-hand axis.

4. Set the color of the bars and specify the line formatting, as shown in the following screen.



5. Save the report to the **My Reports** report menu group. Then, click on the saved report to generate it. The output will look something like this:



Configure a Line Chart SQL Template

Line charts provide an effective way to visualize metrics on a time line. These charts are particularly useful for performance metrics, where you want to compare multiple metrics to assess historical performance.

Example of a Line Chart SQL Template

The following example illustrates a simple bar chart.

1. From the **Template Designer** tabbed window, select **Data Range** and **Array Scope Selector**.

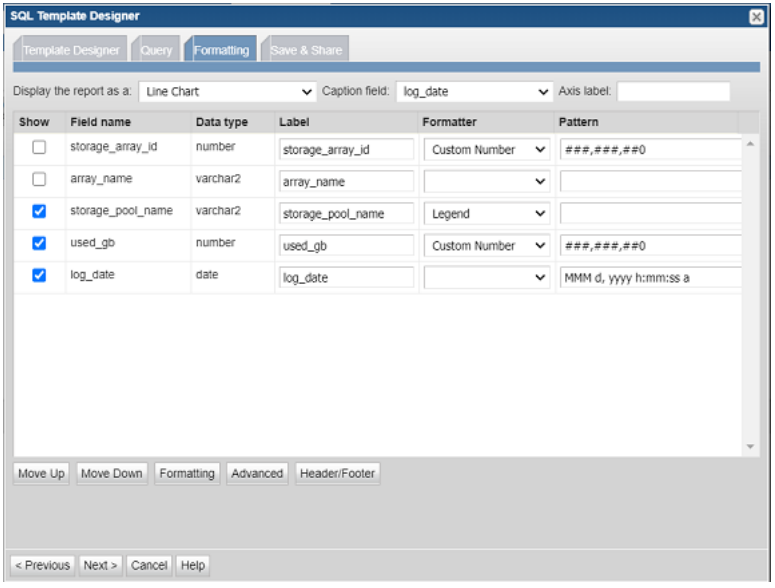
2. Create a query:

```
select storage_array_id, array_name, storage_pool_name,
max(used_capacity_kb/1024/0124) used_gb,
trunc(log_date) log_date
from aps_v_storage_pool_log
where storage_array_id IN (${arrays})
and log_date BETWEEN ${startDate} AND ${endDate}
group by storage_array_id, array_name, storage_pool_name, trunc(log_date)
```

3. Validate the query.
4. In the Formatting tabbed section, choose to display the report as a **Line Chart** and the select log_date as the **Caption** field.

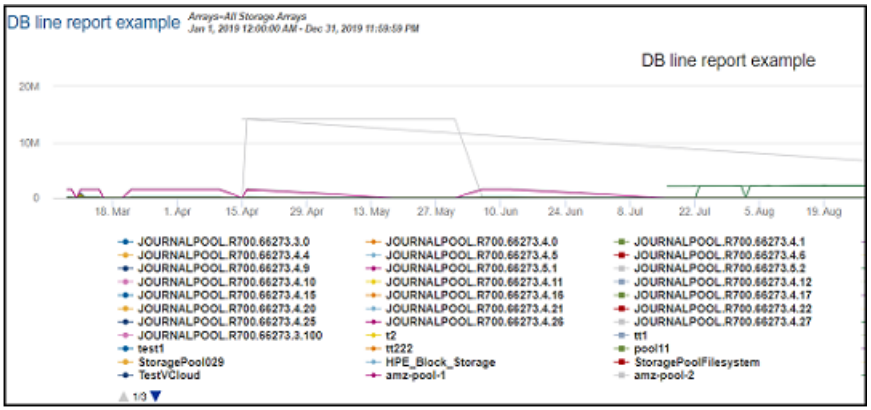
   Expectation for the legend is that only three fields are enabled for the formatter. One should be the Legend, next one should be the data and the last one should be a date.

Note: If there are more than 3 fields, you need to uncheck them.



5. Select select storage_pool_name, used_gb and log_date.
6. For storage_pool_name select **Legend** as the **Formatter**.

7. Save the report. When you generate it from the **Reports** tab, the output will display something like this:



Configure a Donut Chart SQL Template

A Donut Chart is similar to a Pie Chart.

- See "Example of a Donut Chart SQL Template: NetBackup Job Size by Job Type" on page 388.
- See "Configure a Pie Chart SQL Template" on page 398.

Example of a Donut Chart SQL Template: NetBackup Job Size by Job Type

The following example results in a donut chart that represents the various job types--such as, application backups, full backups, incremental backups, and restores--as segments in a donut chart.

1. In the SQL Template Designer, do not check any Template Designer items.
2. In the Query window, enter the following select statement and click **Validate Query**:

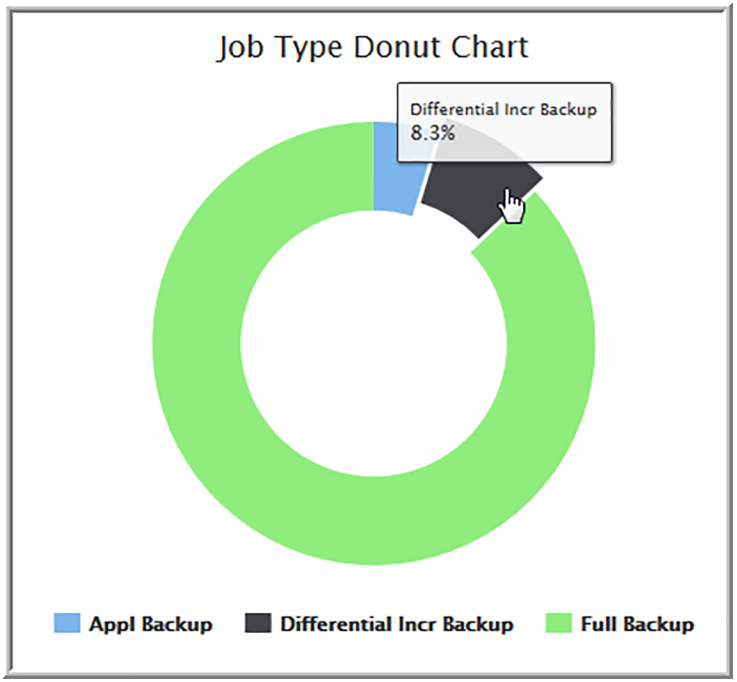
```
SELECT job_type_name,
sum(kilobytes/1024/1024) Job_SIZE_GB
FROM APT_V_NBU_JOB_DETAIL
WHERE
```

```
finish_date > sysdate -7
AND job_type_name IS NOT NULL
GROUP BY job_type_name
ORDER BY job_type_name
```

3. In the Formatting window, select **Donut Chart** from the Display report asdrop-down list.
4. For the Caption field, select **job_type_name** from the drop-down list.
5. Select all the fields to be displayed.

   For a donut chart, you need at least one field to be the caption and another field to be the segment.
6. Click **Next**, enter a report name and Menu Group. Then, click **Finish**.

When you run this report, the output will look something like this:

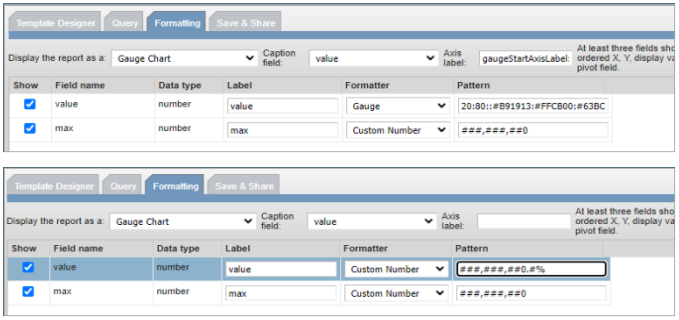


Configure a Gauge Chart SQL Template

The Gauge chart template query must provide two values. The first value is the actual data point - the value appearing between the minimum and maximum, where

the minimum value is always zero. The second value is the maximum value. For percentage-based Gauge charts, the second value must be 100.

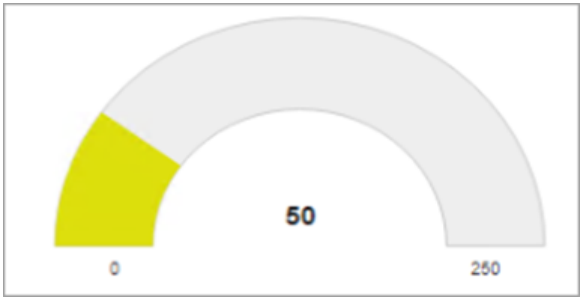
A Gauge chart is considered a percentage-based chart if the format pattern of the first value ends in a percentage, or the formatter is Gauge.



The default colors of a Gauge chart are green starting at zero, yellow at 10% of the maximum value, and red at 90%. If the chart is not percentage based, the values provided by the query are used as is. For example, if the query is as follows:

```
select 50, 250 from apt_v_dual
```

The resulting chart appears as below using the default color palette.

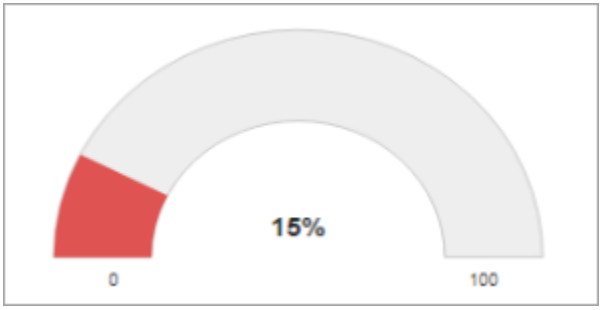


If the formatter selected for the first field is Gauge, then optional format pattern can be provided to define the color ranges. These are referred to as color stops. Two or three color stops can be specified. If no format pattern or an invalid pattern is specified, the default color stops are used.

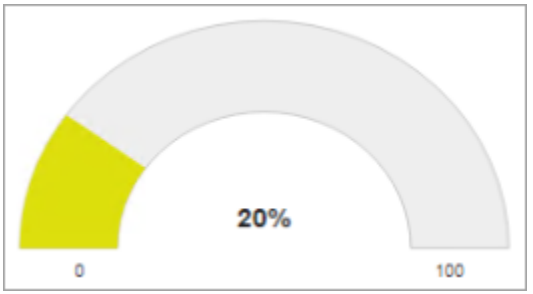
Specify two-color stops

In this scenario, red starts at zero, yellow starts at 20%, and green starts at 80%.

- Example-1: 20:80::red:yellow:green



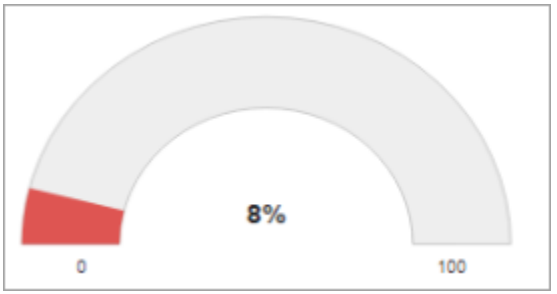
- Example-2: 20:80::#DF5353:#DDDF0D:#55BF3B



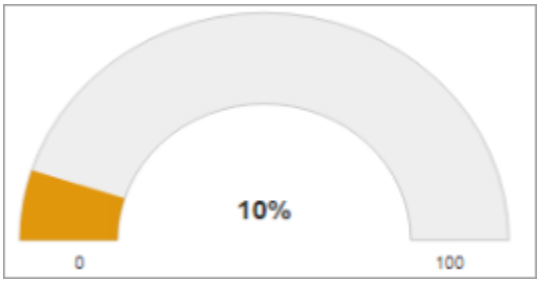
Example: Specify three-color stops

In this scenario, red start at zero, orange start at 10%, yellow starts at 20%, and green starts at 90%.

- Example-3: 10:20:90::red:orange:yellow:green



- Example-4: 10:20:90:: #DF5353:#DFBC0D:#DDDF0D:#55BF3B



Configure a Horizontal Bar Chart SQL Template

A horizontal bar chart, as evident by its name, renders data as individual horizontal bars.

- See "Example of a Horizontal Stacked Bar Chart SQL Template: Available/Allocated Capacity" on page 394.
- See "Configure a Bar Chart SQL Template" on page 384.
- See "Configure a Horizontal Stacked Bar Chart SQL Template" on page 393.
- See "Configure a Stacked Bar Chart SQL Template" on page 403.

Example of a Horizontal Bar Chart SQL Template: Host CPU Performance

1. From the SQL Template Designer tabbed window, select **Date Range** and **Host Groups and Client Scope**.
2. Select **Static Custom Combo Box** and click **Configure** at the bottom of the Query window.
3. Enter the heading for the drop-down list in the Static Custom Combo box--for this example, **Group by:**
4. Enter the following comma-separated list of values and then click **OK**:

   Hour,Day,Week,Month,Quarter,Year

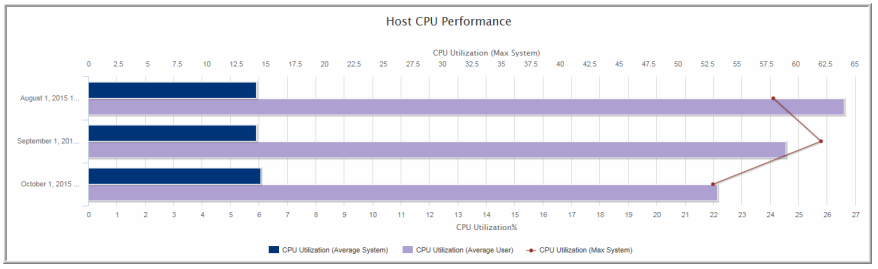
   These will be the options that a user can select when generating a report.
5. Click on the **Query** tab and enter the following query:

```
SELECT
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24','Day',
'DD','Week','WW','Month','MM',
```

```
'Quarter','Q','Year')) the_date,
avg(system_processing_time_pct) avg_system,
avg(user_processing_time_pct) avg_user,
max(system_processing_time_pct) max_system
FROM apt_v_host_cpu_log
WHERE log_date between ${startDate} AND ${endDate}
GROUP BY
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24','Day',
'DD','Week','WW','Month','MM',
'Quarter','Q','Year'))
```

6. Click **Validate Query** and then **Next**.
7. In the Formatting window, select **Horizontal Bar**and set the Caption field to **the_date** with the Axis label set to **CPU Utilization%**
8. Check all four fields: **the_date, avg_system, avg_user,**and **max_system.**
9. One by one, select the bar fields, **avg_system** and **avg_user**, and click the Formatting button (at the bottom of the window) to configure the color of the bar. And, enter a relevant label for each field.
10. Select the max_system field and click the Formatting button (at the bottom of the window) to configure the color of the line.
11. Save the report to a report menu group. Then, click the saved report to generate it.

The output for the horizontal bar chart will look something like this:



Configure a Horizontal Stacked Bar Chart SQL Template

A horizontal stacked bar chart, as evident by its name, renders data as stacked horizontal bars.

- See “Example of a Horizontal Stacked Bar Chart SQL Template: Available/Allocated Capacity” on page 394.
- See “Configure a Bar Chart SQL Template” on page 384.
- See “Configure a Horizontal Bar Chart SQL Template” on page 392.
- See “Configure a Stacked Bar Chart SQL Template” on page 403.

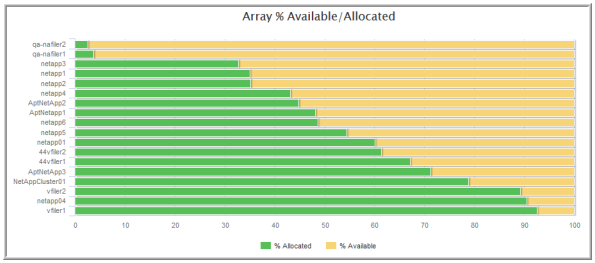
Example of a Horizontal Stacked Bar Chart SQL Template: Available/Allocated Capacity

1. From the SQL Template Designer tabbed window, select the **Static Custom Combo Box** and click **Configure** at the bottom of the Query window.
2. Enter the heading for the drop-down list in the Static Custom Combo box--for this example, **Array Vendor:**
3. Enter the following comma-separated list of values and then click **OK**: **All,EMC,HDS,IBM,NetApp**

   These will be the options that a user can select when generating a report.
4. In the Query tabbed window, enter the following SQL query.

```
SELECT
array_name,
(allocated_kb)/(allocated_kb+available_kb+.001)*100 pct_allocated,
(available_kb)/(allocated_kb+available_kb+.001)*100 pct_available
FROM aps_v_storage_array
WHERE allocated_kb >= 0
AND available_kb >= 0
AND vendor_name LIKE
DECODE('${freeCombo1}','All','%','%${freeCombo1}%')
ORDER BY (available_kb)/(allocated_kb+available_kb+.001)*100 DESC
```

The output for the horizontal stacked bar chart will look something like this:



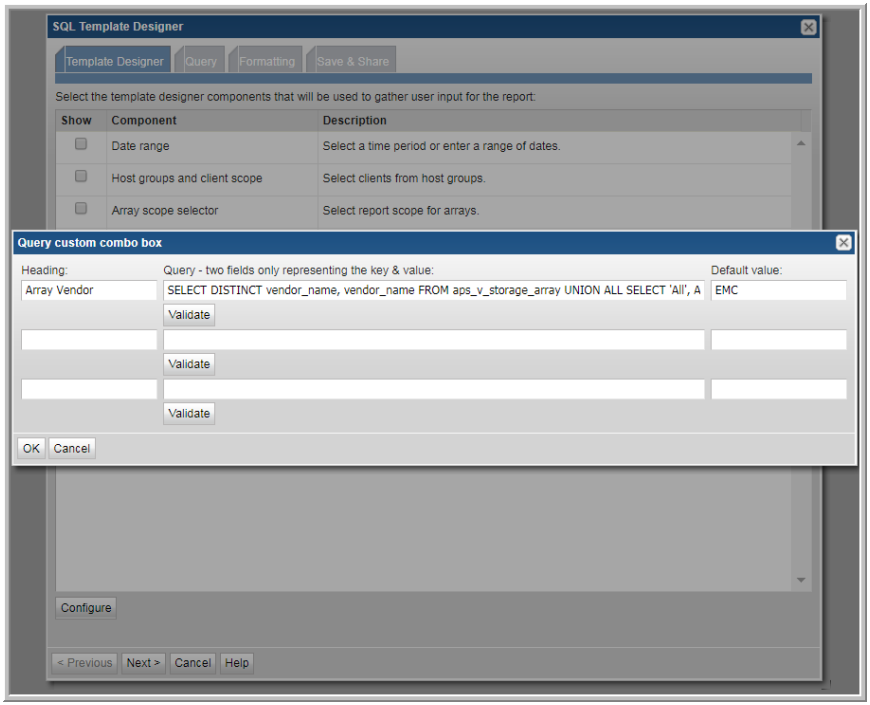
Configure an Area Chart SQL Template

An area chart depicts a time-series relationship. Unlike line charts, area charts can also visually represent volume. Information is graphed on two axes, using data points connected by line segments. Basic area charts depict the data using transparency to overlay the areas. Stacked area charts presents the data areas without overlaying - displaying part-to-whole relations by showing the constituent parts of a whole one over the other.

- See "Configure SQL Template Scope Selector Components" on page 356.
- See "Construct the SQL Query" on page 367.
- See "Number, Size, Date, and Time Formatting" on page 377.
- See "Format the SQL Template Output" on page 376.

Example of Area Chart SQL Template: Allocated Available

1. From the Template Designer tab, click Query custom combo box.



2. With the **Query custom combo box** selected, click **Configure**.
3. In the **Heading** field, enter the heading that will display in the report template's scope selector drop-down list. For example Array Vendor.
4. In the **Query** field, enter the SQL query that returns a list of values that will be available in the combo box:

```
SELECT DISTINCT vendor_name,vendor_name FROM aps_v_storage_array
 UNION ALL SELECT ' All', ' All' FROM dual ORDER BY 1
```

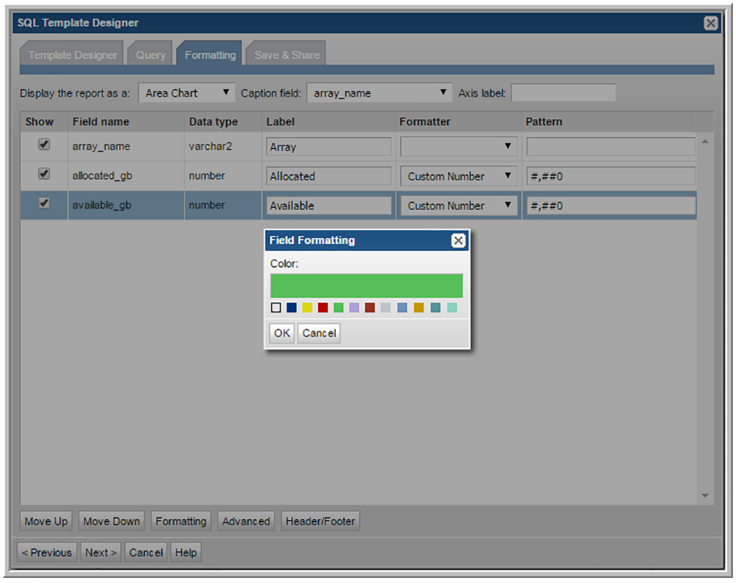
5. Click **Validate** and **OK**.
6. Click the **Query** tab and enter the following. This query will use the value selected in the populated combo box to retrieve the data to be rendered in the report. The custom combo variable is represented as ${queryCombo1} in the query.

```
SELECT
Array_name,
allocated_gb, available_gb
```

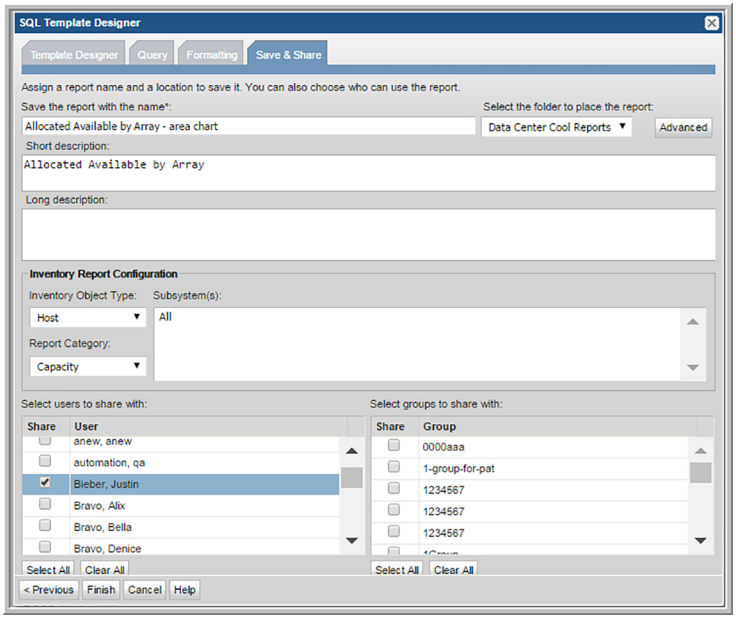
```
FROM aps_v_storage_array
WHERE
'${queryCombo1}' IN
  CASE
    WHEN '${queryCombo1}' NOT IN (' All') THEN
      CASE
       WHEN vendor_name = '${queryCombo1}' THEN '${queryCombo1}'

      END
   ELSE ' All'
END
ORDER BY available_gb DESC
```

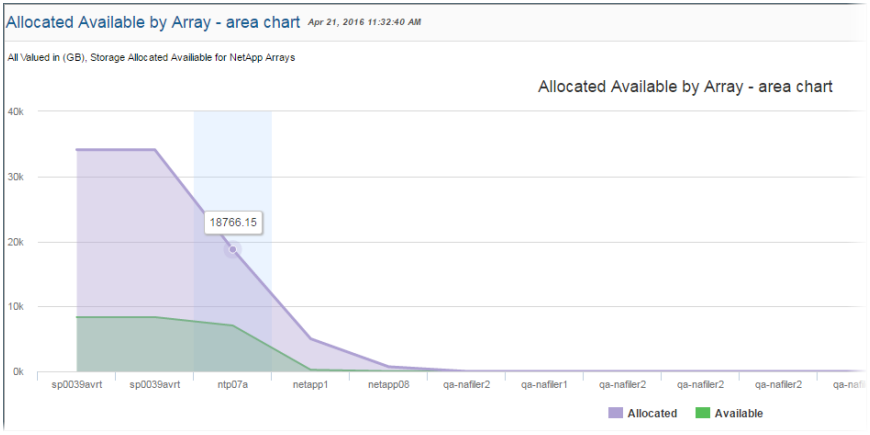
7. Click the **Formatting** tab.
8. In the **Formatting** window, select **Area Chart** or **Stacked Area Chart** from the **Display the report as a** drop-down list.
9. For the **Caption** field, select **array_name** from the drop-down list.
10. Select all the fields to be displayed.
11. Assign a color to each field that will be displayed as an area. In the example, assign a color to the fields **allocated_gb** and **available_gb**.



12. Assign a name for the report template and save it to a report menu group. You can also assign sharing if required.



13. Generate a report from the new template. The output will look similar to the following example:



Configure a Pie Chart SQL Template

A pie chart represents data with pie sectors. This type of chart has two main requirements:

- Select a character field from the Caption field drop-down list.

- For pie charts, use numeric fields as pie chart sectors.
  See "Example of a Pie Chart SQL Template: Overall Job Status Summary" on page 399.

A Donut Chart is similar to a Pie Chart.

See "Configure a Donut Chart SQL Template" on page 388.

Example of a Pie Chart SQL Template: Overall Job Status Summary

The following example creates a pie chart that represents the various status' for a job—such as, success, warning and failure—as sectors (slices) in a pie chart.

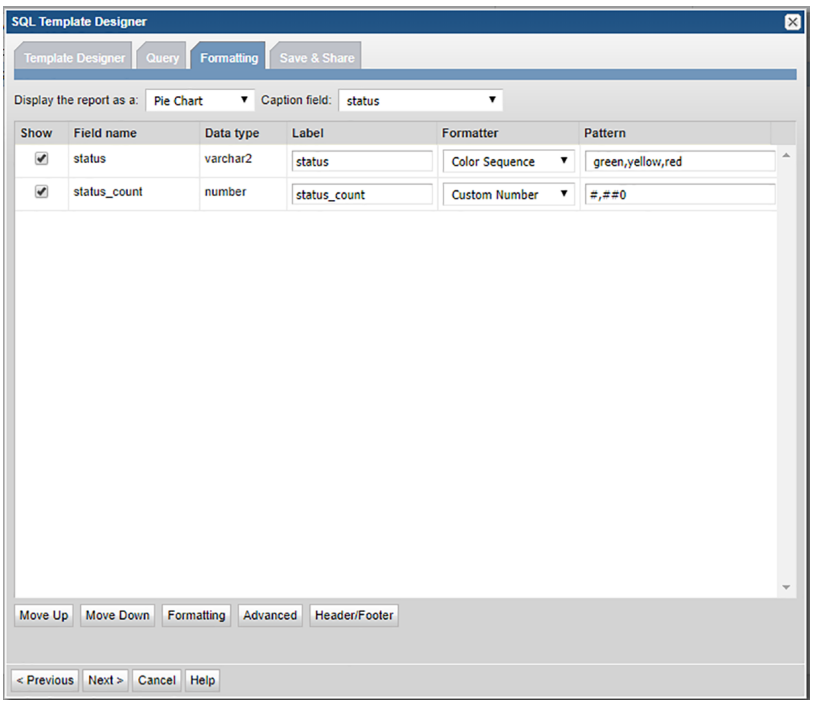
1. In the SQL Template Designer, in the **Template Designer** window, choose **Date range** and **Host groups and client scope**.
2. In the **Query** window, enter the following select statement and click **Validate Query**:

```
WITH t1 AS (
SELECT
0 seq, 'Success' status, count(job_id) status_count
FROM apt_v_job j
WHERE j.client_id IN (${hosts})
AND j.start_date BETWEEN  ${startDate} AND ${endDate}
AND j.summary_status=0
UNION
SELECT
1 seq, 'Partial' status, count(job_id) status_count
FROM apt_v_job j
WHERE j.client_id IN (${hosts})
AND j.start_date BETWEEN  ${startDate} AND ${endDate}
AND j.summary_status=1
UNION
SELECT
2 seq, 'Failed' status, count(job_id) status_count
FROM apt_v_job j
WHERE j.client_id IN (${hosts})
AND j.start_date BETWEEN  ${startDate} AND ${endDate}
AND j.summary_status=2
)
SELECT status, status_count
```

```
FROM t1
ORDER BY seq
```

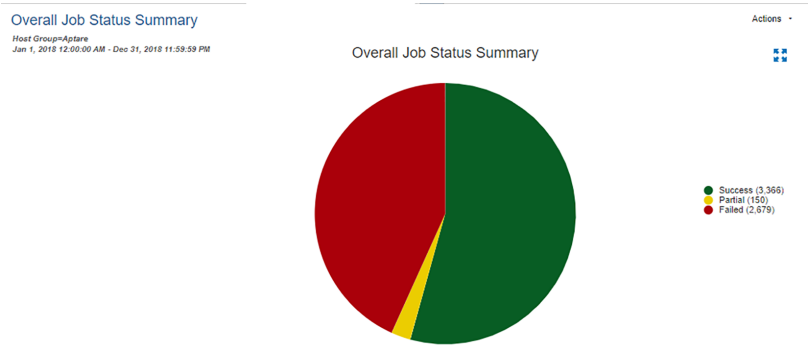
3. In the **Formatting** window, select **Pie Chart** from the **Display report as** drop-down list.
4. For the **Caption** field, select **status** from the drop-down list. Color sequence formatter must be selected on the **Caption** field.
5. Select all the fields to be displayed. For a pie chart, you need at least one field to be the caption and another field to be the sector.
6. Assign custom colors for your pie sectors. Specify the order and first few colors for pie chart sectors by choosing the Color Sequence as the Formatter. Enter the color in the associated Pattern field. Choose from Red, Yellow, Green, Blue, Black, White or Grey. For example, you can make a chart which depicts Success as Green, Partial as Yellow, and Failed as Red.

   Color sequence formatter must be selected on the Caption field. The color values are case-insensitive. If a non-supported color is entered, that color will be ignored. If more sections are present in the pie chart than the colors entered, a random color will be assigned to the additional sections.



For a pie chart, you need at least one field to be the caption and another field to be the sector.

7. Click **Next**, enter a report name and click **Finish**.
8. When you run this report, the output will look something like the following:



Configure a Pivot Table SQL Template

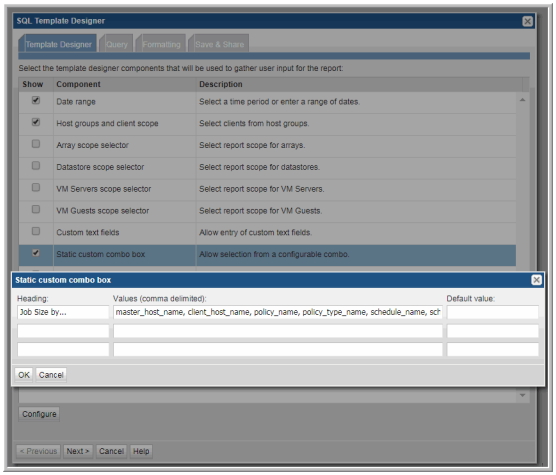
A pivot table provides a useful mechanism for aggregating and summarizing data. In effect, you are taking the flat rows of data and grouping them into a multi-dimensional representation of the data, enabling easier data analysis. When formatting a Pivot Table, only three fields are used (row, column, and data). Use the **Move Up** and **Move Down** feature to position the fields to easily select three fields.

- See “Example of a Pivot Table SQL Template: NetBackup Job Size” on page 401.
- See “Configure a Table in the SQL Template Designer” on page 405.

Example of a Pivot Table SQL Template: NetBackup Job Size

The following example illustrates a key SQL Template Designer formatting feature--Pivot Tables.

1. From the SQL Template Designer tabbed window, select **Date Range** and **Host Groups and Client Scope** and also **Custom Combo Box**.
2. Select **Static custom combo box**and click **Configure** at the bottom of the Query window.

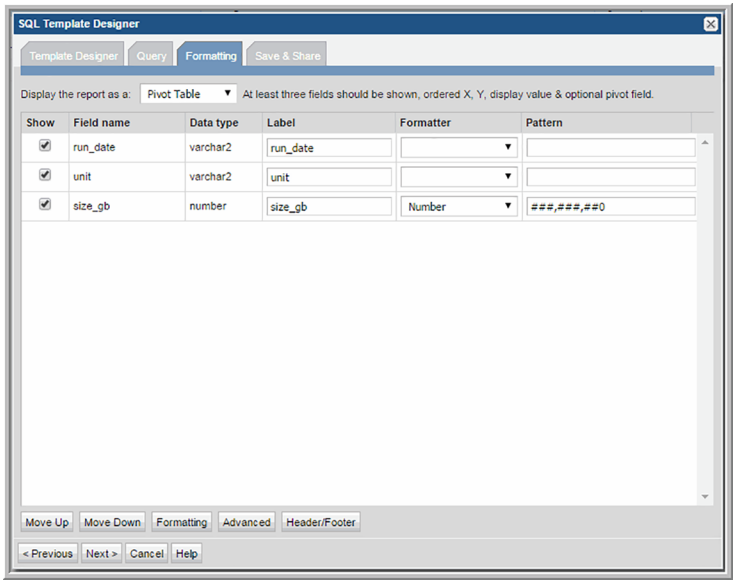


3. Enter the heading for the drop-down combo box--in this example, **Job Size by...**
4. Enter the following comma-separated list of values and then click **OK**: **primary_host_name**, **client_host_name**, **policy_name**, **policy_type_name**, **schedule_name**, **schedule_type_name**, **storage_unit_label**

   These will be the options that a user can select when generating a report.
5. Click on the **Query** tab and enter the following query:

```
SELECT to_char(start_date,'MM/DD/YY') run_date,
to_char(${freeCombo1}) unit,  trunc(sum(kilobytes/1024/1024))
size_gb
FROM apt_v_nbu_job_detail
WHERE start_date BETWEEN ${startDate} AND ${endDate}
AND client_id in (${hosts})
GROUP BY to_char(start_date,'MM/DD/YY'), ${freeCombo1}
```

6. Click **Validate Query** and then **Next**.
7. In the Formatting window, select **Pivot Table** and all three fields: **run_date, unit, and size_gb**



8. Save the report to a report menu group. Then, click on the saved report to generate it.

Configure a Stacked Bar Chart SQL Template

A stacked bar chart represents data with vertical bars. Rather than individual bars per data point, the values are stacked within a single bar. This type of chart can display bars only or it can be configured to include a line that will be charted on the right y-axis. Bar charts are available as two types:

- Dynamic: With this type of chart, you don't need to know how many distinct values are available. Typically, this means that you not select colors for the bars, but let the system select them. In some cases, colors are available for selection for a limited set of known values. A dynamic chart requires 1 Caption, 1 Bar and 1 Legend.
- Static: With this type of chart, you pre-define only the values and colors that you want represented in the chart - typically, a subset of the full set of values. A static bar chart requires 1 Caption and at least 1 Bar.
- See “Example of a Stacked Bar Chart SQL Template: Host CPU Performance” on page 404.
- See “Configure a Horizontal Bar Chart SQL Template” on page 392.
- See “Configure a Horizontal Stacked Bar Chart SQL Template” on page 393.

- See "Configure a Bar Chart SQL Template" on page 384.

Bar charts are all similar and rendering can be easily changed from one orientation to another (horizontal or vertical) or from stacked to individual bars. Once a SQL Template is developed for a bar chart, simply select the

Example of a Stacked Bar Chart SQL Template: Host CPU Performance

1. From the SQL Template Designer tabbed window, select **Date Range** and **Host Groups and Client Scope**.
2. Select **Static Custom Combo Box** and click **Configure** at the bottom of the Query window.
3. Enter the heading for the drop-down list in the Static Custom Combo box--for this example, **Group by:**
4. Enter the following comma-separated list of values and then click **OK**:

   Hour,Day,Week,Month,Quarter,Year

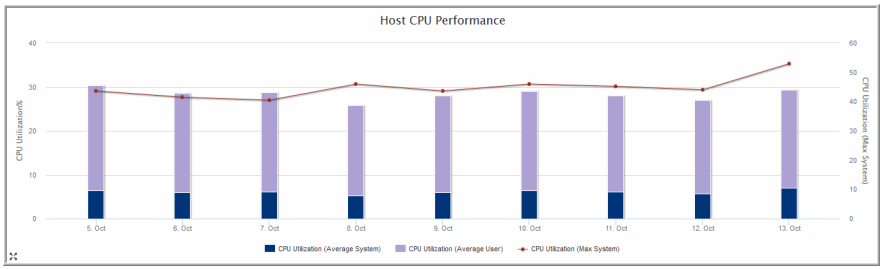
   These will be the options that a user can select when generating a report.
5. Click on the **Query** tab and enter the following query:

```
SELECT
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24',
'Day','DD','Week','WW','Month',
'MM','Quarter','Q','Year')) the_date,
avg(system_processing_time_pct) avg_system,
avg(user_processing_time_pct) avg_user,
max(system_processing_time_pct) max_system
FROM apt_v_host_cpu_log
WHERE log_date between ${startDate} AND ${endDate}
GROUP BY
trunc(log_date,DECODE('${freeCombo1}','Hour','HH24',
'Day','DD','Week','WW','Month',
'MM','Quarter','Q','Year'))
```

6. Click **Validate Query** and then **Next**.
7. In the Formatting window, select **Stacked Bar**and set the Caption field to **the_date** with the Axis label set to **CPU Utilization%**. To create a Dynamic Stacked Bar chart, select the **Caption** field and choose a different String field as a **Legend**. Legend can be selected from the **Formatter**.

8. Check all four fields: **the_date, avg_system, avg_user, and max_system.**
9. One by one, select the bar fields, **avg_system** and **avg_user**, and click the Formatting button (at the bottom of the window) to configure the color of the bar. And, enter a relevant label for each field.
10. Select the max_system field and click the Formatting button (at the bottom of the window) to configure the color of the line.
11. Save the report to a report menu group. Then, click on the saved report to generate it.

The output for the stacked bar chart will look something like this:



Configure a Table in the SQL Template Designer

A table is the most basic chart used to represent data. Use the **Move Up** and **Move Down** feature to position the fields to easily select three fields.

Example of a Table SQL Template: Host Group Membership List

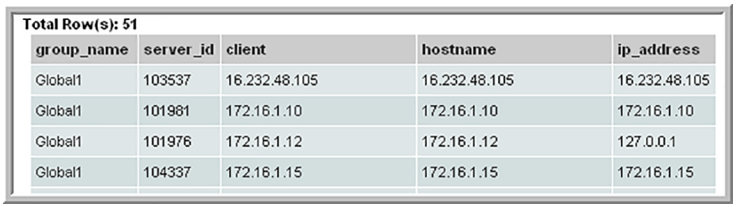
The following example reports on host group membership, in display name order.

1. In the SQL Template Designer, check only **Host Groups and Client Scope**.
2. In the Query window, enter the following select statement and click **Validate Query**:

```
SELECT g.group_name, s.server_id, s.display_name client,
s.hostname, s.ip_address
  FROM apt_v_group_member m, apt_v_group g, apt_v_server s
 WHERE g.group_id = m.group_id
   AND m.child_type = 2 --Server (excludes children that are host
 groups)
   AND m.child_id = s.server_id
```

```
   AND s.server_id IN (${hosts})
ORDER BY s.display_name, g.group_name
```

3. In the Formatting window, select the fields to be displayed. For this example, it makes sense to select all and display the report as a **Table**.
4. Click **Next**, enter a report name and select a Menu Group. Then, click **Finish**.
5. When you run this report, you can modify the scope to generate the report for a specific host group. The output will look something like this:



Total Row(s): 51

group_name	server_id	client	hostname	ip_address
Global1	103537	16.232.48.105	16.232.48.105	16.232.48.105
Global1	101981	172.16.1.10	172.16.1.10	172.16.1.10
Global1	101976	172.16.1.12	172.16.1.12	127.0.0.1
Global1	104337	172.16.1.15	172.16.1.15	172.16.1.15

6. Note that the column header “client” was substituted for the field name, as specified in the query. You can specify similar substitutions using the Formatting window in the SQL Template Designer.
7. Since this query is ordered by display name and then group name, it may be more reasonable to display the output in a similar fashion. Use the Formatting tab to re-arrange the table columns (Move Up, Move Down).

See “Format the SQL Template Output” on page 376.

Example of a Table SQL Template: Exposed Clients

The following example lists the clients within the report scope that have not been backed up within the selected time frame.

1. In the SQL Template Designer, check both **Date Range**and **Host Groups and Client Scope**.
2. In the Query window, enter the following select statement and click **Validate Query**:

```
SELECT apt_v_job.client_id, apt_v_job.client_name,
apt_v_job.server_id,
apt_v_job.server_name
  FROM apt_v_job
 WHERE apt_v_job.client_id IN (${hosts})
   AND apt_v_job.summary_status IN (0,1) -- Success or Warning
HAVING MAX(start_date) < ${startDate}
GROUP BY apt_v_job.client_id, apt_v_job.client_name,
```

```
apt_v_job.server_id,
apt_v_job.server_name
```

3. In the Formatting window, select the fields to be displayed. For this example, it makes sense to select all and display the report as a **Table**.
4. Click **Next**, enter a report name and click **Finish**.
5. When you run this report, specify the start and end dates or a time period. So, for example, you could select Last 72 Hours and it will be substituted for **startDate**, listed in the query. You can modify the scope to generate the report for a specific host group. The output will look something like this:

Total Row(s): 6

client_id	client_name	server_id	server_name
305726	fiji	305697	everest
305818	vietnam	305697	everest
305822	mercury	305697	everest
305823	kuwait	305697	everest
305852	malta	305697	everest
305870	uganda	305697	everest

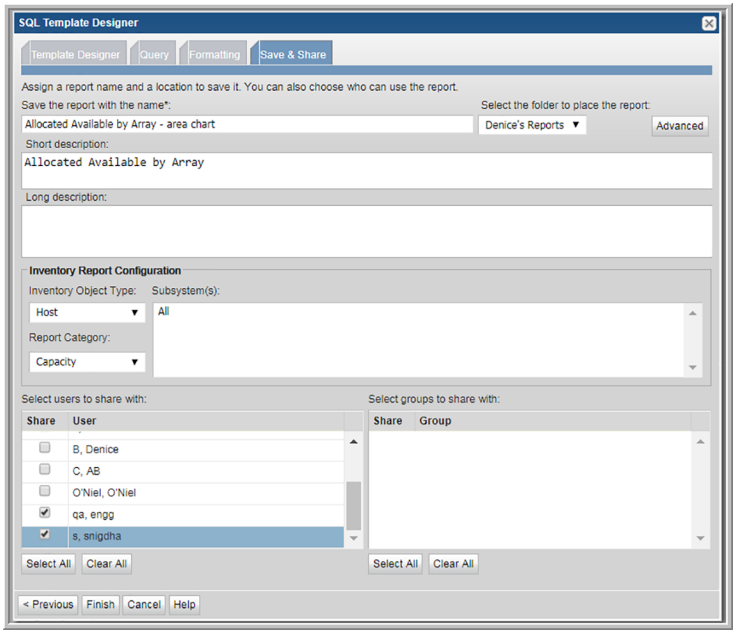
Save and Share Report Templates

When you save a report template, the following fields must be configured:

- Report Name (duplicate names are allowed)
- Folder (where the report will reside)
- Short Description (including key words that will aid searching)
- Long Description (including details that can also help in searching)

Optionally, you can assign:

- See “Inventory Report Configuration” on page 416.
  to set how report templates are categorized in the Inventory.
- Users and user groups for sharing the template.



Advanced SQL Report Template Options

- See "Creating Drilldowns" on page 408.
- See "Using SQL Templates for Drilldowns" on page 411.
- See "Create mouse-over hovers" on page 411.
- See "Add an External Reference" on page 413.

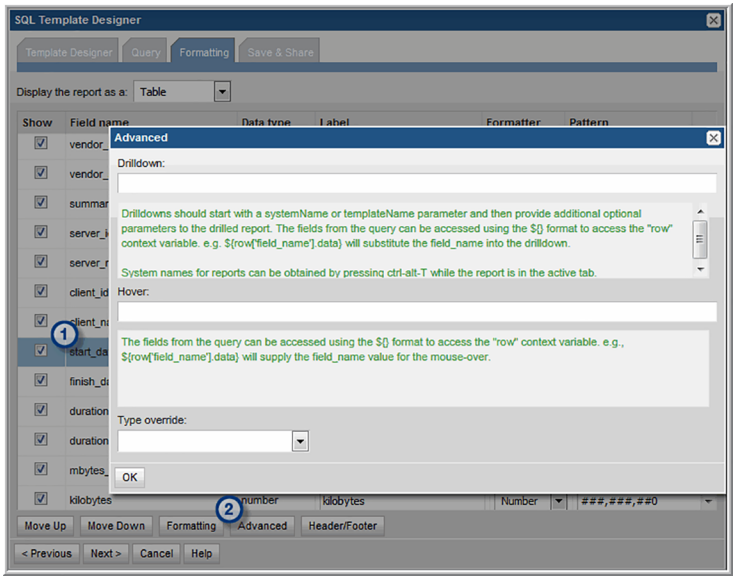
Creating Drilldowns

Just as the out-of-the-box reports contain links that drill down to other reports, Custom Reports can be configured to include drilldowns.

Fields from the query can be used in the drill down syntax.

In the SQL Template Designer, in the Formatting window:

1. Select the field that will become the drill down link.
2. Click **Advanced** to launch the Drilldown window where you will enter the syntax required to drilldown to the report.



3. In the Drilldown window, enter the details that the Portal will need to link to the report. The following example and the accompanying descriptions illustrate the required components for a drilldown specification.

   Examples

   ```
   systemName=displayServerDetail&serverId=${row['server_id'].data}
   templateInstanceId=400&serverId=${row['server_id']}
   ```

   Where:

systemName	Drilldowns should start with a systemName or templateName parameter and then provide additional optional parameters to the target report. The fields from the query can be accessed using the ${} format to access the "row" context variable. e.g. ${row['field_name'].data} will substitute the field_name into the drilldown. System names for reports can be obtained by pressing **CTRL+ALT+T** while the report is in the active tab. Use either systemName, templateInstanceId, or templateName in the syntax, based on the following use cases. **systemName** - The required prefix for drilldowns to factory-shipped, out-of-the-box reports that cannot be customized with either the SQL Template Designer or the Dynamic Template Designer. **templateInstanceId** - The required prefix for drilldowns to Report Templates created with the Dynamic Template Designer. **templateName** - The required prefix for drilldowns to Report Templates created with the SQL Template Designer.
<displayReportName>&<fieldName>	Insert the specific report **systemName**, **templateInstanceId,** or **templateName**along with the &<field Name> that is to be supplied by the parent report. To identify the system name or ID of an existing report, generate that report and in the active browser window type:**CTRL+ALT+T** Two types of names are displayed: ■ **System Name (systemName)**: For factory-shipped reports that cannot be customized, use this name for the **displayReportName.** ■ **Dynamic Template ID (templateInstanceId)**: For Report Templates created with the Dynamic Template Designer, use this number in place of **displayReportName**. Note that for a Report Template that was create with the SQL Template Designer, **CTRL+ALT+T** does not display the **templateName**. A Template Name must be configured. See “Using SQL Templates for Drilldowns” on page 411.
=${row[‘<field_name>’]}.data	The syntax required for the field name specification.

4. Click **OK** in the drill down window.

Using SQL Templates for Drilldowns

Report Templates that have been created with the SQL Template Designer can be configured to be the target of a drill down. In this example, a Report Template is created with the SQL Template Designer to serve as the parent report. Another Report Template is then created with the SQL Template Designer to be the drill down details.

1. Using the SQL Template Designer, create a parent report that includes a query that lists a group of hosts, using the apt_v_server published view (example query: select * from apt_v_server). For the purpose of this example, we'll call this the Parent Report.
2. Using the SQL Template Designer, create a second report that will be the Drilldown Report. For the purpose of a simplified example, you can use the same query that you used for the Parent Report.
3. Customize the Parent Report and in the Formatting tab, select the hostname field and click the **Advanced** button at the bottom of the SQL Template Designer window.

   See "Advanced SQL Report Template Options" on page 408.

   - Enter the following syntax in the Drilldown field and click **OK** before saving the template:

     ```
     templateId=<Drilldown_Report_Id>=${row['hostname']}.data
     ```

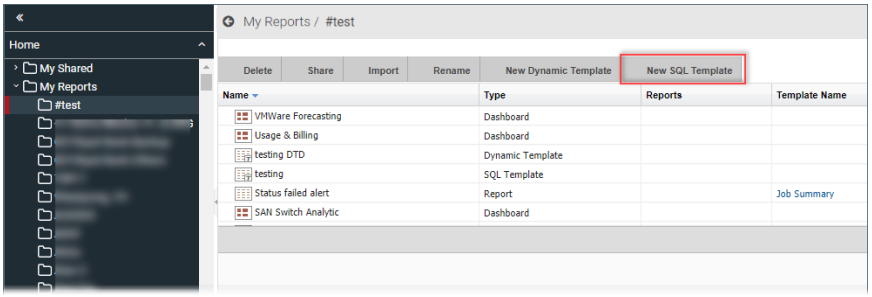
     Replace `<Drilldown_Report_Id>` with the actual report ID in the above syntax.

     Note: To get the template ID, run the report and use the shortcut CTRL+ALT+T.

Create mouse-over hovers

Custom Bar Charts, created with the SQL Template Designer, can be configured to have a mouse-over on the bars to display the charted values.

1. Select **Reports > My Reports> Your Custom Name Folder**.
2. Click the **New SQL Template** button.

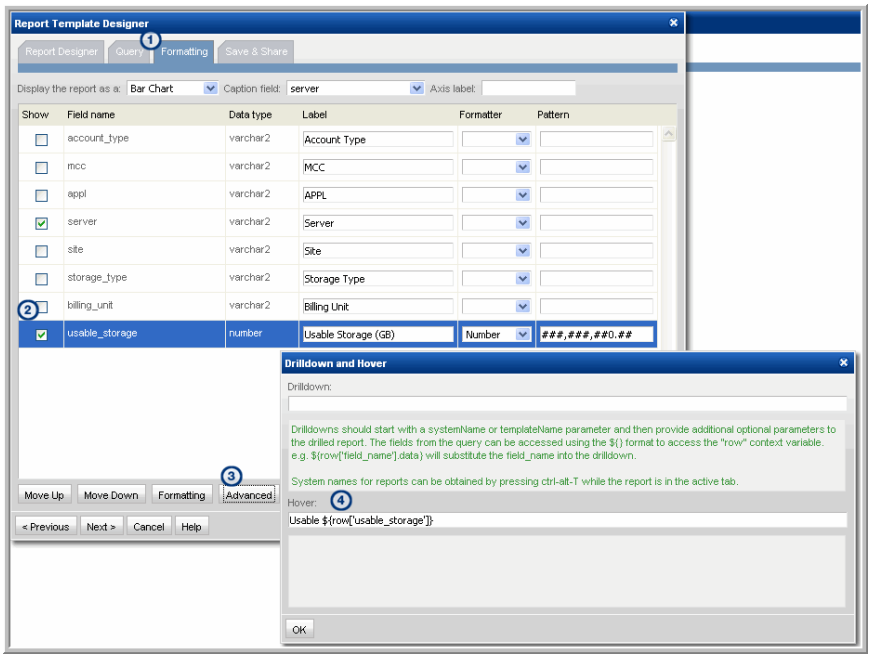


3. Click the **Formatting** tab in the SQL Template Designer window.
4. Select the field that will become the displayed as a bar in the chart.
5. Click **Advanced** to launch the **Drilldown and Hover** window where you will enter the syntax required for the mouse-over option.
6. In the Hover input box, enter the syntax for the value.

   The fields from the query can be accessed using the ${} format to access the row context variable.

   EXAMPLE

   ${row['field_name'].data} will supply the field_name value for the mouse-over.



Add an External Reference

Using the SQL Template Designer, you can create an tabular report template with a frame that can be populated from another URL providing HTML5 content. This special report template can then be placed on a dashboard and used in conjunction with other reports.

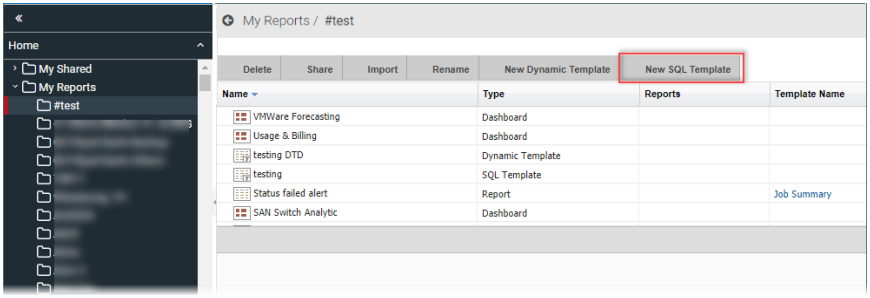
See “Create new dashboards and add reports” on page 197.

Note: Not all sites will display. If the site included in the iframe returns the X-Frame-Options header, the linked site may not allow its inclusion in the iframe.

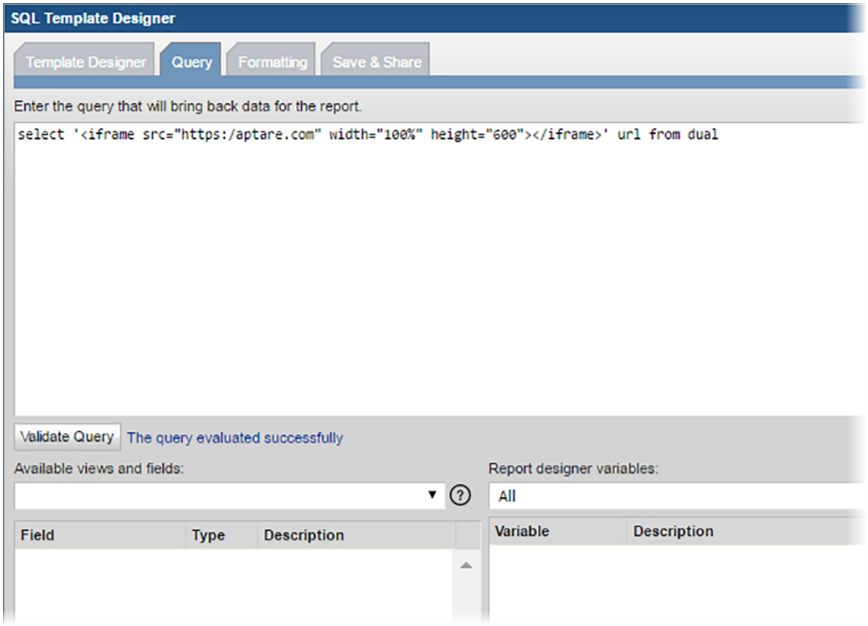
To create an external reference

1 Select **Reports > My Reports> Your Custom Name Folder**.

2 Click the **New SQL Template** button.



3 Click the **Query** tab. No action is required on the **Template Designer** tab.



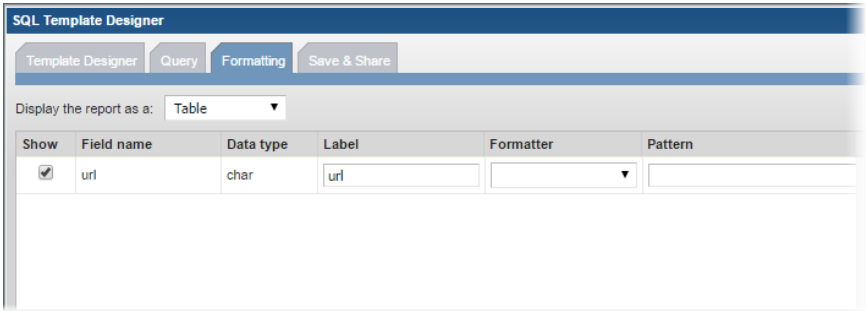
4 Enter the following select statement:

```
select '<iframe src="https:/aptare.com" width="100%"
height="600"></iframe>' url from dual
```

5 Replace **"https:/aptare.com"** with the destination URL.

6 Click **Validate Query**.

7 Click the **Formatting** tab. This page is pre-populated with the required field.

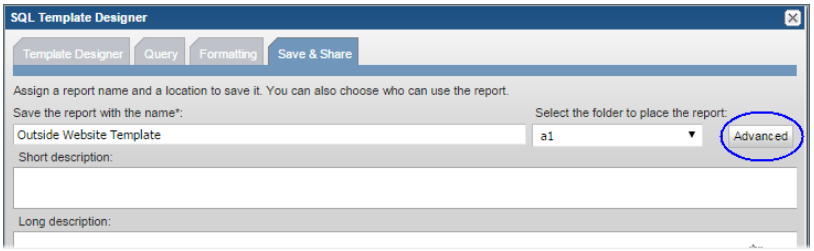


8 Verify the report is set to display as a **Table**.

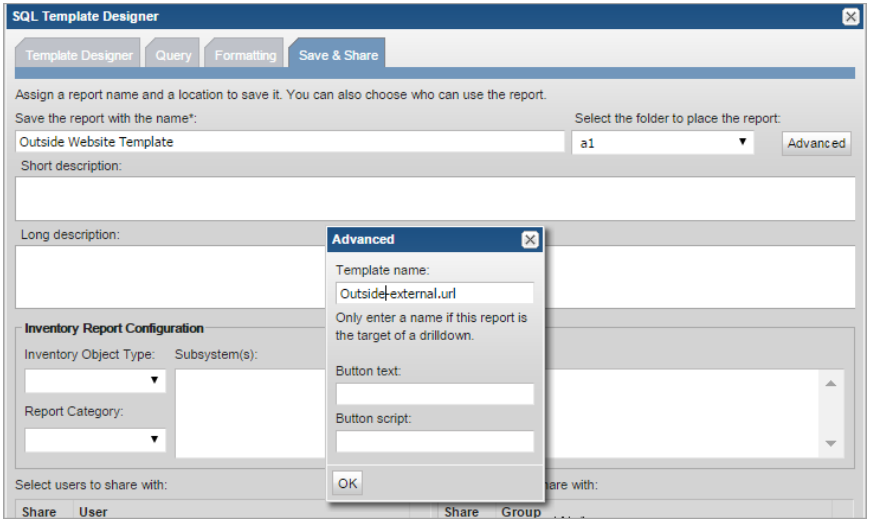
9 Click the **Save & Share** tab.

10 Enter a name for the template and select the folder to save it in. The template name must be unique.

11 Optionally, enter a description.



12 Click **Advanced**.



13 Enter **xxx-external.url** in the Template name field.

14 Replace **xxx** with a name.

15 Click **OK**.

16 Click **Finish** when back on the **Save & Share** tab.

17 Locate the report template in your specified folder and run it. You can add this to a dashboard.

Inventory Report Configuration

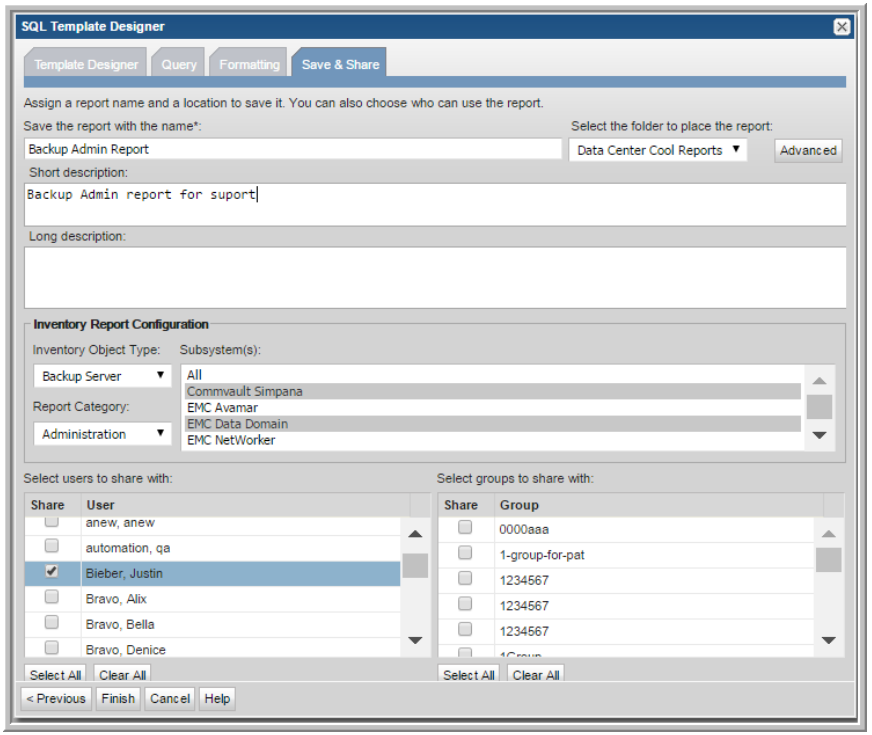
Assign the classification for report templates to display in the **Inventory**. This optional setting, can be assigned to report templates for arrays, hosts and backup servers. Assign an Inventory report category, object type and subsystem.

When displayed in the Inventory, templates are sorted by object type (Array, Backup Server or Host) and information categories such as performance, storage or forecasting. For templates created using the SQL Template Designer, you can assign values for these as you create the template or you can customize an existing one.

If a category has not been selected, these templates are displayed under the heading **Uncategorized**.

To assign the Inventory Report Configuration

1. Navigate to the **Save & Share** tab. You can either be creating a report template or editing one.
2. Select an **Inventory Object Type** and associated **Subsystem** vendor. You can also select **All**.



3. Select a **Report Category**. Choose from **Administration**, **Backup**, **Billing**, **Capacity**, **Forecasting, Management**, **Overview**, **Performance**, **Risk Mitigation**, or **Storage Optimization**.

About the Reports in the Inventory

- SQL Templates categorized for Hosts are displayed for all host groups and all individual hosts.
- SQL Templates categorized for Backup Servers and specific subsystems are displayed only for the groups of that subsystem and individual servers of that subsystem.

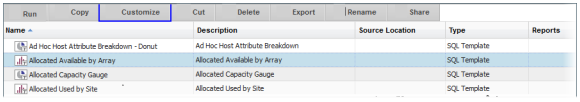
- SQL Templates categorized for Backup Servers and All subsystems are displayed for all groups of subsystems for Backup Servers and individual servers.
- SQL Templates categorized for Arrays are displayed for all array groups and all individual arrays.
- SQL Templates categorized for VM Guests are displayed for all VM Guest groups and all individual VM Guests.
- SQL Templates categorized for Datastores are displayed for all datastore groups and all individual datastores.
- SQL Templates categorized for VM Servers are displayed for all VM Server groups and all individual VM Servers.

Edit a Custom Report Template

Once you have saved a custom-designed report template, you can return to it to edit its configuration. For example, you can modify the SQL query or share it with additional users.

To edit a custom report template

1. Select the report template in the Reports tab.
2. Click **Customize** on the action bar. The relevant Template Designer is displayed to edit the template.



3. Refer to these sections for details about the various designers:

   See “Dynamic Template Designer Overview” on page 220.

   See the section called “CSV Format Specifications (Host Import)” on page 523.

Export/Import SQL Templates

In support of community report sharing, custom report templates can be exported to a file (.rtd) and then imported into another environment. This feature broadens the access to already developed, unique reports that may be ideal for your environment.

Note that IT Analytics Portal exports digitally signed templates and allows import of only of digitally signed templates.

Export a SQL Template

Note that IT Analytics Portal exports only digitally signed templates.

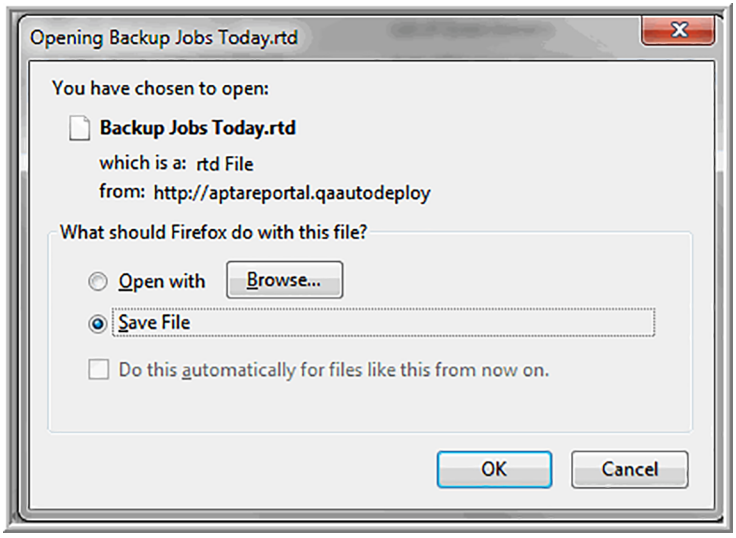
To export a custom-designed report template from your Portal to an .rtd file:

1. Create and save a custom report template using the instructions in the following.

   See "Dynamic Template Designer Overview" on page 220.

   See "SQL template designer overview" on page 351.

2. In the **Reports** content pane, select the report template and click **Export**in the Action bar above the View Panel.



3. Click **Save File** and **OK** to save this .rtd file so that it can be shared and imported by others.

See "Import a SQL Template" on page 419.

Import a SQL Template

Report templates are version-specific and are not backwardly compatible. Therefore, you cannot export a report template that was created in a later Portal software version and then import it into an earlier Portal version. The reverse, however, is supported--a template exported from an earlier version can be imported into a later Portal software version, but the template may experience some loss of formatting when imported.

Note that IT Analytics Portal allows import of only digitally signed templates.

To import a custom-designed report template, follow these steps:

1. Click **Reports**.
2. In the navigation panel on the left, select the folder where you want the imported report template to be saved.
3. Click **Import** in the Action bar above the View Panel.

   If the template you are importing is not digitally signed, the portal displays an error stating the template is not signed. See Enable import of unsigned report template to import the template.
4. In the **Report Template Import** window, click **Choose file** to find the saved template (`.rtd`) file on your local system. You can import multiple template files (`.rtd`) by consolidating them into a folder, zipping/compressing the folder and selecting the folder to import.
5. Select the `.rtd` file and click **OK**. The imported template or templates will appear in the folder you specified and now can be generated.

See "Export a SQL Template" on page 419.

Recommendation: Visit the Report Library for additional report templates: http://reportlibrary.aptare.com/

Enable import of unsigned report template

IT Analytics Portal allows import of only digitally signed report templates by default. To enable import of an unsigned report template, you must set the portal.allowUnsignedReportImport attribute to True from the portal.

To enable import of unsigned report template:

1. Click **Admin**.
2. In the navigation panel on the left, select **Advanced** > **System Configuration** and open the **Custom Parameters** tab.
3. Set the value of **portal.allowUnsignedReportImport** parameter to **True**.
4. Click **Save and Apply**.

   Import of unsigned report template is now enabled.

Chapter 10

Pipelined functions for report query building

This chapter includes the following topics:

- About Pipelined Functions
- General Functions
- APTlistOfDates
- aptStringConcat
- collectString
- getLicenseClientDetail
- getServerAttributeValue
- getObjectAttributeValue
- getChildServerGroupContextById
- getServerGroupContextById
- secsToHoursMinSecs
- Backup Manager Functions
- APTgetJobTypeName
- APTgetTapeDriveStatusName
- getFullPathname
- listJobSummaryAfterRestart

- listJobSummaryAfterRestartNBW
- listJobSummaryAfterRestart for NetWorker Backup Jobs
- listOfBackupWindowDates
- listOfBackupWindowDates (by Backup Window ID)
- Policy Auditing Functions
- listClientChanges
- listPathnameChanges
- listPolicyChanges
- listScheduleChanges
- Capacity Functions
- listChargebackCatByVOLSDetail
- listChargebackCatByNcVolDetail
- listChargebackCatByFSDetail (for HNAS)
- listChargebackCatByFSDetail (for EMC Isilon)
- listChargebackByLUNSummary
- listChargebackByLUNDetail
- listChargebackCatByLUNSummary
- listChargebackCatByLUNDetail

About Pipelined Functions

The product provides a set of published read-only database views of the underlying data within the database and related object model. The purpose of these views is to provide a read-only reporting or data extraction mechanism for advanced users of the product.

Typically, these views will be accessed via the SQL Template Designer, where SQL queries can be constructed to develop custom reports. Using these read-only views, you can write your own reports or stored procedure handlers to query the reporting database. These views provide a simple and fast mechanism to access a read-only view into the raw data maintained within the report database.

Pipelined functions can be used to query the database to further expand the capabilities of the custom report templates built with the SQL Template Designer.

General Functions

Function Name	Description
See "APTlistOfDates" on page 423.	Provides a list of dates based on the start date, end date, and grouping aggregation interval to a query.
See "aptStringConcat" on page 425.	Concatenates string values returned from a table.
See "getLicenseClientDetail" on page 426.	Determines backup license usage details. It returns the host ID, host name, backup vendor, host type, and the last updated date.
See "getServerAttributeValue" on page 426.	Returns a string--the host/server attribute value.
See "getObjectAttributeValue" on page 427.	Returns a string--the object's attribute value.
See "getChildServerGroupContextById" on page 428.	Provides the capability for getting data for a host group that is not at the top level of the host group hierarchy. This limits the scope of the query to a sub-group.
See "getServerGroupContextById" on page 429.	Provides the capability for getting specific host group data.
See "secsToHoursMinSecs" on page 429.	Useful for converting job duration to a readable format.

APTlistOfDates

This function enables you to provide a list of dates to a query.

```
FUNCTION APTlistOfDates(
startDateIN DATE,
endDateIN DATE,
groupByIN NUMBER)
```

This function returns a character string.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

Example Query

```
select * from table(rtd.APTlistOfDates(to_Date('10012008','MMDDYYYY'),
 to_date('10032008','MMDDYYYY'), 11))
```

The output from this example:

```
THE_DATE
--------------------
01-Oct-2008 00:00:00
02-Oct-2008 00:00:00
03-Oct-2008 00:00:00
```

The third argument--groupBy--controls the granularity of the results, as follows:

Value	Description
1	Group by 15 minutes
2	Group by 30 minutes
3	Group by 5 minutes.
4	Group by 10 minutes.
5	Group by 1 minute.
10	Group by hour
11	Group by day
12	Group by week
13	Group by month
14	Group by quarter
15	Group by year

aptStringConcat

Use this function in a SQL Template Designer query or the Dynamic Template Designer to concatenate string values returned from a table.

Example Query

```
select aptStringConcat(hostname) from apt_v_server where rownum < 10
```

This query in the SQL Template Designer will generate output similar to:



aptstringconcat(hostname)
aptaredev1,esx1,aptarexen1,aptaredev3.corp,flora,sdchtmp01,aptarew2003_old,hds-sun1,aptbuetest01

collectString

Use this function in a SQL Template Designer query or in the Method Designer to concatenate distinct values. This can be especially useful for including a sparkline chart in a tabular report.

```
SELECT rtd.collectString(SET(CAST(COLLECT(<column name>) AS
stringListType)), <delimeter>) from <table name>;
```

The following restrictions apply:

- While concatenating NUMBER fields, explicitly cast the column to VARCHARs() as shown in the following example. While casting, we need to specify the size of the column.
  Example:

  ```
  SELECT rtd.collectString(SET(CAST(COLLECT(cast(client_id as
  varchar2(10))) AS stringListType)),', ') name from apt_v_job;
  ```

- For STRING or NUMBER columns, this function is restricted to up to 512 characters.
  Example:

  ```
  SELECT rtd.collectString(SET(CAST(COLLECT(hostname) AS
  stringListType)),', ') name from apt_v_server;
  ```

getLicenseClientDetail

This function can be used to determine backup license usage details. It returns the host ID, host name, backup vendor, host type, and the last updated date.

```
select * from TABLE(rtd.getLicenseClientDetail)
```

getServerAttributeValue

This function returns a string--the host/server attribute value.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

```
getServerAttributeValue(
hostID,
attributeName)
```

Argument Name	Type
hostID	Number
attributeName	String in single quotes

Example Query

```
SELECT server_id, status
FROM (
SELECT server_id, rtd.getserverattributevalue(server_id,'Status')
Status
FROM apt_v_server s
)
WHERE status IS NOT NULL
ORDER BY server_id
```

The function listed above provides a reliable method of reporting on host/server attributes via the SQL Template Designer.

The apt_v_server_attribute is a dynamically created view, specific to your environment.

getObjectAttributeValue

This function returns a string--the object's attribute value.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

```
getObjectAttributeValue(
objectID,
attributeName,
attributeType)
```

Argument Name	Type
objectID	Number
attributeName	String in single quotes
attributeType	String in single quotes

Example Query

In the following SQL Template Designer query, it assumes that you have configured a Location attribute for your storage arrays and you want to return the location of the array's data center for your report.

```
SELECT
rtd.getObjectAttributeValue(sa.storage_array_id,'Location','A') AS
" Data Center Location",
sa.array_name AS "Array Name",
sa.allocated_gb AS "Allocated (GB)"
FROM aps_v_storage_array sa
WHERE
rtd.getObjectAttributeValue(sa.storage_array_id,'Location','A') IS
NOT NULL
```

The resultant report will display a table with three columns: Data Center, Array Name, and Allocated (GB).

getChildServerGroupContextById

This function provides the capability for getting data for a host group that is not at the top level of the host group hierarchy. This limits the scope of the query to a sub-group.

```
FUNCTION getChildServerGroupContextById(
groupID,
clientID,
depthLevel)
```

This function returns a character string.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

Specify the following arguments:

Argument Name	Type	Length	Description
groupID	Number	6	This is the highest level host group to be accessed.
clientID	Number	6	Specify a specific host ID.
depthLevel	Number		This number indicates the number of levels down the hierarchical host group tree you want the query to search.

Example Query

```
SELECT rtd.getChildServerGroupContextById(100209,server_id,3) FROM
apt_v_server
```

getServerGroupContextById

This function provides the capability for getting specific host group data.

```
FUNCTION getServerGroupContextById(
groupID,
clientID,
depthLevel)
```

This function returns a character string.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

Specify the following arguments:

Argument Name	Type	Length	Description
groupID	Number	6	This is the highest level host group to be accessed.
clientID	Number	6	Specify a specific host/server ID.
depthLevel	Number		This number indicates the number of levels down the hierarchical host group tree you want the query to search.

Example Query

```
SELECT rtd.getServerGroupContextById(100000,server_id,3) FROM
apt_v_server
```

secsToHoursMinSecs

This function is useful for converting job duration to a readable format. For example, 622 seconds would return 00:10:22.

```
secsToHoursMinSecs(totalSecs IN NUMBER)
```

Example Query

```
SELECT rtd.secsToHoursMinSecs (duration_secs) FROM apt_v_job
```

Backup Manager Functions

Function	Description
See "APTgetJobTypeName" on page 431.	Used to query database tables related to backup job details, such as apt_v_job and apt_v_nbu_job
See "APTgetTapeDriveStatusName" on page 432.	Use this function to access data from: apt_v_tape_drive. This function returns a character string.
See "getFullPathname" on page 433.	Use this function to access data from: apt_v_client_file. This function returns a character string.
See "listJobSummaryAfterRestart" on page 433.	Returns a list of NetBackup jobs. It could be used, for example, to determine the ultimate success for NetBackup jobs within a backup window. For example, if the backup window was set for 4:00 p.m. to 4:00 p.m. the next day, if the restart completed before the end of the backup window, it is considered to be successful.
See "listJobSummaryAfterRestart for NetWorker Backup Jobs" on page 436.	Returns a list of NetWorker jobs that do not have a backup window. It could be used, for example, to determine ultimate success for NetWorker jobs. If the job restart finishes within the timeframe (startDate - finishDate), it is considered successful.
See "listJobSummaryAfterRestartNBW" on page 435.	Returns a list of NetBackup jobs that do not have a backup window. It could be used, for example, to determine ultimate success for NetBackup jobs.
See "listOfBackupWindowDates" on page 438.	Two versions of this function enable the following functionality: ■ returns a list of backup windows explicitly supplied in the function ■ returns a list of backup windows gleaned from the windows defined in the Portal

Function	Description
See "listOfBackupWindowDates (by Backup Window ID)" on page 439.	Based on the backup window definition, the function returns a list of the adjusted start and finish dates.

APTgetJobTypeName

This function can be used to query database tables related to backup job details, such as apt_v_job and apt_v_nbu_job.

```
FUNCTION APTgetJobTypeName(
productType,
jobType,
vendorJobType)
```

This function returns a character string.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

Argument Name	Type	Length
productType	Number	2
jobType	Number	4
vendorJobType	Number	

Example Query 1

```
SELECT rtd.APTgetJobTypeName(j.product_type, j.job_type, NULL)
FROM apt_v_job j
```

Example Query 2

```
SELECT rtd.APTgetJobTypeName(j.product_type, j.job_type,
n.vendor_job_type)
FROM apt_v_nbu_job n, apt_v_Job j
WHERE j.job_id = n.job_id
```

The output from this example:

```
Full Backup
Appl Backup
Appl Backup
```

APTgetTapeDriveStatusName

Use this function to access data from:

```
apt_v_tape_drive
```

This function returns a character string.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

```
FUNCTION APTgetTapeDriveStatusName(
vendorProductType,
vendorDriveStatus)
```

Argument Name	Type	Length
vendorProductType	Number	
vendorDriveStatus	Character	1

Example Query

```
SELECT rtd.APTgetTapeDriveStatusName(d.product_type,
d.vendor_drive_status) FROM apt_v_tape_drive d
```

The output from this example:

```
Up
In-Use
In-Use
Mounting
```

getFullPathname

Use this function to access data from:

```
apt_v_client_file
```

This function returns a character string.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

```
FUNCTION getFullPathname(
clientID,
clientFileID)
```

Argument Name	Type	Length
`clientID`	Number	6
`clientFileID`	Number	10

Example Query

```
SELECT rtd.getFullPathname(d.client_id, d.client_file_id) FROM
apt_v_client_file d
```

listJobSummaryAfterRestart

This function returns a list of NetBackup jobs. It could be used, for example, to determine ultimate success for NetBackup jobs within a backup window. For example, if the backup window was set for 4:00 p.m. to 4:00 p.m. the next day, if the restart completed before the end of the backup window, it is considered to be successful.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

```
listJobSummaryAfterRestart(
```

```
startDate,
finishDate,
backupWindowId,
clientList,
excludeStatusList,
excludePolicyList,
excludeClientList)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector
backupWindowId	Numeric
clientList	numberListType Example: numberListType(1,2,3)
excludeStatusList	numberListType Example: numberListType(1,2,3)
excludePolicyList	numberObjectListType Example: numberObjectListType(numberObjectType(1), numberObjectType(2))
excludeClientList	numberObjectListType Example: numberObjectListType(numberObjectType(1), numberObjectType(2)))

Example Query

```
SELECT job_id, nbu_job_id, c.hostname client, s.hostname server,
start_date,
finish_date, kilobytes, DECODE(overall_status,0,
'Successful',1,'Partial',2,'Queued', 3,'Running',4,'Failed', NULL)
overall_status,
DECODE(was_restarted,0,'No','Yes') was_restarted, vendor_status,
orig_vendor_status,
file_pathlist, window_start_date, window_finish_date
FROM
TABLE(nbu_rtd.listJobSummaryAfterRestart(${startDate},${endDate},100000,${spHosts},
```

```
null,null,null)) t, apt_v_server s, apt_v_server c
WHERE t.server_id = s.server_id
AND t.client_id = c.server_id
ORDER BY t.start_date
```

listJobSummaryAfterRestartNBW

This function returns a list of NetBackup jobs regardless of a backup window. It could be used, for example, to determine ultimate success for NetBackup jobs. If the job restart finishes within the timeframe (startDate -finishDate), it is considered successful.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

```
listJobSummaryAfterRestartNBW
startDate,
finishDate,
clientList,
excludeStatusList,
excludePolicyList,
excludeClientList,
jobTypeList)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector
clientList	numberListType Example: numberListType(1,2,3)
excludeStatusList	numberListType Example: numberListType(1,2,3)
excludePolicyList	numberObjectListType Example: numberObjectListType(numberObjectType(1), numberObjectType(2))

Argument Name	Type
excludeClientList	numberObjectListType Example: numberObjectListType(numberObjectType(1), numberObjectType(2)))
jobTypeList	numberListType Example: numberListType(1,2,3)

Example Query

```
SELECT job_id, nbu_job_id, c.hostname client, s.hostname server,
start_date,
finish_date, kilobytes, DECODE(overall_status,0,
'Successful',1,'Partial',2,'Queued', 3,'Running',4,'Failed', NULL)
overall_status,
DECODE(was_restarted,0,'No','Yes') was_restarted, vendor_status,
orig_vendor_status,
file_pathlist, window_start_date, window_finish_date
FROM
TABLE(nbu_rtd.listJobSummaryAfterRestartNBW(${startDate},${endDate},${spHosts},
null,null,null,null)) t, apt_v_server s, apt_v_server c
WHERE t.server_id = s.server_id
AND t.client_id = c.server_id
ORDER BY t.start_date
```

listJobSummaryAfterRestart for NetWorker Backup Jobs

This function returns a list of NetWorker jobs within a backup window. It could be used, for example, to determine ultimate client backup success for NetBackup jobs. If the job restart finishes within the time frame (startDate - finishDate), it is considered successful.

Note: In a SQL Template Designer query, the function name must be prefaced with: **leg_rtd.**

```
listJobSummaryAfterRestart(
```

```
startDate,
finishDate,
backupWindowId,
clientList,
excludeStatusList,
excludePolicyList,
excludeClientList)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector
backupWindowId	If no backup window is defined, the default is 12 hours from the start date. To configure a Backup Window via the Portal: Admin > Reports > Backup Windows and click Add. To determine the backupWindowId, See "listOfBackupWindowDates (by Backup Window ID)" on page 439.
clientList	numberListType Example: numberListType(1,2,3)
excludeStatusList	numberListType Example: numberListType(1,2,3)
excludePolicyList	numberObjectListType Example: numberObjectListType(numberObjectType(1), numberObjectType(2))
excludeClientList	numberObjectListType Example: numberObjectListType(numberObjectType(1), numberObjectType(2)))

Example Query

Note: In the following sample Report Template query, spHosts is the same as Hosts, but is specifically for use in stored procedures.

```
SELECT job_id, c.hostname client, s.hostname server, sc.schedule_name,
 start_date, finish_date, kilobytes, DECODE(overall_status,0,
'Successful',1,'Partial',3,'Running',4,'Failed', NULL) overall_status,
DECODE(was_restarted,0,'No','Yes') was_restarted, vendor_status,
orig_vendor_status, client_resource_name, window_start_date,
window_finish_date
FROM
TABLE(leg_rtd.listJobSummaryAfterRestart(${startDate},${endDate},100000,
${spHosts},null,null,null)) t, apt_v_server s, apt_v_server c,
apt_v_leg_schedule sc
WHERE t.server_id   = s.server_id
  AND t.client_id   = c.server_id
  AND t.schedule_id = sc.schedule_id(+) ORDER BY t.start_date
```

listOfBackupWindowDates

Two versions of this function enable the following functionality:

- returns a list of backup windows explicitly supplied in the function
- returns a list of backup windows gleaned from the windows defined in the Portal -

See "listOfBackupWindowDates (by Backup Window ID)" on page 439.

Note: In the SQL Template Designer query, the function name must be prefaced with: **rtd.**

Explicitly Supply the Backup Windows

```
FUNCTION listOfBackupWindowDates(
startDate,
finishDate,
backupWindowListIN apt_BackupWindowListType)
```

Argument Name	Type
startDate	Date
finishDate	Date
backupWindowList	

Based on the backup window definition, the above function returns a list of the adjusted start and finish dates.

Example Query

```
SELECT start_date, finish_date
FROM TABLE(rtd.LISTOFBACKUPWINDOWDATES(TO_DATE('01012008','MMDDYYYY'),
 to_DATE('01072008','MMDDYYYY'),
apt_BackupWindowListType(
APT_BACKUPWINDOWTYPE('Mon', 40, 64),
APT_BACKUPWINDOWTYPE('Tue', 64, 88),
APT_BACKUPWINDOWTYPE('Wed', 88, 112),
APT_BACKUPWINDOWTYPE('Thu', 112, 136),
APT_BACKUPWINDOWTYPE('Fri', 136, 160),
APT_BACKUPWINDOWTYPE('Sat', 160, 184),
APT_BACKUPWINDOWTYPE('Sun', 16, 40)
 )
 ))
```

The output from this example:

```
12/31/2007 4:00:00 PM1/1/2008 3:59:59 PM
1/1/2008 4:00:00 PM1/2/2008 3:59:59 PM
1/2/2008 4:00:00 PM1/3/2008 3:59:59 PM
1/3/2008 4:00:00 PM1/4/2008 3:59:59 PM
1/4/2008 4:00:00 PM1/5/2008 3:59:59 PM
1/5/2008 4:00:00 PM1/6/2008 3:59:59 PM
1/6/2008 4:00:00 PM1/7/2008 3:59:59 PM
```

listOfBackupWindowDates (by Backup Window ID)

Retrieve Backup Windows (Defined in the Portal) with a Backup Window ID:

```
FUNCTION listOfBackupWindowDates(
startDate,
finishDate,
windowGroupIDIN apt_date_window.windows_group_id%TYPE)
```

Argument Name	Type
startDate	Date
finishDate	Date
windowGroupID	

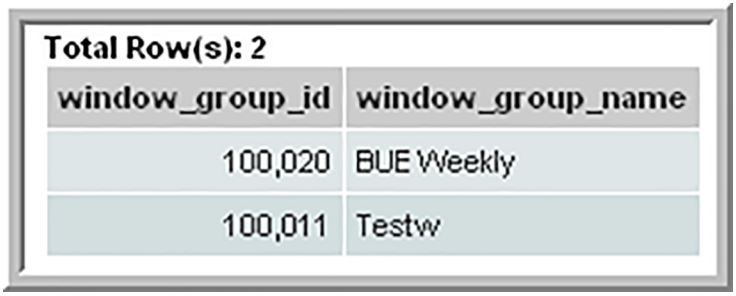
Based on the backup window definition, the above function returns a list of the adjusted start and finish dates.

Example Query

The following examples provide a basic idea of the function. These queries should serve as only an example of how the function might be incorporated into a more sophisticated query.

In a SQL Template Designer query, first determine the Backup Window IDs:

```
select window_group_id, window_group_name from apt_v_date_window
```

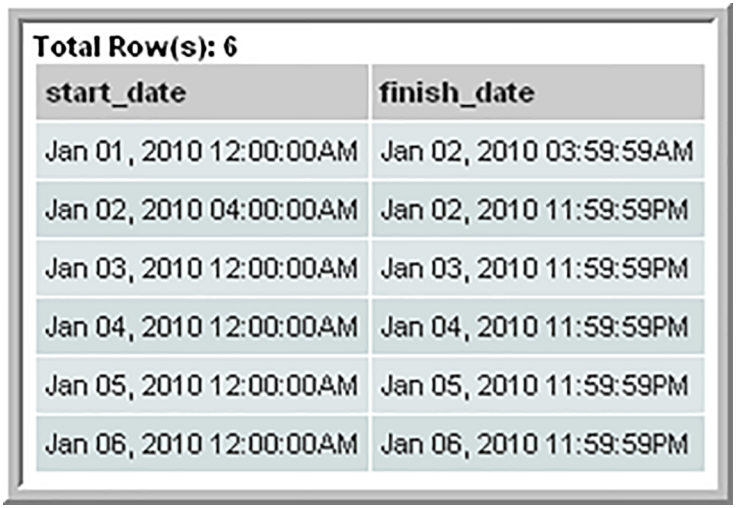


Total Row(s): 2

window_group_id	window_group_name
100,020	BUE Weekly
100,011	Testw

Then, use the Backup Window ID (window_group_id value from above example) in the following query:

```
select * from
TABLE(rtd.listOfBackupWindowDates(TO_DATE('01012010','MMDDYYYY'),
to_DATE('01072010','MMDDYYYY'), 100011 ))
```



Total Row(s): 6

start_date	finish_date
Jan 01, 2010 12:00:00AM	Jan 02, 2010 03:59:59AM
Jan 02, 2010 04:00:00AM	Jan 02, 2010 11:59:59PM
Jan 03, 2010 12:00:00AM	Jan 03, 2010 11:59:59PM
Jan 04, 2010 12:00:00AM	Jan 04, 2010 11:59:59PM
Jan 05, 2010 12:00:00AM	Jan 05, 2010 11:59:59PM
Jan 06, 2010 12:00:00AM	Jan 06, 2010 11:59:59PM

Policy Auditing Functions

Use the following functions to monitor NetBackup policy changes.

Function	Description
See "listClientChanges" on page 441.	Returns a list of NetBackup policies for which clients have changed.
See "listPathnameChanges" on page 442.	Returns a list of NetBackup policies for which path names have changed.
See "listPolicyChanges" on page 443.	Returns a list of NetBackup policies that have changed.
See "listScheduleChanges" on page 444.	Returns a list of NetBackup schedules that have changed.

listClientChanges

This function returns a list of NetBackup policies in which clients have changed.

Note: In the SQL Template Designer query, the function name must be prefaced with: **nbu_rtd.**

```
listClientChanges(
startDate,
finishDate,
clientList,
policyList,
groupID,
cascade)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector
clientList	numberListType Example: numberListType(1,2,3)
policyList	numberListType Example: numberListType(1,2,3)
groupID	Host group ID.
cascade	Specify 1 to designate that you want to cascade through the host group hierarchy; Otherwise, set to 0.

listPathnameChanges

This function returns a list of NetBackup jobs for which policy path names have changed.

Note: In the SQL Template Designer query, the function name must be prefaced with: **nbu_rtd.**

```
listPathnameChanges(
startDate,
finishDate,
clientList,
```

```
policyList,
groupID,
cascade)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector
clientList	numberListType Example: numberListType(1,2,3)
policyList	numberListType Example: numberListType(1,2,3)
groupID	Host group ID.
cascade	Specify 1 to designate that you want to cascade through the host group hierarchy; Otherwise, set to 0.

listPolicyChanges

This function returns a list of NetBackup jobs for which policies have changed.

Note: In the SQL Template Designer query, the function name must be prefaced with: **nbu_rtd.**

```
listPolicyChanges(
startDate,
finishDate,
clientList,
policyList,
groupID,
cascade)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector

Argument Name	Type
clientList	numberListType Example: numberListType(1,2,3)
policyList	numberListType Example: numberListType(1,2,3)
groupID	Host group ID.
cascade	Specify 1 to designate that you want to cascade through the host group hierarchy; Otherwise, set to 0.

listScheduleChanges

This function returns a list of NetBackup jobs for which policy schedules have changed.

Note: In the SQL Template Designer query, the function name must be prefaced with: **nbu_rtd.**

```
listScheduleChanges(
startDate,
finishDate,
clientList,
policyList,
groupID,
cascade)
```

Argument Name	Type
startDate	DATE with the value derived from the Scope Selector
finishDate	DATE with the value derived from the Scope Selector
clientList	numberListType Example: numberListType(1,2,3)
policyList	numberListType Example: numberListType(1,2,3)

Argument Name	Type
groupID	Host group ID.
cascade	Specify 1 to designate that you want to cascade through the host group hierarchy; Otherwise, set to 0.

Capacity Functions

Use the following functions to determine capacity chargebacks.

Function	Description
See "listChargebackCatByVOLSDetail" on page 445.	Provides the capability for listing capacity chargebacks for NetApp 7-Mode storage systems.
See "listChargebackCatByNcVolDetail" on page 447.	Provides the capability for listing capacity chargebacks for NetApp Cluster-Mode storage systems.
See "listChargebackCatByFSDetail (for HNAS)" on page 449.	Provides the capability for listing capacity chargebacks for Hitachi NAS (HNAS) storage systems.
See "listChargebackCatByFSDetail (for EMC Isilon)" on page 450.	Provides the capability for listing capacity chargebacks for EMC Isilon storage systems.
See "listChargebackByLUNSummary" on page 451.	Provides the capability for listing capacity chargebacks by LUNs.
See "listChargebackByLUNDetail" on page 453.	Provides the capability for listing capacity details for chargebacks by LUNs.
See "listChargebackCatByLUNSummary" on page 455.	Provides the capability for listing capacity chargebacks by LUNs by category; for example, for each storage tier, list the capacity chargeback. This function provides much of the same data that is displayed in the Chargeback Policy Capacity report.
See "listChargebackCatByLUNDetail" on page 456.	Provides the capability for listing capacity chargebacks by LUNs for array capacity (without regard to host usage). This function provides much of the same data that is displayed in the Chargeback Array Capacity report.

listChargebackCatByVOLSDetail

This function provides the capability for listing capacity chargebacks for the NetApp 7-Mode storage systems.

This function provides much of the same data that is displayed in the Chargeback Array Capacityreport.

Policy types supported for this pipelined function include:

- Array Name
- Array Type
- Array Family
- RAID Type
- Drive Speed
- Drive Capacity
- Drive Type
- Domain

Use the following syntax for this pipelined function.

```
FUNCTION listChargebackCatByVOLSDetail(
userID,
listOfDomains,
listOfArrays,
listOfPolicy)
```

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT *
FROM table (srm_rtd.listChargebackCatByVOLSDetail(100000,
numberListType(),
numberListType(),
numberListType()
))
```

Specify the following arguments:

Argument Name	Type	Length	Description
userId	Number		
listOfDomains			numberListType Example: numberListType(1,2,3)
listOfArrays			numberListType Example: numberListType(1,2,3)
listOfPolicy			numberListType Example: numberListType(1,2,3)

This function returns rows that contain: storage array ID, array name, policy ID, policy name, unit cost, total KB, unallocated KB, allocated KB, HDP capacity KB, non-HDP capacity KB, HDP free capacity KB, array group capacity KB, array group used KB, array group available KB, array group virtual capacity KB, array group PDEV capacity KB.

listChargebackCatByNcVolDetail

This function provides the capability for listing capacity chargebacks for NetApp Cluster-Mode storage systems. This function provides much of the same data that is displayed in the Chargeback Array Capacityreport.

Policy types supported for this pipelined function include:

- Array Name
- Array Type
- Array Family
- RAID Type
- Drive Speed
- Drive Capacity
- Drive Type
- Domain

Use the following syntax for this pipelined function.

```
FUNCTION listChargebackCatByNcVolDetail(
userID,
listOfDomains,
listOfArrays,
listOfPolicy)
```

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT *
FROM table (srm_rtd.listChargebackCatByNcVolDetail(100000,
numberListType(),
numberListType(),
numberListType()
))
```

Specify the following arguments:

Argument Name	Type	Length	Description
userId	Number		
listOfDomains			numberListType Example: numberListType(1,2,3)
listOfArrays			numberListType Example: numberListType(1,2,3)
listOfPolicy			numberListType Example: numberListType(1,2,3)

This function returns rows that contain: storage array ID, array name, policy ID, policy name, unit cost, total KB, unallocated KB, allocated KB, HDP capacity KB, non-HDP capacity KB, HDP free capacity KB, array group capacity KB, array group

used KB, array group available KB, array group virtual capacity KB, array group PDEV capacity KB.

listChargebackCatByFSDetail (for HNAS)

This function provides the capability for listing capacity chargebacks for Hitachi NAS (HNAS) storage systems. This function provides much of the same data that is displayed in the Chargeback Array Capacity report.

Policy types supported for this pipelined function include:

- Array Name
- Array Type
- Array Family
- Pool Name
- Tiering Policy
- Domain

Use the following syntax for this pipelined function.

```
FUNCTION listChargebackCatByFSDetail(
userID,
listOfDomains,
listOfArrays,
listOfPolicy)
```

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT *
FROM table (srm_rtd.listChargebackCatByFSDetail(100000,
numberListType(),
numberListType(),
numberListType()
))
```

Specify the following arguments:

Argument Name	Type	Length	Description
userId	Number		
listOfDomains			numberListType Example: numberListType(1,2,3)
listOfArrays			numberListType Example: numberListType(1,2,3)
listOfPolicy			numberListType Example: numberListType(1,2,3)

This function returns rows that contain: storage array ID, array name, policy ID, policy name, unit cost, total KB, unallocated KB, allocated KB, HDP capacity KB, non-HDP capacity KB, HDP free capacity KB, array group capacity KB, array group used KB, array group available KB, array group virtual capacity KB, array group PDEV capacity KB.

listChargebackCatByFSDetail (for EMC Isilon)

This function provides the capability for listing capacity chargebacks for EMC Isilon storage systems. This function provides much of the same data that is displayed in the Chargeback Array Capacity report.

Policy types supported for this pipelined function include:

- Array Name
- Array Type
- Array Family
- Domain

Use the following syntax for this pipelined function.

```
FUNCTION listChargebackCatByFSDetail(
userID,
listOfDomains,
listOfArrays,
listOfPolicy)
```

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT *
FROM table (srm_rtd.listChargebackCatByFSDetail(100000,
numberListType(),
numberListType(),
numberListType()
))
```

Specify the following arguments:

Argument Name	Type	Length	Description
userId	Number		
listOfDomains			numberListType Example: numberListType(1,2,3)
listOfArrays			numberListType Example: numberListType(1,2,3)
listOfPolicy			numberListType Example: numberListType(1,2,3)

This function returns rows that contain: storage array ID, array name, policy ID, policy name, unit cost, total KB, unallocated KB, allocated KB, HDP capacity KB, non-HDP capacity KB, HDP free capacity KB, array group capacity KB, array group used KB, array group available KB, array group virtual capacity KB, array group PDEV capacity KB.

listChargebackByLUNSummary

This function provides the capability for listing capacity chargebacks by LUNs.

```
FUNCTION listChargebackByLUNSummary(
userID,
listOfHosts,
listOfDomains,
listOfArrays,
listOfHostGroups,
cascade)
```

This function returns a list of hosts and related chargeback policy data, similar to what is displayed in the Chargeback By Host report.

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT * FROM table (srm_rtd.listChargebackByLUNSummary(100000,
numberListType(),
numberListType(),
numberListType(),
numberListType(${serverGroups}),
1))
```

Specify the following arguments:

Argument Name	Description
userId	Number
listOfHosts	numberListType Example: numberListType(1,2,3)
listOfDomains	numberListType Example: numberListType(1,2,3)
listOfArrays	numberListType Example: numberListType(1,2,3)
listOfHostGroups	numberListType Example: numberListType(1,2,3)

Argument Name	Description
cascade	Specify 1 to designate that you want to cascade through the host group hierarchy; Otherwise, set to 0.

This function returns rows that contain: host ID, host name, policy ID, policy name, unit cost, total GB, total cost.

listChargebackByLUNDetail

This function provides the capability for listing capacity details for chargebacks by LUNs.

```
FUNCTION listChargebackByLUNDetail(
userID,
listOfHosts,
listOfDomains,
listOfArrays,
listOfHostGroups,
cascade,
listOfPolicy)
```

This function returns a list of each host LUN chargeback policy. This is the detail that would be available as a drill down from the Chargeback By Host report.

See “listChargebackByLUNSummary” on page 451.

Chargeback in Cluster and Virtualization Environments

- In a cluster environment, if multiple hosts share a LUN, only one host will be charged.
- In virtualization environments, all VMs share the same capacity from the ESX server, so this function will return an estimation (calculated) for each LUN and for the ESX server, it will return the total capacity of the LUN.
- The following example shows the ESX server last in the list.

```
Host Name      LUN ID                                    Capacity

HQmkting1      133200                                    20
```

```
HQsales1        133200                                  10

HQfinance1      133200                                  10

HQesx10         133200                                  510
```

Example Query

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

```
SELECT *
FROM table (srm_rtd.listChargebackByLUNDetail(100000,
numberListType(),
numberListType(),
numberListType(),
numberListType(${hostGroups}),
1,
numberListType()
))
```

Specify the following arguments:

Argument Name	Description
userID	Number
listOfHosts	numberListType Example: numberListType(1,2,3)
listOfDomains	numberListType Example: numberListType(1,2,3)
listOfArrays	numberListType Example: numberListType(1,2,3)
listOfHostGroups	numberListType Example: numberListType(1,2,3)

Argument Name	Description
cascade	Specify 1 to designate that you want to cascade through the host group hierarchy; Otherwise, set to 0.
listOfPolicy	numberListType Example: numberListType(1,2,3)

This function returns rows that contain: host ID, host name, policy ID, chargeback policy name, storage array ID, storage array name, LUN ID, LUN name, unit cost, total GB, total cost.

listChargebackCatByLUNSummary

This function provides the capability for listing capacity chargebacks by LUNs by category; for example, for each storage tier, list the capacity chargeback. This function provides much of the same data that is displayed in the Chargeback Policy Capacityreport.

```
FUNCTION listChargebackCatByLUNSummary(
userID,
listOfDomains,
listOfArrays)
```

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT * FROM table (srm_rtd.listChargebackCatByLUNSummary(100000,
numberListType(),
numberListType()
))
```

Specify the following arguments:

Argument Name	Type	Length	Description
userId	Number		

Argument Name	Type	Length	Description
listOfDomains			numberListType Example: numberListType(1,2,3)
listOfArrays			numberListType Example: numberListType(1,2,3)

This function returns rows that contain: policy ID, chargeback policy name, policy ID, unit cost, total KB, unallocated KB, allocated KB, HDP capacity KB, non-HDP capacity KB, HDP free capacity KB, array group capacity KB, array group used KB, array group available KB, array group virtual capacity KB, array group PDEV capacity KB.

listChargebackCatByLUNDetail

This function provides the capability for listing capacity chargebacks by LUNs for array capacity (without regard to host usage). This function provides much of the same data that is displayed in the Chargeback Array Capacityreport.

```
FUNCTION listChargebackCatByLUNDetail(
userID,
listOfDomains,
listOfArrays,
listOfPolicy)
```

Note: In the SQL Template Designer query, the function name must be prefaced with: **srm_rtd.**

Example Query

```
SELECT *
FROM table (srm_rtd.listChargebackCatByLUNSDetail(100000,
numberListType(),
numberListType(),
numberListType()
))
```

Specify the following arguments:

Argument Name	Type	Length	Description
userId	Number		
listOfDomains			numberListType Example: numberListType(1,2,3)
listOfArrays			numberListType Example: numberListType(1,2,3)
listOfPolicy			numberListType Example: numberListType(1,2,3)

This function returns rows that contain: storage array ID, array name, policy ID, chargeback policy name, unit cost, total KB, unallocated KB, allocated KB, HDP capacity KB, non-HDP capacity KB, HDP free capacity KB, array group capacity KB, array group used KB, array group available KB, array group virtual capacity KB, array group PDEV capacity KB.

Chapter 11

Alert configuration

This chapter includes the following topics:

- Overview
- Alert rules
- Managing alert policies
- ServiceNow: Job finalized alert rule
- Managing alert notifications
- Report-Based alerts
- Syslog configuration
- About using SNMP
- Detect alerts in the inventory
- Suppress alert notifications

Overview

IT Analytics empowers you to intelligently and pro-actively ensure operational wellness. Through real-time symptom detection and notification, you can quickly spot problems across your datacenter, rapidly identify their causes, and minimize service degradation and disruption. Alert policies enable you to define watch conditions within your datacenter, and create custom alerts to trigger based on set thresholds and time periods. By defining thresholds, you can pro-actively manage situations before they occur and take action.

Watch this video on how to configure the Alert Notifications, monitor policies usage using Alert Policy Administration, and much more.

http://video.symantec.com/services/player/bcpid292374537001?bckey=AQ~~,AAAABuIiy9k~,I8Bhas-Vwr9zYL9V36WFi86fR_NoePscn&bctid=6301526630001

Alert rules

IT Analytics provides a pre-defined set of rules to examine common trouble spots within your enterprise. These rules serving as templates, include configurable parameters to isolate specific conditions relevant for your environment. Once configured, you can save the custom instance of the rule and create the Alert Policy. While various use cases will drive how you configure a rule, the goal is to help you quickly and pro-actively identify problem areas and trends that require attention, before they become unmanageable. Categories to alert on include:

- Data Collection
- Data Protection
- Performance
- Storage
- Storage Forecast
- System Administration
- Virtualization
- Virtualization Forecast

These rules are monitored in real-time as the data is collected or on a scheduled basis if the condition is time-based. Rule availability is based on your product installation.

Using Alert rules to create a policy, 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 criteria whenever required.

See *Alert reports* section in the *Report Reference Guide* for a variety of ways to examine the data you have selected to monitor.

Scheduled Rules vs. Run Time Rules

Each Alert Rule dictates the schedule for a policy and they are either time-based or dynamic.

Some rules are time-based, because they examine a group of events and have a pre-defined schedule (you can modify the schedule). The data collected requires

a time period to extract meaningful values. For example, to derive a percentage of failures, you must collect more than one backup job. It is the time period in between runs that is examined for the conditions you set in the Alert Policy.

Some rules are dynamic - every time data is collected that matches the criteria set in the Alert Policy, the alert is triggered. These real-time Alert Rules deliver instant results using the notification method defined in the policy.

Alert rules sample set

The following table lists a sample set of Alert rules by category. This is not a complete list of rules.

Table 11-1 Alert rules sample set

Category	Rule	Description
Data Collection	Data Collector Upgrade Failure	Alerts on selected data collectors when they have failed to upgrade.This is applicable for both upgrade manager and .jar.
Data Collection	Data Collector Offline	Alerts if the selected data collectors are in an offline state. Offline indicates the collector is shut down, or the collector cannot connect to the Data Receiver.

Table 11-1 Alert rules sample set *(continued)*

Category	Rule	Description
Data Protection	High Job Failure Rate	Monitors a high job failure percentage for any primary server belonging to any product. A specific job type can also be selected. This is a schedule-based rule where you set a schedule to run the rule using a CRON expression. The last run date will be saved in the database and every time when the rule is set to run, only the jobs finished between last run date of the rule and the present date will be considered. For the first run time, the date range for jobs will be last 2 hours.
Data Protection	Primary Server Connection	Alerts on Primary Servers when there are any connection issues. **Note:** This rule is only applicable if Host Collection is setup on the selected Primary Server.
Data Protection	Job Finalized	Alerts on any job of the selected job type, product and primary server which has completed with a specified Exit Status. Job Type, Primary Server and Product are defined by in the Alert Policy and Exit Status is part of the scope.

Table 11-1 Alert rules sample set *(continued)*

Category	Rule	Description
Data Protection	NetBackup Catalog Not Backed Up	Alerts when a NetBackup Catalog is not backed up for a Primary Server since the last run. For first time configurations, the catalog backup will be checked for last 2 days on the selected Primary Servers.
Performance	LUN Performance	Alert on selected Logical Units (LUNs) when their performance (I/O, Throughput, etc.) reach a specific threshold.
Performance	Switch FC Port Performance	Alerts on selected Switch Fibre Channel (FC) Ports when their performance (receive rate, transmit rate, etc.) reach specific thresholds.
Storage	LUN Capacity	Alerts on selected Logical Units (LUNs) when their used capacities reach a specific threshold.
Storage	Storage Pool Capacity	Alerts on selected Storage Pools when their capacity reaches specific thresholds for areas such as Array Group Allocated, Data Reduction, Raw Used, or Thin Pool Allocated.
Storage Forecast	Storage Pool Capacity Forecast	Alerts on selected Storage Pools when their used capacity forecast reaches specific percentage thresholds.

Table 11-1 Alert rules sample set *(continued)*

Category	Rule	Description
System Administration	License Usage	Alerts on usage overage percentage and approaching expiration dates for licensed modules. See "License Usage Alert Rule" on page 477.
Virtualization	VM Datastore Capacity	Alerts on selected VMWare Datastores when their used capacity and used capacity forecast reaches specific thresholds.
Virtualization	VMWare VM Guest Disk	Alerts on selected VMWare VM Guests when their used capacity reaches specific thresholds.
Virtualization Forecast	VMWare Datastore Capacity Forecast	Alerts on selected VMWare Datastores when their used capacity forecast reaches specific percentage thresholds.

Managing alert policies

Alert policies enable you to define watch conditions within your datacenter, and create custom alerts to trigger based on set thresholds and time periods.

Policy Administration includes the following.

- See "Add/Edit an Alert Policy" on page 463.
- See "Copying an Alert Policy" on page 480.
- See "Deleting an Alert Policy" on page 604.

Add/Edit an Alert Policy

Alert policies are intended to be independently activated and alerted on. Once you choose what to watch in your datacenter, you can receive actionable information through alert notifications. This same information can be viewed on-demand through

the different Alert reports - where events can be captured and drilled into for root-cause analysis. The following covers an overview of the steps:

1. Select an alert rule.
2. Name the policy and define the time intervals.
3. Define the scope to monitor.
4. Set the threshold for the alert rule.

See “Select an Alert Rule” on page 464.

See “Name the Policy and Define Time Intervals” on page 465.

See “Define the Scope to Monitor” on page 470.

See “Set the Threshold for an Alert Rule” on page 476.

Select an Alert Rule

Choosing the Alert Rule is the first step in creating an Alert Policy. Alert Rules serve as templates and once you include the configurable parameters to isolate specific conditions, save the custom instance of the rule and create the Alert Policy.

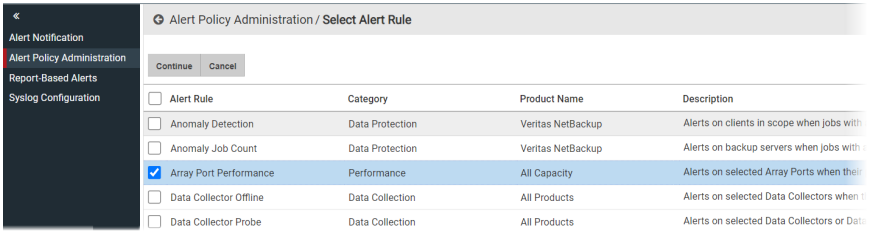
See “Alert rules” on page 459.

1. Navigate to **Alerts > Alert Policy Administration**.
2. Specify **Domain** from the list.
3. Click **Add**.

Note: User can edit and then select multiple alerts and add to **Alert Notification**.

IT Analytics provides a pre-defined set of monitoring rules to examine areas within your enterprise to monitor. These rules include configurable parameters to isolate specific conditions relevant for your environment. Many Alert Policies can be created from each Alert Rule.

Once the rule is selected, customize the Policy details to define 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.



Name the Policy and Define Time Intervals

To start creating the Alert Policy, assign a name, define the frequency of the symptom you are monitoring, and modify the schedule if required.

Table 11-2 Name the Policy and define Time Intervals

Field	Description
Alert Policy Name	Enter a name for the Alert Policy. This a required field.
Active	Click toggle button to activate the policy. To avoid un-necessary notifications, deactivate Alert Policies when there are situations such as planned outages or scheduled maintenance. You can also suppress alerts from the Alert Details report. See the *Report Reference Guide* for details in this report.

Table 11-2 Name the Policy and define Time Intervals *(continued)*

Field	Description
Symptom Frequency	

Table 11-2 Name the Policy and define Time Intervals *(continued)*

Field	Description
	It refers to the frequency at which the alert is generated. The options are explained below: ■ **Every time**: This is the default setting and indicates that the alert will be generated every time alert conditions are met for the selected scope. ■ **Consecutive X**: This symptom frequency implies that the alert will be generated if the alert condition is met for *X* times. For Example: In Job Finalized rule policy, the scope has condition for exit status 99 and symptom frequency as 2, then an alert is generated if there are 2 or more than 2 consecutive failed jobs with the error code 99. Host Job Exit Status Outcome Client A 99 - Client A 0 - Client A 99 - Client A 99 Alert Generated Client B 99 - Client B 0 - Client A 99 Alert Generated Client A 156 - Client A 99 - ■ **X out of Y**: It implies that the an alert will be generated if there are at least *X* total alerts in last *Y* attempts. Maximum value for *Y* can be 10. For example if symptom frequency is 3 out of 5, then the following is true: ■ Always last *Y* attempts are analyzed. In this example, last 5 attempts will be looked back. ■ No alert is generated until 5 attempts are complete (since the value of *Y* is 5). Hence, even if 4 failures are observed, an alert is not raised until the fifth one is processed. ■ After 5 (*Y*) attempts, a look back calculates if there are at least 3 (*X*)or more failures out of the 5 attempts. For example, the table below explains the alert generation for Symptom Frequency: 3 out of 5 and Scope: Client A with Exit Status 99. Host Job Exit Status Outcome Client A 0 - Client A 99 - Client A 0 - Client A 99 - Client A 99 Alert Generated Client A 0 Alert Generated Client B 0 - Client B 99 -

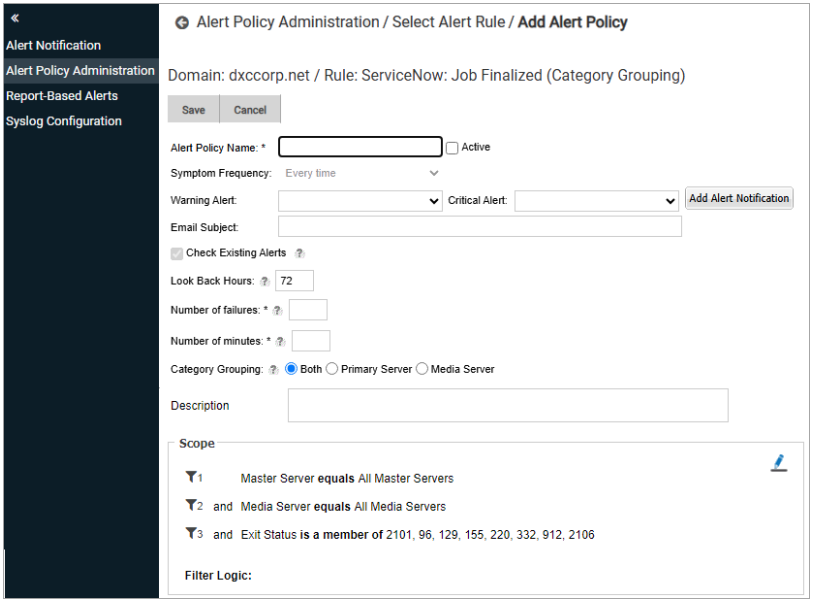
Table 11-2 Name the Policy and define Time Intervals *(continued)*

Field	Description
	`Client A 156 Alert NOT Generated as out of last 5 attempts 3 have failed with Exit Status 99.` `Client A 99 Alert Generated`
Schedule	Some Rules are based on a schedule. If the selected Rule is based on a pre-defined schedule, that default schedule is displayed. Click the icon to overwrite the default. See "Scheduled Rules vs. Run Time Rules" on page 459.
Warning Alert	Select the pre-defined Alert to issue for a warning condition. This would be triggered when the threshold value for a Warning condition is reached. See "Adding an Alert Notification Delivery Method" on page 481.
Critical Alert	Select the pre-defined Alert to issue for a critical condition. This would be triggered when the threshold value for a Critical condition is reached. See "Adding an Alert Notification Delivery Method" on page 481.
Check Existing Alert	When enabled, an alert is generated only if there is no existing alert found for the alert policy and alert object combination.
Number of Failures	The minimum count of failures that must occur in the previous 'n' minutes (where 'n' is the value specified in Number of Minutes field) before an alert is triggered. **Note:** For more details, refer to the example given at the end of this table.
Number of Minutes	The rolling time window within which the 'Number of failures' threshold must be met in order for an alert to be generated. **Note:** For more details, refer to the example given at the end of this table.
Look Back Hours	Configurable only if **Check Existing Alerts** is enabled and indicates the number of hours after which another alert can be generated provided the alert-triggering event has not resolved. The default number of hours are inherited from the system parameter Existing alert default look back (hours) that is defined under **Admin** > **Advanced** > **System Configuration** > **Portal** > **General**
Include NetBackup policy for existing alert check	Configurable only if **Check Existing Alerts** is enabled and applicable only for alerts integrated with ServiceNow, such as ServiceNow: Job Finalized. Checks for the uniqueness of client and NetBackup policy. If the NetBackup policy is not found (such as for Restore jobs), then the existing alerts / tickets are checked by including Job Type instead of NetBackup Policy

Table 11-2 Name the Policy and define Time Intervals *(continued)*

Field	Description
Enable NetBackup Backup Window Lookup	Select this parameter to process NetBackup specific ServiceNow job failure. Selecting this option, a ticket is raised in ServiceNow when there is no successful backup for a client after the close of its Backup Window associated with client's backup policy and schedule. If, at the time of Backup Window processing is over, there is a running job found, then a ticket is created if the job finishes as failure irrespective of the exit status of the jobs finished with in the Backup Window. **Note:** Selecting **NetBackup Backup Window Lookup** option, the **Symptom Frequency** option is disabled and do not have impact on the alerts or ServiceNow tickets processing. The **Symptom Frequency** and **Enable NetBackup Backup Window lookup** options are mutually exclusive.
Add Alert Notification	Click to setup the alert notification method as you are adding the Alert Policy. This can also be add outside of adding an alert policy. See "Adding an Alert Notification Delivery Method" on page 481.
Description	Add an optional description for the Alert Policy.
Category Grouping	Choose how the alerts must be grouped for matching events of primary and media servers: ■ **Both** (Default): Separate alert for each media server and primary server. ■ **Primary Server Only**: Single alert for each primary server, regardless of how many media servers match. ■ **Media Server Only**: Single alert for each media server, without grouping by the primary server.

Note: The following is an example of an alert rule which will be triggered with the values specified in **Number of Failures** and **Number of Minutes** fields.



The alert rule is triggered when the count of actual failed jobs is more than or equal to 'X' and no alert is triggered previously, within the last 'Y' minutes.

Given the value of 'X' for **Number of Failures** and 'Y' for **Number of Minutes** fields

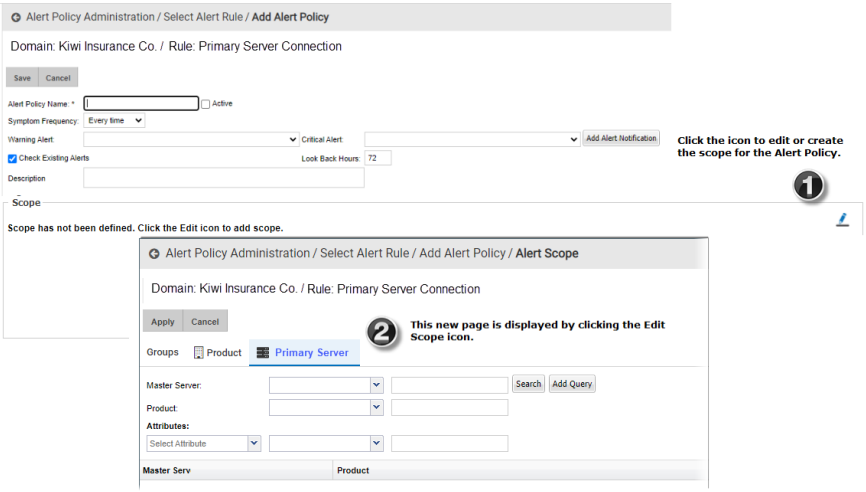
Table 11-3 Define Time Intervals

Number of Failures 'X'	Number of Minutes 'Y'	Actual Number of failures	Category Grouping ticket	Category Grouping ticket generated?
10	30	5	No	**No** (Number of actual failures is less than 10)
10	30	15	No	**Yes** (Number of actual failures is more than 10 and open ticket NOT found for the server.)
10	30	11	Yes	**No** (Number of actual failures is more than 10 and open ticket found for the server)

Define the Scope to Monitor

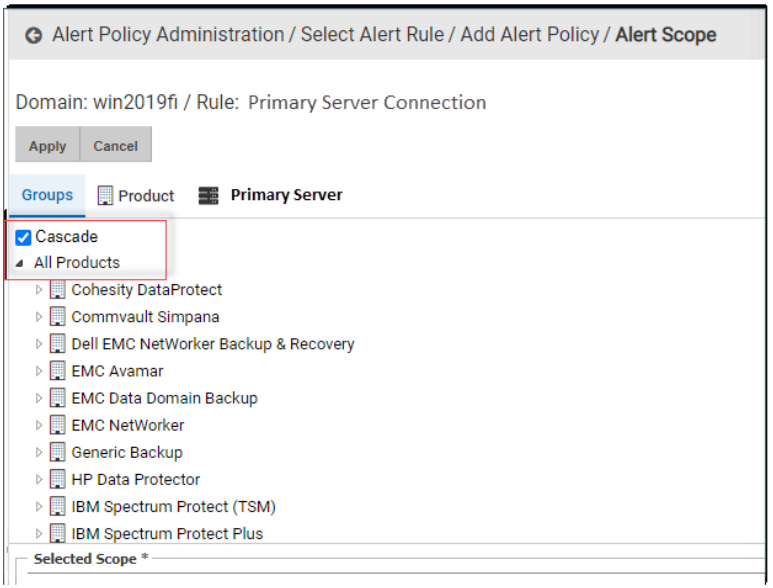
Policy scope can be set to alert on a wide range of conditions or a specific granular one. This window enables you to:

- Double-click to add an object to the scope
- Double-click to remove an object from the scope - This removes an object that has been place in the report scope.
- Drag and drop objects into the scope - Click the object and move it until you see a red dotted rectangle. Drop it into the pane.
- Search for objects to add to the scope.



Once in the new screen, use the **Groups** tab to browse for a broad scope such as everything from a particular vendor. The benefit of a broad scope is that when new objects are added to your datacenter and they fit the scope criteria, they are picked up automatically by the Alert Policy.

Note: Alerting is configured at the Domain level, however for multi-tenancy environments, the **Cascade** option in the **Groups** tab impacts what objects are available to monitor. Select **Cascade** to display objects at your Domain level and those domains that are children to your level. De-select **Cascade** to only display and monitor objects from your Domain.



Use the searchable tabs to refine the scope to a more granular level, such as watch a specific set of primary servers. For each object represented in the Groups tab, there is a corresponding searchable tab. The searchable tab selection is dynamic and derived from the Alert rule you select. For example, if you select **Primary Server Connection**, the **Groups** tab displays all discovered products in your datacenter. The searchable tabs enable you to find **Products** and **Primary Servers**. By using each of these components, you can construct a query to monitor exactly the scenario you require.

Attributes are also available to use as a scoping filter for most Rules. Attributes are user-defined characteristics of the objects.

See “Scope Example: Data Protection: Monitor for High Job Failure Rate” on page 473.

See “Scope Example: Data Protection: Monitor Backup Jobs” on page 474.

Using Attributes in the Scope

Attributes, either user-defined or system attributes, can be used to define the scope of your alert. Attributes are available for selection from most Rules. For example, you have set up a “location” attribute that you associated with the Primary Servers. With that set, you can generate an alert for Primary Servers in a particular location that also fit your other scope criteria. An attribute can be added as a part of the monitor query.

See “ About attributes ” on page 536.

Use **Show Resolved Scope** to display a list of items in your datacenter that currently meet the selected criteria. This provides validation that the scope you set is capturing what you intended.

Scope Example: Data Protection: Monitor for High Job Failure Rate

Rule	High Job Failure Rate
Scenario	Monitor all backup products in your datacenter that meet the threshold criteria for a failure percentage. This is a scheduled policy.

1. Navigate to **Alert>Alert Configuration>Alert Policy Administration**.
2. Click **Add**.
3. Choose the Alert Rule: **High Job Failure**. Double-click or click **Continue** to advanced to the next screen.
4. Name the Policy and set the **Symptom Frequency** to **Every time**.

   See “Name the Policy and Define Time Intervals” on page 465.
5. Choose or create the notification for a **Warning** and **Critical Alert**. Alert Notifications must be defined before the Alert Policy is created to select them for alerts. You can also click **Add Alert Notification** and define it when you configure the Alert Policy.

   See “Adding an Alert Notification Delivery Method” on page 481.
6. Under the **Scope** heading, click the **Edit** icon.
7. Click **Groups** and select **All Products**. The benefit of selecting **All Products** is that when new vendors are added to your datacenter they are picked up automatically by the Alert Policy.
8. Under the **Threshold Setting** heading, click the **Edit** icon.
9. Set the threshold and click **Active** toggle button to define the trigger for an alert. In this example, choose the **Operator** as **greater than or equal** to and we will set a value of 50 to indicate a Job failure rate 50% is considered a **Warning** threshold and 80 to indicate a **Critical** threshold.
10. Enable the Alert Policy by clicking the **Active** checkbox beside the Alert Policy Name field.
11. Click **Save**. The Alert Policy is displayed on the **Alert Policy Administration** page.

Scope Example: Data Protection: Monitor Backup Jobs

Table 11-4 Data Protection: Monitor Backup Jobs

Rule	Job Finalized
Scenario	Monitor all backup jobs for your Tech Pubs department whose Primary Server name contains NBU, and have an Exit Status Code of 123, 456, 789. However, you only care about the Exit Status Code if it happens 3 out of 5 times. This is a run-time policy. For the Job Finalized Alert Rule, Exit Status is a mandatory entry.

1. Navigate to **Alert>Alert Configuration>Alert Policy Administration**.
2. Click **Add**.
3. Choose the Alert Rule: **Job Finalized**. Double-click or click **Continue** to advanced to the next screen.
4. Name the Policy and set the **Symptom Frequency to X out of Y**.

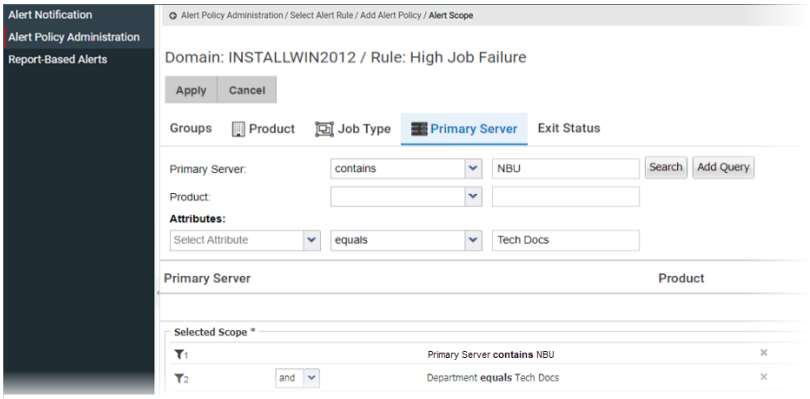
   See "Name the Policy and Define Time Intervals" on page 465.
5. Enter 3 of 5 to indicate you only want to trigger an alert if the symptom occurs 3 times out of 5.
6. Choose or create the notification for a **Warning** and **Critical Alert**. Alert Notifications must be defined before the Alert Policy is created to select them for alerts. You can also click **Add Alert Notification** and define it when you configure the Alert Policy.

   See "Adding an Alert Notification Delivery Method" on page 481.

Setup the Primary Server Scope

1. Under the **Scope** heading, click the **Edit** icon.
2. Click the **Primary Server** tab.
3. Choose the Operator **contains** and enter **NBU** in the field. Advanced Filter operators are available.

   See "Advanced filter operators" on page 46.
4. Click **Search** to display all discovered Primary Servers that contain NBU in the name field.



5. Click **Add Query**. This adds your filter selection as the first row in the **Selected Scope**. You could also drag and drop, or double click an individual Primary Server to add it to the **Selected Scope**. For this example, we are adding all the values found in the **Search**.
6. Select **Department** under **Attributes**. Choose the Operator **contains** and enter **Tech Docs** in the field.

Note: Because the filter value is broad (NBU instead of an individual Primary Server) when new Primary Servers are added to your Tech Docs department and given the same NBU naming convention, they are automatically added to this Alert Policy scope.

Setup the Exit Code Scope

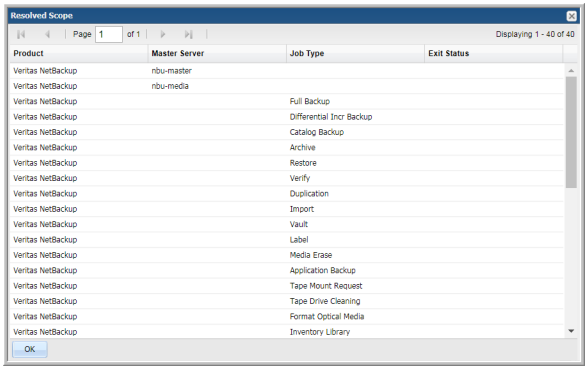
1. Click the **Exit Status** tab. For the Job Finalized Alert Rule, Exit Status is mandatory.
2. Choose the Operator **equals** and enter **123, 456,789** in the field. Advanced Filter operators are available.

   See “Advanced filter operators” on page 46.
3. Click **Add Query**. This adds your filter selection as the next row in the Alert Policy **Selected Scope**.

   If your query is multiple lines, you can use the **Filter Logic** field to customize the filter expression order and the operators. Logic defined in this field will override any setting established on the top of the dialog. Use the numbers on the left of the filter expressions to construct your Filter Logic.

Edit the logic using the filter numbers and by adding parentheses or changing the operators. For example, you can change “1 AND 2 OR 3” to “1 AND (2 OR 3)”.

4. Click **Show Resolved Scope**. This lists everything in your datacenter that meets the criteria defined in your Selected Scope.



5. Enable the Alert Policy by clicking the **Active** checkbox beside the Alert Policy Name field.
6. Click **Save**. The Alert Policy is displayed on the **Alert Policy Administration page**.

Set the Threshold for an Alert Rule

Some Rules have specific trigger thresholds to set for an alert. You can enter values that indicate a warning or a critical alert. The threshold availability is dependent on the Rule. If a threshold is displayed, enter the values to indicate what constitutes a warning alert and what is a critical alert. Thresholds must also be explicitly activated.

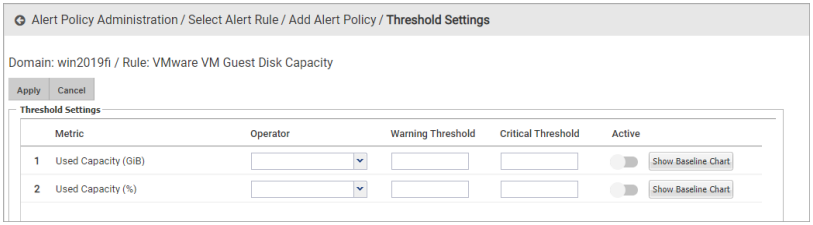
To help determine threshold values, you can view a baseline chart based on collected historical data to see where the numbers currently fall. Use this information to create your own threshold values or apply the historical baseline directly to your policy.

Show Baseline Example

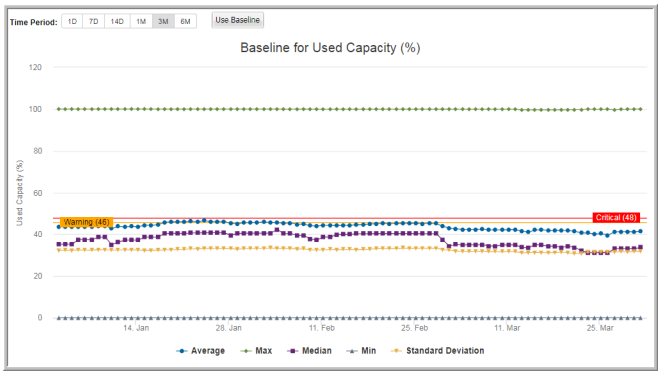
The following example uses the Alert Rule VMware VM Guest Disk Capacity.

1. Enter a percentage value for the Warning Threshold. For example, if you’d like to know when your Used Capacity reaches 80%, enter 80 to trigger a warning alert. If the Alert policy is setup to recognize a warning threshold, when this value is reached, an alert could be issued.

2. Enter a percentage value for the Critical Threshold. For example, if you'd like to know when your Used Capacity reaches 90%, enter 90 to trigger a critical alert. If the Alert policy is setup to recognize a critical threshold, when this value is reached, an alert could be issued.



3. Click **Show Baseline Chart** to display a chart of historical data for the selected threshold for VM Guest Disks. The chart is based on your selected scope. You can adjust the Time Period as required to examine the historical results. This helps you to determine the numbers for your warning and critical thresholds. You can derive your own numbers from the data the create threshold values or apply the baseline numbers directly.



License Usage Alert Rule

The License Usage rule can be used to created a policy to alert on usage percentage and approaching expiration dates for licensed modules. License usage awareness is important because once limits have been reached, data for any new devices will not be captured and in the case of expiration, Portal access can be denied until renewal.

This System Administration rule is a special category of alert rule and only available to Super Users in the root domain. In general, once rules are configured, you save the custom instance and create the Alert Policy. An initial instance of the License Usage policy is set to Active and comes configured with:

- Default name

- Schedule
- Scope
- Threshold settings for license expiration reminders and usage percentages for overage

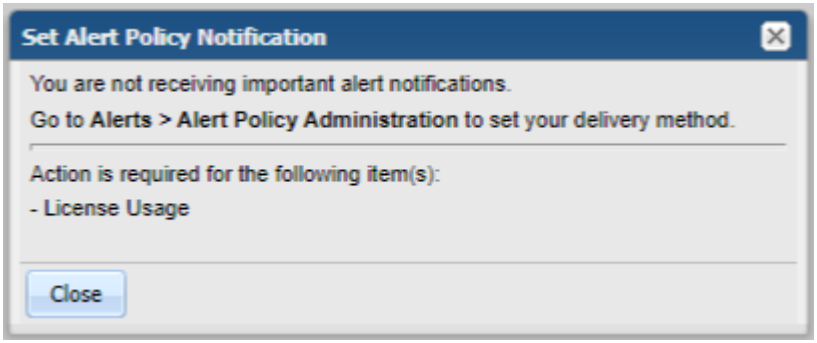
These settings allow IT Analytics to report on and ultimately display notifications for critical licensing issues.

Setting the Alert Notification

License usage information is available in the reports License Summary and Alert Summary, but you can choose to be notified when the thresholds are met or nearing. To do so, you must define how the alert is delivered. This is not part of the pre-configured settings.

See "Adding an Alert Notification Delivery Method" on page 481.

When you log into the Portal, a pop-up is displayed until the notification method is configured. The pop-up can also be disabled by de-activating the License Usage alert policy or deleting it entirely. A new License Usage policy can be added at any time.



See "Deleting an Alert Policy" on page 604.

See "Add/Edit an Alert Policy" on page 463.

ServiceNow: Job finalized alert rule

This alert rule, specifically for ServiceNow integrations, enables users to easily setup a rule for failing backup jobs on primary servers, policies and/or schedules and create a ServiceNow incident.

Once the backup job is successfully executed, the ServiceNow incident is automatically tracked and marked as Resolved. However, a successful backup job

means the first time the job ran successfully after one or more failures, for the same client-server pair that triggered the alert.

The alert rule, ServiceNow: Job Finalized, alerts on jobs of the selected product vendor and primary server that have completed with a specified Exit Status. For Veritas NetBackup, the scope can also include policy and/ or schedule. This alert also enables incident creation within ServiceNow for the selected scope.

Note: For the ServiceNow: Job Finalized alert rule, Exit Status is a mandatory selection.

With this specific alert rule, there are some special considerations.

- Policies and Schedules can be selected as part of the scope for data collectors, such as Veritas NetBackup and Backup Exec, where they exist. You can choose any product and their Primary Servers.
- When defining the scope, you can choose a **Server Type**. This means that if Client is selected, then only those jobs will be alerted on where the Primary Server is not backing up itself. If any other Server is selected, then only those jobs will be monitored where the server is also the client (backing up itself). If this is skipped, then all the jobs will be reported as per the other scope filters.

Scope Example: Data Protection: ServiceNow Jobs

Rule	ServiceNow: Job Finalized
Scenario	This scenario assumes that the ServiceNow for Backup Solution app is installed and configured. Monitor all Veritas NetBackup backup jobs for your Info Dev department for the policy named FirstCatalogPolicy and have an Exit Status Code of 123, 456, or 789. If this alert is triggered, a ServiceNow incident is created within the ServiceNow application.

1. Navigate to **Alerts** > **> Alert Policy Administration**.
2. Click **Add**.
3. Choose the **Alert Rule: ServiceNow: Job Finalized**. Double-click or click **Continue** to advanced to the next screen.
4. Name the Policy and set the **Symptom Frequency** to **Every time.**

5. Choose or create the notification for a **Warning** and **Critical Alert**. Alert Notifications must be defined before the Alert Policy is created to select them for alerts. You can also click **Add Alert Notification** and define it when you configure the Alert Policy.
6. Under the **Scope** heading, click the **Edit** icon.
7. On the **Groups** tab, filters are displayed. Click the required filters to expand.

   Note: For example: Click **All Products** and double-click **Veritas NetBackup** to add to the scope.

8. Click the **Exit Status** tab. For the **ServiceNow:Job Finalized** alert rule, **Exit Status** is mandatory.
9. To include comma separated list of **Exit Status** in the **Alert Scope**, choose comparison operator *is member of* and add comma separated list of Exit Statuses in the text box such as 19,203,112.
10. Click **Add Query**. This add your filter selection as the next row in the Alert Policy **Selected Scope**. If your query is multiple lines, you can use the **Filter Logic** field to customize the filter expression order and the operators. Logic defined in this field will override any setting established on the top of the dialog. Use the numbers on the left of the filter expressions to construct your Filter Logic.
11. Edit the logic using the filter numbers and by adding parentheses or changing the operators. For example, you can change "1 AND 2 OR 3" to "1 AND (2 OR 3)".
12. Click **Show Resolved Scope**. This lists everything in your datacenter that meets the criteria defined in your Selected Scope.
13. Enable the Alert Policy by clicking the **Active**checkbox beside the Alert Policy Name field.
14. Click **Save**. The Alert Policy is displayed on the **Alert Policy Administration** page.

Copying an Alert Policy

Sometimes when you add an Alert Policy with a number of specific query parameters, you may want to modify one or two parameters without having to recreate the entire policy. Copy allows you use an existing Alert Policy as a template and create a new version that you can modify and save with another name.

1. Navigate to **Alerts>Alert Configuration>Alert Policy Administration**.

2. Select an Alert Policy and click **Copy**.
3. Enter a name for the new Alert Policy. The new Policy is displayed on the **Alert Policy Administration** page.
4. Select the new Alert Policy and click **Edit**.
5. Modify the parameters as required.

Deleting an Alert Policy

1. Navigate to **Alerts>Alert Configuration>Alert Policy Administration**.
2. Select an Alert Policy and click **Delete**. A confirmation dialog is displayed.

Managing alert notifications

Alert Notifications define how and when you want to be notified when thresholds are met. Alerts can be sent to you when a specific set of criteria is met. Once defined, they are saved with a name and can be assigned to any Alert Policy. For example, a warning alert might only send an email to an Administrator and a critical alert would send an email to the Administrator and the engineering team.

The Inventory also displays badging to indicate the alerts on the specific objects with the ability to view relevant reports. Even if you do not configure Alert notifications, monitoring still continues and this same information can be viewed on-demand through the different Alert reports - where events can be captured and drilled into for root-cause analysis.

- See “Adding an Alert Notification Delivery Method” on page 481.
- See “Deleting an Alert Notification Delivery Method” on page 483.
- See “Suppress alert notifications” on page 498.

Adding an Alert Notification Delivery Method

1. Select **Alerts>Alert Configuration>Alert Notification**.
2. Click **Add**.
3. Enter a **Name** and **Description**.
4. Select the delivery method for the alert. Choose from:
   - **Email**: Enter a valid email address. This can be a comma-separated list of email addresses. One email is sent for entire report.
   - **SNMP Trap**: When you check this box choose V2 or V3. Entry fields change based on your selection.

- For **V2**, enter the **Server Name**, **Port**, and **Community**. One alert will be sent for each row.
- For **V3**, enter the **Server Name**, **Port**, **Security Name** and select a **Security Level**

SNMP Trap V2 Field Descriptions

Server: Enter the IP address of the server which will be used to receive trap.

Port: Port on which SNMP server is running.

Community: A type of shared password between the portal and the SNMP server, which is used to authenticate the portal.

SNMP Trap V3 Field Descriptions

Server: Enter the IP address of the server which will be used to receive trap.

Port: Port on which SNMP server is running.

SNMP Engine ID - The Engine ID is used by SNMPv3 entities to uniquely identify them.

Security Name: A unique name and similar to username for authentication. This Security Name is used during the creation of users in the SNMP server:

For example: if the Security Name is "demo_user", use following command to create a user:

```
createUser -e
0x8000137001a9fe97108f0f4f2f demo_user
MD5 demo_password DES demo_password
```

where:

```
0x8000137001a9fe97108f0f4f2f = engineId
demo_user = security name
```

Security Level:

- None
- Authentication: Provide authentication protocol and authentication passphrase
- Authentication & Privacy: Provide authentication protocol, authentication passphrase and privacy protocol and privacy Passphrase.

- **Execute Script**: The user-created script must reside in: /opt/aptare/portal/user_script. Check the box and enter a shell script name (Linux). If a path name is included, it will be appended to the above path. For example, when the alert is triggered.

- **Syslog**: Select this and then attach a Syslog configuration from the list that must perform the alert notification. The list is populated provided a configuration is added under **Syslog Configuration** for the domain.
  See "Syslog configuration" on page 487.
- **Webhook**: Select the webhook configuration that you want to use for the alert notification. This list is populated with the webhooks configured under **Admin** > **Advanced** > **Webhook**.

5. Optionally, click **Test** to trigger the action and validate the delivery.
6. Click **Save**.

Deleting an Alert Notification Delivery Method

Alert notifications are associated with Alert Policies. When you delete notifications, they are removed from the policies.

1. Navigate to **Alerts>Alert Configuration>Alert Notification**.
2. Select an Alert Notification and click **Delete**.

Report-Based alerts

When a tabular report has been populated with data, you can set an alert to notify you. For example, you can save a Job Summary Report for Failed Events and then configure an alert for this report. The Portal will check for a report that contains data according to the schedule you select. You can configure alerting for any report that contains a single table. A report must be saved before you can set an alert.

See "Setting Up Alerts for Tabular Reports" on page 186.

For convenience, your reports with an alert configured are displayed in this section of the Alerts tab. These report-based alerts are not system wide and only belong to the logged in user.

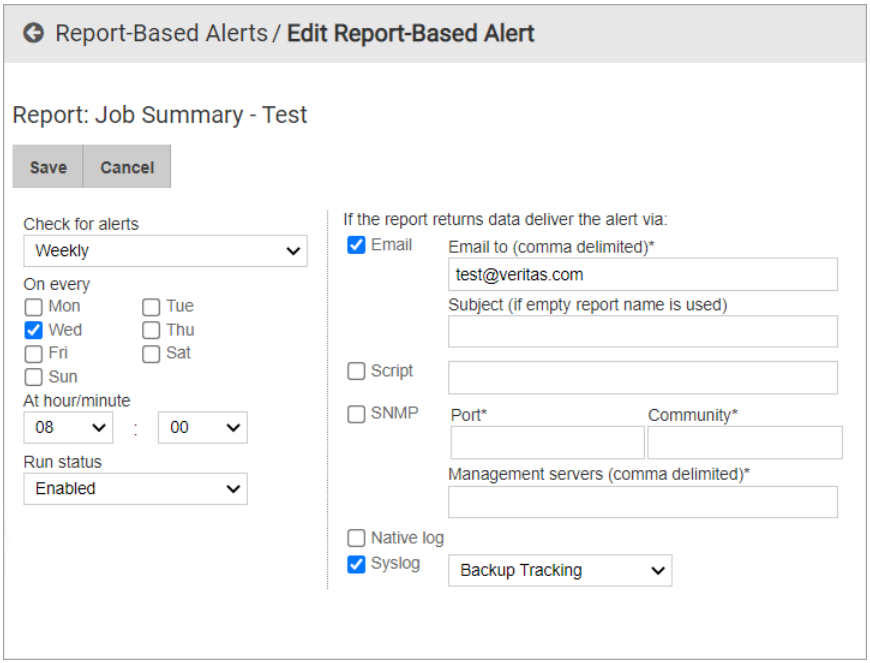
- See "Edit notifications for Report-Based Alerts" on page 483.
- See "Run a Report-Based Alert On-Demand" on page 486.
- See "Delete Report-Based Alert" on page 487.

Edit notifications for Report-Based Alerts

For convenience, your reports with an alert configured are displayed in this section of the Alerts tab and you can edit how you are notified of the alert.

1. Select **Alerts** > **Alert Configuration** > **Report-Based Alerts**. A list of tabular reports with an alert configured are displayed.

2. Select a report and click **Edit**.



3. Modify the schedule and notification method using the following fields:

Check for alerts/Schedule	The following options are displayed in the drop-down list: ■ On a defined schedule- Primary schedules can be configured and then applied to reports. Modifications to a primary schedule will automatically be applied to all the reports associated with that primary schedule. See "Configure primary schedules" on page 606. ■ Frequency in Minutes - Select from every 5, 15, or 30 minutes. ■ Hourly - Select 1, 2, 3, 4, 6, or 12 hours. ■ Daily - At: hour/minute. Select a specific time. ■ Weekly ■ On every. The day(s) on which the report will be checked. ■ At: hour/minute. Select a specific time. ■ Monthly ■ On the. The day on which the report will be checked. ■ At: hour/minute. Select a specific time. The Portal can check the same report multiple times in a single day. ■ Cron Expression - Enter a CRON expression to fine tune the alert schedule. for details about working with CRON expressions.
Run Status	Select **Enabled** or **Disabled**. This selection enables or disables the schedule for the report to be checked for alerts.
Email	Check the box and provide a comma-separated list of email addresses. One email will be sent for the entire report.
Subject	Enter the subject. The report name is used if the field is left blank.

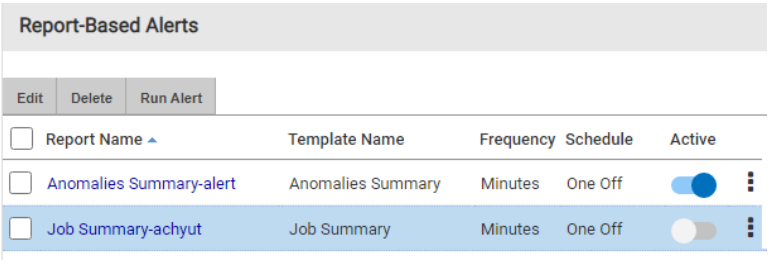
Script	The user-created script needs to reside in: `/opt//portal/user_scripts` Check the box and enter a shell script name (Linux). If a path name is included, it will be appended to the above path. For example, filter a report to include only the columns of information that you need. When the alert is triggered, a .csv file of the report is generated and the path to that file is made available to the script to take whatever actions you want with this report data. Administrators: To enable Script delivery refer to the following section. See “Add/Configure a domain” on page 766.
SNMP	When you check this box, the Port, Community string, and Management servers fields will be populated from the configured defaults. One alert is sent for each row. To override the defaults, overwrite any or all of the three SNMP fields. Administrators: To enable SNMP delivery and to configure SNMP default values refer to the following. See “Add/Configure a domain” on page 766.
Native Log	When this box is checked, a log entry is written to the OS-specific log: either the Windows event log or the Linux syslog. Administrators: To enable Native Log delivery refer to the following section. See “Add/Configure a domain” on page 766.
Syslog	When this box is checked, alert notifications are routed through the Syslog configuration specified from the Syslog list. See “Syslog configuration” on page 487.

Run a Report-Based Alert On-Demand

Run report-based alerts from the **Alerts** tab to:

- Validate the alert is working as defined
- Launch an immediate run of the alert without waiting for the scheduled run

1. Navigate to **Alerts>Alert Configuration>Report-Based Alerts**.



2. Select the **Report Name**.
3. Click **Run Alert**.

Delete Report-Based Alert

Delete report-based alerts from the **Alerts** tab. This action removes the report from the Report-Based Alerts list and deletes the alert from the report. The action does not delete the report.

1. Navigate to **Alerts > Report-Based Alerts**.
2. Select the **Report Name**.
3. Click **Delete**.

Syslog configuration

Configure a Syslog server in your environment before setting Syslog as the alert notification method. With this configuration Syslog notification becomes available for default alerts and report-based alerts.

The following Syslog server details are required for the configuration:

- Server IP
- Name
- Port number
- Application name and protocol

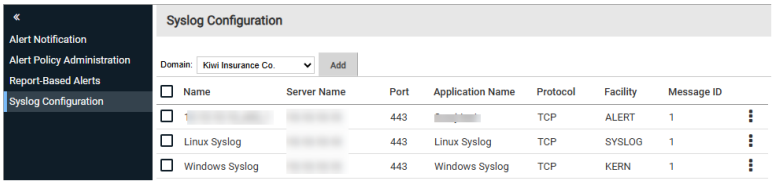
The rest of the details can be specified during the configuration.

Configure Syslog notification

Before you configure the Syslog notification, you must ensure you have Syslog server already configured in your environment and you have its IP, name, port number, application name, and protocol details. See “Configure alerting for a domain” on page 768.

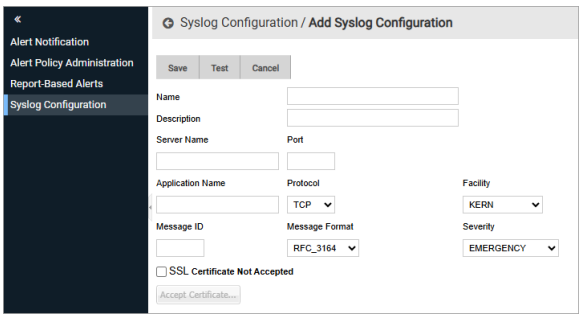
To configure Syslog notification:

1. Go to **Alerts** > **Syslog Configuration**. If you are adding the first notification, the page appears blank. During subsequent configurations, the page displays a list of previously configured notifications.



2. Select a domain from the **Domain** list. The **Domain** list is displayed only if you have access permissions to more than one domains.

3. Click **Add** to open the configuration screen.



4 Configure Syslog notification based on the descriptions below:

Field Name	Description
Name	Name of your Syslog configuration. This name appears in the list displayed on the Syslog Configuration page and in the Syslog list when you configure an alert notification.
Description	A short description of your Syslog configuration.
Server Name	Enter the IP address of the Syslog server.
Server Port	Port number on which the Syslog server is running.
Application Name	The device or application that generated the message.
Message ID	The identification number of the message.
Protocol	Protocol used to send the message. Choose from TCP or UDP.
Message Format	Choose from RFC_3164, RFC_5424, or RFC_5424.
Facility	Facility represents the machine process that created the Syslog event. For example, is the event created by the kernel, by the mail system, or by security/authorization processes, and so on. In the context of this field, it is a kind of filter, instructing SMS to forward to the remote Syslog Server only those events whose facility matches the one defined in this field.
Severity	Select the level of severity that will trigger the notification.

Field Name	Description
SSL	SSL works between the portal and the Syslog server. To communicate securely both need each other's certificates. Hence, Portal certificate must be present in the portal keystore before configuring Syslog with SSL. See *Add a Certificate into the portal keystore* in the *IT Analytics System Administrator guide*. Click **Accept Certificate** to fetch the Syslog server's certificate and display it for verification. After accepting the certificate, it is used to validate the identity of Syslog server.

5 Click **Test** to verify the configuration and connectivity on the Syslog server.

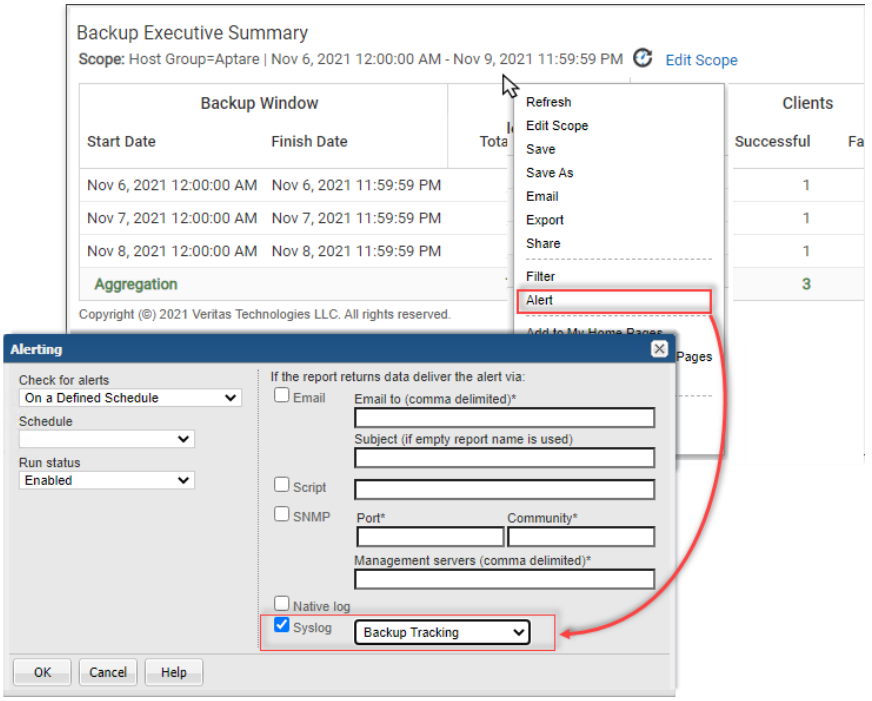
6 Click **Save**. Your configuration appears in the **Syslog Configuration** list. Once saved, the Syslog configuration name also appears in the list against the **Syslog** checbox of the Add Alert Notification screen.

Enable Syslog notification for report-based alerts

Prerequisite: Before you enable Syslog notification for report-based alerts, enable Syslog alerting for the domain from **Admin** > **Domains** > **Domains** > **Alerting** tab. See "Configure alerting for a domain" on page 768.

To enable Syslog notification for report-based alerts:

1. Open the report for which you want to enable Syslog alert notification.
2. Right-click on the report and select **Alert** from the menu. The **Alerting** screen is displayed.



3. Define the alert frequency and select **Syslog** as the alert delivery mode. Also, specify the Syslog configuration from the list.
4. Click **OK** to save the changes. The portal uses the specified Syslog configuration to send notifications of the report-based alerts.

Delete configuration for Syslog notification

You can delete the configuration for Syslog notification only if it is no longer attached to any report or alert. The portal displays an error message stating the Syslog is configured with a report or alert when you attempt to delete an attached configuration.

To delete a Syslog configuration from the portal:

1. Make sure the none of the alerts or reports are attached to the Syslog configuration.
2. Go to **Alerts** > **Syslog Configuration**. The page displays a list of configured notifications.
3. Select the configuration and click **Delete**.

About using SNMP

This section provides information about SNMP and how IT Analytics uses SNMP.

See "About SNMP" on page 492.

See "About SNMP versions" on page 493.

See " SNMP versions supported in IT Analytics" on page 493.

See "About the Management Information Base (MIB) and IT Analytics support" on page 493.

See "SNMP traps" on page 494.

About SNMP

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. It is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite. SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth.

SNMP is based on the manager model and agent model. This model consists of a manager, an agent, a database of management information, managed objects, and the network protocol.

The manager provides the interface between the human network manager and the management system. The agent provides the interface between the manager and the physical devices being managed.

The manager and agent use a Management Information Base (MIB) and a relatively small set of commands to exchange information. The MIB is organized in a tree structure with individual variables, such as point status or description, being represented as leaves on the branches. A numeric tag or object identifier (OID) is used to distinguish each variable uniquely in the MIB and in SNMP messages.

About SNMP versions

Many versions of SNMP are available.

The versions of SNMP protocol are as follows:

- SNMPv1
  The SNMPv1 version is the first version of the protocol and is defined by RFC 1157. This document replaces the earlier versions that were published as RFC 1067 and RFC 1098. Security is based on community strings.
- SNMPv2
  It was created as an update of SNMPv1 adding several features. The key enhancements to SNMPv2 are focused on the SMI, manager-to-manager capability, and protocol operations.
  SNMPv2c combines the Community-based approach of SNMPv1 with the protocol operation of SNMPv2 and omits all SNMPv2 security features.
  - The original SNMPv2 (SNMPv2p)
  - Community-based SNMPv2 (SNMPv2c)
  - User-based SNMPv2 (SNMPv2u)
  - SNMPv2 star (SNMPv2*).
- SNMPv3
  This version of the protocol is a combination of user-based security and the protocol operations and data types from SNMPv2p, and support for proxies. The security is based on that found in SNMPv2u and SNMPv2*. RFC 1905, RFC 1906, RFC 2261, RFC 2262, RFC 2263, RFC 2264, and RFC 2265 define this protocol.

SNMP versions supported in IT Analytics

IT Analytics supports the following SNMP versions:

- SNMPv2c
- SNMPv3

About the Management Information Base (MIB) and IT Analytics support

Each SNMP element manages specific objects with each object having specific characteristics. Each object and characteristic has a unique object identifier (OID) that is associated with it. Each OID consists of the numbers that are separated by decimal points (for example, 1.3.6.1.4.1.2682.1).

These OIDs form a tree. The MIB associates each OID with a readable label and various other parameters that are related to the object. The MIB then serves as a data dictionary that is used to assemble and interpret SNMP messages.

SNMP traps

This section explains the content of an SNMP trap that is sent from IT Analytics.

MIB files for report based alert can be found at `<APTARE_HOME>/goodies/`. For example if <APTARE_HOME> is `/opt/aptare`, then MIB files will be present at `/opt/aptare/goodies`.

See "SNMP traps for system alerts" on page 494.

See "SNMP traps for report-based alerts" on page 496.

SNMP traps for system alerts

Each IT Analytics trap contains 2 standard object identifiers and 13 IT Analytics-specific object identifiers. An object identifier (or OID) is a numeric string that is used to uniquely identify an object.

The following table shows the contents of a trap that is sent from IT Analytics. A total of 15 bindings (or 15 name-value pairs) are present in each trap that is sent from IT Analytics. Each binding associates a particular Management Information Base (MIB) object instance with its current value.

IT Analytics trap binding shows the name-value pairs that the traps pass to the SNMP manager if SNMP payload format is selected as **Enhanced**.

Table 11-5 IT Analytics system alerts trap binding

Name	Value
1.3.6.1.2.1.1.3.0	This field is the time (in hundredths of a second) between when IT Analytics server service starts and the IT Analytics trap is sent. See Request for Comment (RFC) 1905 and 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1173792454

Table 11-5 IT Analytics system alerts trap binding *(continued)*

Name	Value
1.3.6.1.6.3.1.1.4.1.0	This field is the unique identifier for this trap. See RFC 1905 and RFC 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1.3.6.1.4.1.15622.3.12.10.2.0.4
1.3.6.1.4.1.15622.3.12.10.1.1 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch. alertApplicationName)	This field is the Application Name. Example: IT Analytics
1.3.6.1.4.1.15622.3.12.10.1.2 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertDomainId)	This value specifies the Domain Id of the alert. Example: 100000
1.3.6.1.4.1.15622.3.12.10.1.3 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertDescription)	This field is the alert description.
1.3.6.1.4.1.15622.3.12.10.1.4 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertPolicyName)	This field is the alert policy name.
1.3.6.1.4.1.15622.3.12.10.1.5 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertObjectTypeName)	Defines the object type for which the alert was set. Object type dictates which drill down report is displayed.
1.3.6.1.4.1.15622.3.12.10.1.6 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertRootObjectName)	Displays the object on which the alert was issued.
1.3.6.1.4.1.15622.3.12.10.1.7 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertParentObjectTypeName)	If a hierarchy structure exists, displays the object type of the parent for the alert object.
1.3.6.1.4.1.15622.3.12.10.1.8 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertParentObjectName)	If a hierarchy structure exists, displays the name of parent of the alert object.
1.3.6.1.4.1.15622.3.12.10.1.9 (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertLastUpdated) (iso.org.dod.internet.private.enterprises.it-analytics.products.Itaportal.itaSystemAlert.itaSystemAlertTrapVarsBranch.alertLastUpdated)	Displays the date the alert condition occurred.

Table 11-5 IT Analytics system alerts trap binding *(continued)*

Name	Value
1.3.6.1.4.1.15622.3.12.10.1.10 (iso.org.dod.internet.private.enterprises.ianalytics.products.ia.portalIaSystemAlertIaSystemAlertTrap.vasBranchalertCount)	The number of occurrences of alert.
1.3.6.1.4.1.15622.3.12.10.1.11 (iso.org.dod.internet.private.enterprises.ianalytics.products.ia.portalIaSystemAlertIaSystemAlertTrap.vasBranchalertSeverity)	Indicates the severity of the alert.
1.3.6.1.4.1.15622.3.12.10.1.12 (iso.org.dod.internet.private.enterprises.ianalytics.products.ia.portalIaSystemAlertIaSystemAlertTrap.vasBranchalertLastAlertDate)	Displays the latest date the alert condition occurred.
1.3.6.1.4.1.15622.3.12.10.1.13 (iso.org.dod.internet.private.enterprises.ianalytics.products.ia.portalIaSystemAlertIaSystemAlertTrap.vasBranchticketNumber)	ServiceNow ticket number.

The first two OIDs listed in the table are standard SNMP OIDs. The other OIDs starting from 1.3.6.1.4.1.15622.3.12.10.1.1 to 1.3.6.1.4.1.15622.3.12.10.1.12 are IT Analytics OIDs. As per SNMPv2c trap definition, the two standard SNMP OIDs must be present as part of every trap.

All the 12 IT Analytics OIDs are defined in the IT Analytics MIB files. However, the two standard OIDs are not defined in the IT Analytics MIB files.

SNMP traps for report-based alerts

Each IT Analytics report-based alert trap contains two standard object identifiers and 2N IT Analytics-specific object identifiers where value of N depends on number of columns present in each row in the report. An object identifier (or OID) is a numeric string that is used to uniquely identify an object.

The following table shows the contents of a trap that is sent from IT Analytics for report-based alerts. Four bindings (or 4 name-value pairs) are present in each trap that is sent from IT Analytics and 2N bindings are present for each column present in report. For each column, there are 2 bindings, one for column name and other for its value. Each binding associates a particular Management Information Base (MIB) object instance with its current value.

IT Analytics trap binding shows the name-value pairs that the traps pass to the SNMP manager if SNMP payload format is selected as Enhanced.

Table 11-6 IT Analytics report-based alert trap binding

Name	Value
1.3.6.1.2.1.1.3.0	This field is the time (in hundredths of a second) between when IT Analytics server service starts and the IT Analytics trap is sent. See Request for Comment (RFC) 1905 and 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1173792454
1.3.6.1.6.3.1.1.4.1.0	This field is the unique identifier for this trap. See RFC 1905 and RFC 2576 for a detailed definition. http://www.ietf.org/rfc/rfc1905.txt http://www.ietf.org/rfc/rfc2576.txt Example: 1.3.6.1.4.1.15622.3.12.11.2.0.4
1.3.6.1.4.1.15622.3.12.11.1.1 (iso.org.dod.internet.private.enterprises.analytics.products.datareport.AlertReport.AlertTrap.vasBranchreportName)	It displays the report name
.1.3.6.1.4.1.15622.3.12.11.1.2 (iso.org.dod.internet.private.enterprises.analytics.products.datareport.AlertReport.AlertTrap.vasBranchreportDataCount)	It displays number of columns present in report. This value can be used to iterate over bindings which contains column name and column value. For example, if the value is 2 then it means 2 columns are present in the report and there are 4 bindings for column name and value.
1.3.6.1.4.1.15622.3.12.11.1.3.1.1.1 (iso.org.dod.internet.private.enterprises.analytics.products.datareport.AlertReport.AlertTrap.vasBranchreportDatareportEntry.columnName.1)	It displays name of column 1 of the report.
1.3.6.1.4.1.15622.3.12.11.1.3.1.2.1 (iso.org.dod.internet.private.enterprises.analytics.products.datareport.AlertReport.AlertTrap.vasBranchreportDatareportEntry.columnValue.1)	It displays value of column 1 of the report.
1.3.6.1.4.1.15622.3.12.11.1.3.1.1.N (iso.org.dod.internet.private.enterprises.analytics.products.datareport.AlertReport.AlertTrap.vasBranchreportDatareportEntry.columnName.N)	It displays name of column N of the report.

Table 11-6 IT Analytics report-based alert trap binding *(continued)*

Name	Value
1.3.6.1.4.1.15622.3.12.11.1.3.1.2.N (.iso.org.dod.internet.private.enterprises.analytics.productData.reportAlert.reportAlertTrap.vasBranch.reportDataReportEntry.columnName.N)	It displays value of column N of the report.

Note: The name-value pair in the report-based alert trap binding will be from 1 to N, where N is number of columns present in the report.

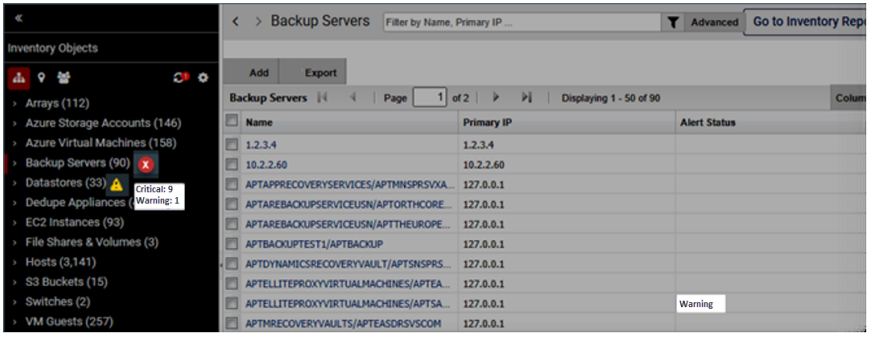
Detect alerts in the inventory

Alerts can be delivered through a defined notification method.

See "Adding an Alert Notification Delivery Method" on page 481.

You can also quickly detect alerts in the Inventory with badging. Icons displayed in the Inventory Objects panel identify which objects have triggered a critical or warning alert using the thresholds you defined. Click the badge to display Alert Details. If Alerts have been suppressed using the Alert Detail report they do not show in the count displayed in the badging. By selecting the object, you can display reports that help you to identify where the problems are occurring and take action.

See the *Report Reference Guide* for information on Alert Detail report.



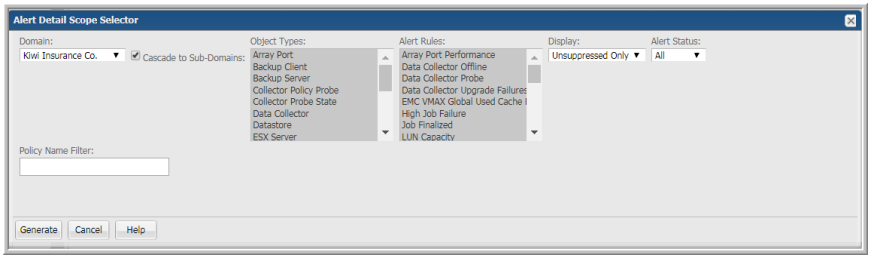
Suppress alert notifications

Use the report to suppress alert notifications system-wide. This report is also useful for viewing alerts once they are suppressed. Alert Suppression is privilege-based.

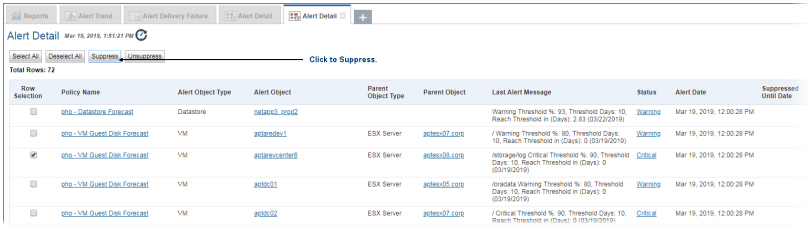
See the *Report Reference Guide* for information on the Alert Detail report.

1. Search for **Alert Detail**.

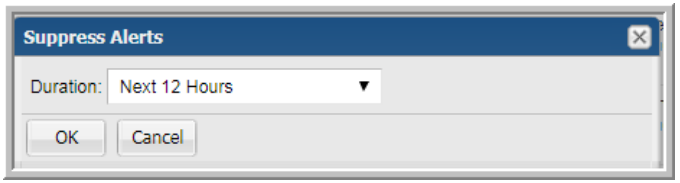
2. Select to display **Unsuppressed Only** or **All** in the **Alert Detail Scope Selector**.



3. Generate the report. The **Alert Detail** is displayed.



4. Select the row of the alert to suppress.
5. Click **Suppress**.

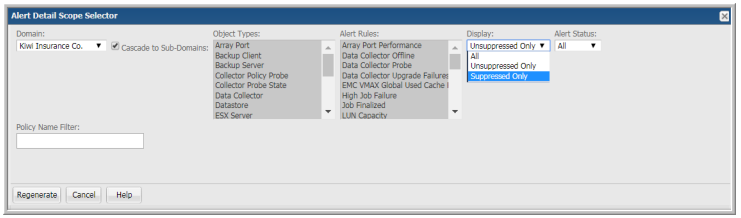


6. Select the duration for suppression.

View suppressed alerts

Use the Alert Detail report to view suppressed alerts.

1. Search for **Alert Detail**.
2. Select to display **Suppressed Only** or **All** in the **Alert Detail Scope Selector**.



3. Generate the report. The **Alert Detail** is displayed.

Chapter 12

Manage hosts, backup servers, and host groups

This chapter includes the following topics:

- About hosts, backup servers, and host goups
- Plan your host group hierarchy
- Manage host groups
- NetBackup Primary Servers
- About the Primary Servers Group
- Identifying Inactive Clients (For Veritas NetBackup)
- Managing hosts and backup servers
- Manage host aliases

About hosts, backup servers, and host goups

Hosts and Backup Servers can be added to the portal by:

- Data collection
- Importing
- Manual addition

Once hosts and backup servers are in the Portal, you can organize them into Host Groups to serve a number of functions:

- Organize hosts to create groupings relevant to your business.

- Control report scope to tailor report output to include only a specific subset of host data. In the Portal, this data filtering capability is known as the report scope.
- Limit user access to hosts and data by assigning a home host group to a user, limiting access to only the data in that group and its sub-groups. This also serves as a security mechanism. Domains are also used to limit access.
  See “Root folder and domains” on page 765.
  for additional information.

Plan your host group hierarchy

The Portal installation sets up a default Host Group hierarchy. This default hierarchy includes a group that represents a physical grouping of hosts. You can create additional host groups for hosts or backup servers.

Veritas NetBackup servers are given special considerations.

See “Default Host Group Hierarchy for Veritas NetBackup” on page 518.

When you create host groups, consider the following:

Report Scope	Reporting requirements drive how you organize your host group hierarchy. ■ Organize host groups so that when you generate reports, you can easily select the servers or a list of hosts that you want included in the report scope. ■ Create sub-groups of hosts to match your enterprise’s structure--that is, business units, departments, etc. ■ Account for multiple reporting views. For example, if you need to report on backups by geographical location and you also want to list backup status by operating system type, you will need to have hosts represented in both of these host groups.
User Access	Host groups provide the framework for limiting a user’s access to data. You will use home groups to limit access. See “Assign a User Home Host Group” on page 516.

Note: If you have a large number of hosts, refer to the following.

See “Automated Tool to Create Host Groups” on page 506.

Manage host groups

- See “Add Host Groups and Sub-Groups” on page 503.

- See “Host Group Membership” on page 507.
- See “Rename Host Groups” on page 511.
- See “Delete Host Groups” on page 512.
- See “Move Host Groups” on page 513.
- See “Assign a User Home Host Group” on page 516.

Hosts can be assigned to more than one group. When creating host groups, you start from the host or backup server and assign a group or you can start from the group and assign members.

For approaches to host group organization and a list of typical host group management tasks refer to the following.

See “Plan your host group hierarchy” on page 502.

See “Default Host Group Hierarchy for Veritas NetBackup” on page 518.

Add Host Groups and Sub-Groups

Host groups can contain zero or more hosts/backup servers and zero or more sub-groups. However, a host or backup server must always belong to a host group. You can create your host groups and sub-groups based on physical and logical relationships. For example a physical grouping could be location. Logical sub-groups under location could be HQ, East Coast, West Coast and Europe.

You can add Host Groups from the Host Group view and/or during the host/backup server creation process.

See “Plan your host group hierarchy” on page 502.

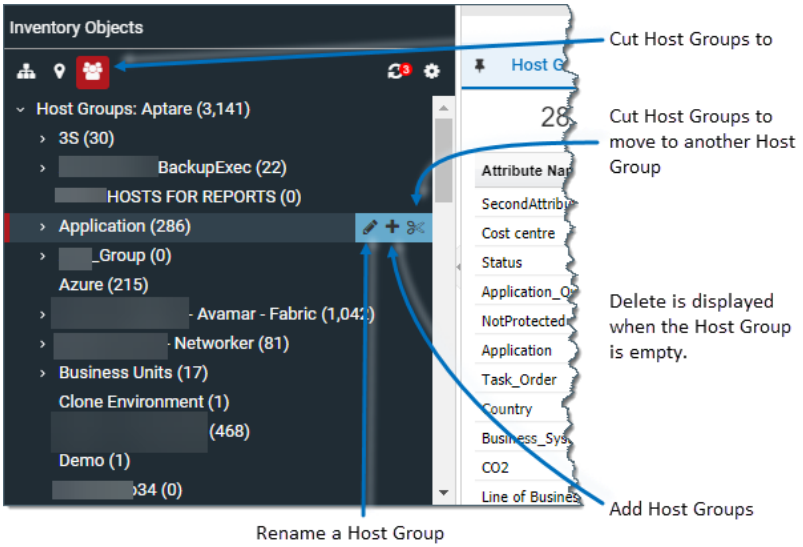
See “Adding and Editing Hosts and Backup Servers ” on page 521.

About Host Group Operations

In the Inventory, Host Group operations are available inline while in the Host Group view of the hierarchy panel.

- See “Add Host Groups” on page 504.
- See “Host Group Membership” on page 507.
- See “Rename Host Groups” on page 511.
- See “Delete Host Groups” on page 512.
- See “Move Host Groups” on page 513.

Delete is only available when the Host Group is empty. You can also roll over a Host Group to view the Group ID.



Add Host Groups

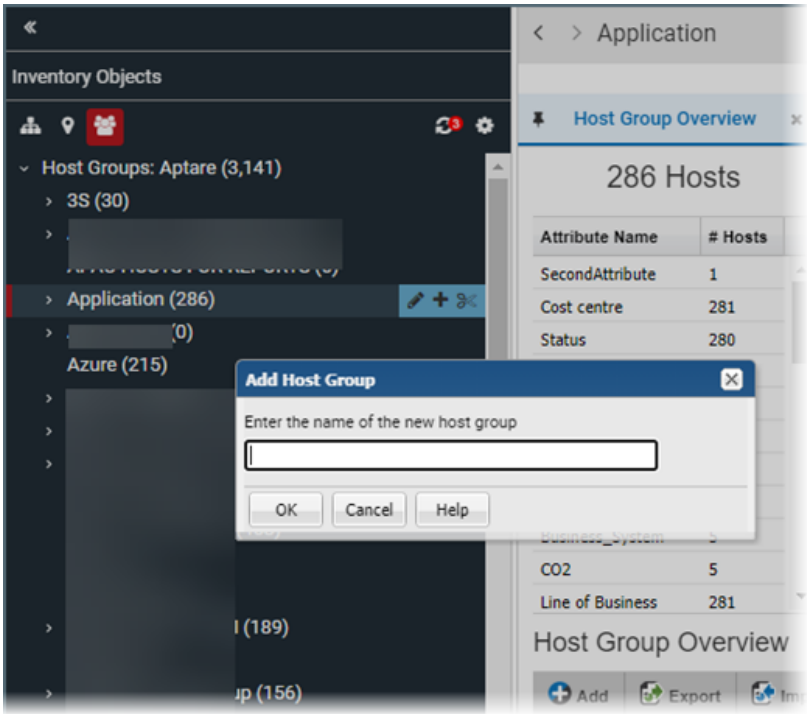
Once the Host Group is created, membership is assigned through Host management operations.

To add a host group

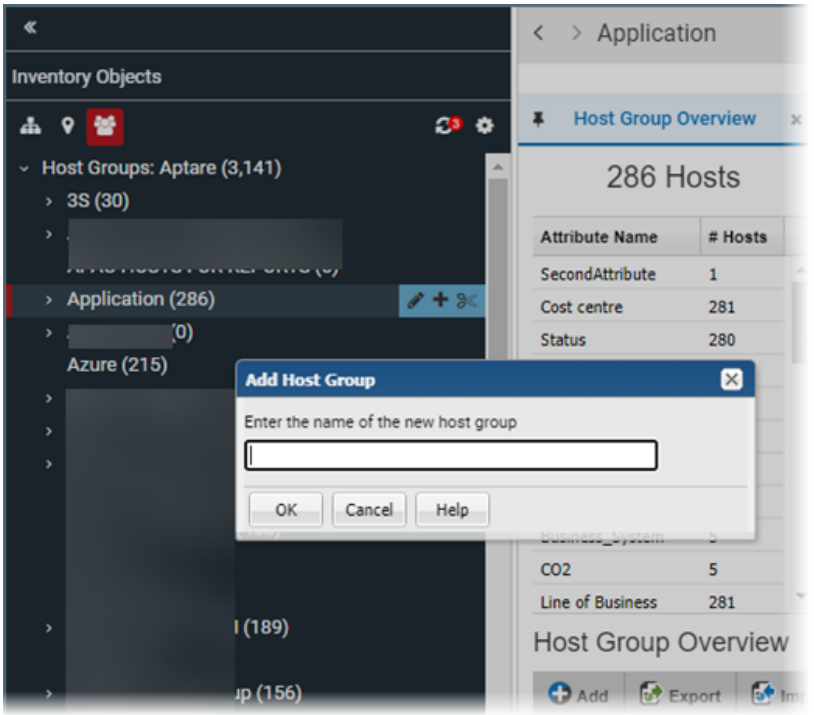
1. Click **Inventory.**
2. Change the view to Host Groups.

   See "Use host groups to organize your data" on page 73.
3. Choose a host group folder.

4 Click the **Add Host Group** icon. These icons are displayed inline beside the host group.You can also right-click. The Add Host Group dialog is displayed.



5 Enter a name for the new Host Group. Within the folder structure, the name must be unique. Folders on the same level cannot have the same name. The full path must be unique.



Note: The host group you selected in the hierarchy panel will be the parent of the host group you are about to add.

6 Click **OK**.

You can also add Host Groups during the host/backup server creation process.

See “Adding and Editing Hosts and Backup Servers ” on page 521.

Automated Tool to Create Host Groups

If you have a large number of hosts/backup servers, you can use the provided PL/SQL utilities for bulk processing. Instead of manually creating and organizing host groups through the Portal, you can run PL/SQL utilities to do the work for you. For more information about batch processing, search the documentation for Bulk Load Utilities.

Host Group Membership

You can assign hosts and backup servers to host groups from any view in the Inventory that displays the objects. A Host Group has no limit to the number of hosts it can contain. When you add a host you must assign a Host Group. After adding a host, you can edit Host Group membership.

- See "To assign a new host/server to a Host Group" on page 507.
- See "To copy a host/server into a Host Group" on page 508.
- See "To move a host/server from a Host Group into another Host Group" on page 509.
- See "To remove members from a host group" on page 510.

You can add Host Groups from the Host Group view and/or during the host/backup server creation process.

See "Adding and Editing Hosts and Backup Servers " on page 521.

To assign a new host/server to a Host Group

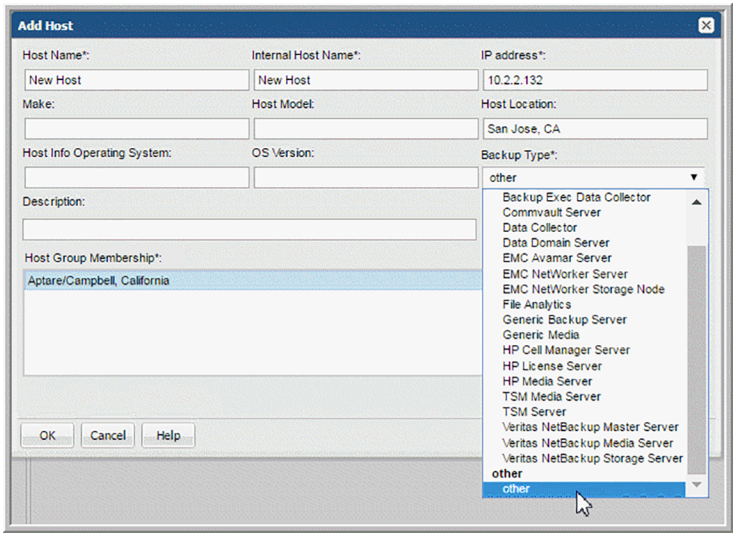
1. Click **Inventory.**
2. Change the view to Host Groups.

   See "Use host groups to organize your data" on page 73.
3. Choose a host group.
4. Switch to **List View** if required.

   See "Managing hosts and backup servers" on page 520.
5. Click **Add** to add a host to the selected host group. The **Add Host** dialog is displayed.
6. Specify the required information.
   - The Host Name is the name that will be displayed in reports. This is a required field.
   - The Internal Host Name must be an exact match for the name of the host as it is recognized by the product from which it was collected. This is a required field.
   - IP address of the host. This is a required field.
   - The Make, Host Model, Host Location, Host Info Operating System, and OS Version, if known.

   The **Backup Type**that you select from the drop-down list should characterize a specific host. If the host is not a backup server or backup client, select **Other**. This is a required field. The **Time Zones** field is displayed when the server is

designated as a Primary Server. The Time Zone setting is only available only for a host that is configured as a NetBackup Primary server.



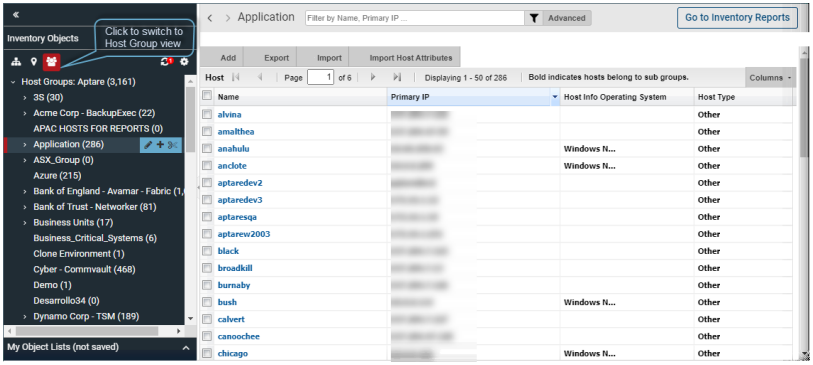
- Description - Enter a description relevant for your environment.

7 Click **OK**. The host is assigned to the Host Group you selected and the hierarchy panel is updated.

To copy a host/server into a Host Group

A host/server can belong to multiple Host Groups. When you copy a host/server from one Host Group to another, this retains the existing Host Group membership and adds the new membership.

1 In the Inventory, use the Hierarchy Toolbar and switch to Host Group view.



2 Search for or navigate to your hosts or backup servers.

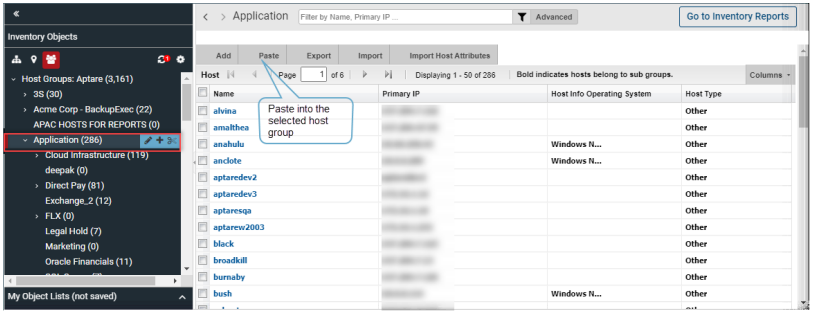
See “Navigating with search” on page 35.

3 Alternatively, select the Host Group to edit in the Hierarchy Panel. The members of the group are displayed in the grid.

4 Select the Hosts or Backup Servers you want to copy into/assign to a new Host Group.

5 Click **Copy**.

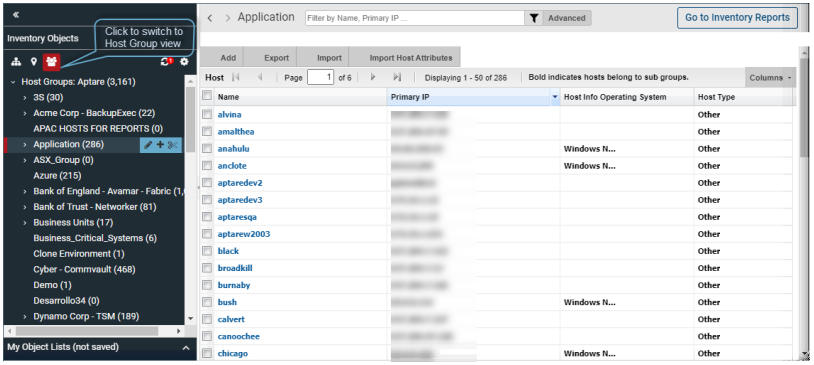
6 Select a Host Group in the Hierarchy Panel and click **Paste**. The hosts/backup servers you selected are added to the selected Host Group.



To move a host/server from a Host Group into another Host Group

Use Cut and Paste to move hosts/servers from one Host Group to another. This removes the existing Host Group membership and adds the new membership. Note, a host or backup server must always be assigned to at least one Host Group.

1 In the Inventory, use the Hierarchy Toolbar and switch to Host Group view.



2 Search for or navigate to your hosts or backup servers.

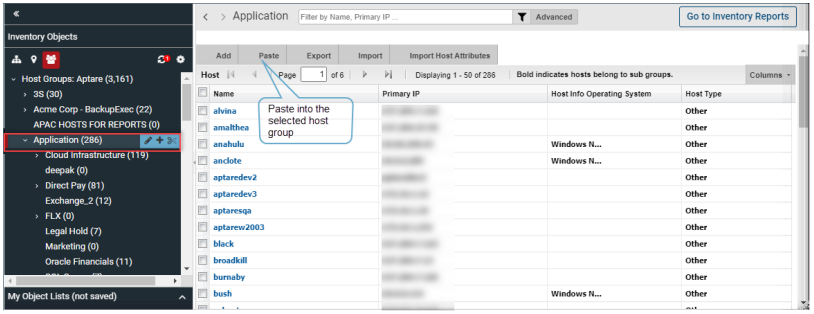
See “Navigating with search” on page 35.

3 Alternatively, select the Host Group to edit in the Hierarchy Panel. The members of the group are displayed in the grid.

4 Select the Hosts or Backup Servers you want to copy into/assign to a new Host Group.

5 Click **Cut**.

6 Select a Host Group in the Hierarchy Panel and click **Paste**. The hosts/backup servers you selected are added to the Host Group and removed from the original Host Group.



To remove members from a host group

1 In the Inventory, use the Hierarchy Toolbar and switch to Host Group view.

2 Select the Host Group to edit in the Hierarchy Panel.

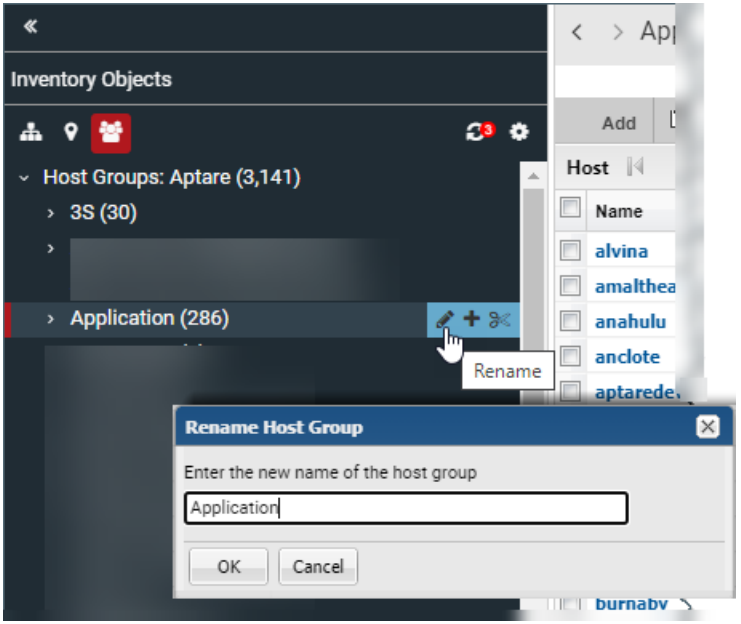
3. Select the Hosts or Backup Servers you want to **Remove** from a Host Group. The action bar displays available operations.
4. Click **Remove. Remove** unsubscribes the host/backup server from membership to the group. It does not delete the Host/Server or the Host Group. Hosts and backup servers must belong to at least one host group.

Rename Host Groups

To rename host groups:

1. Click **Inventory**
2. Change the view to Host Groups.

   See "Use host groups to organize your data" on page 73.
3. Choose a host group folder. Inline icons are displayed for Host Group operations.
4. Click the rename icon. The **Rename Host Group** dialog is displayed.



5. Enter the new name for the Host Group.
6. Click **OK.**

Delete Host Groups

You must delete or move individual hosts into another host group before you can delete a Host Group.

To delete a host group

1. Click **Inventory.**
2. Change the view to Host Groups.

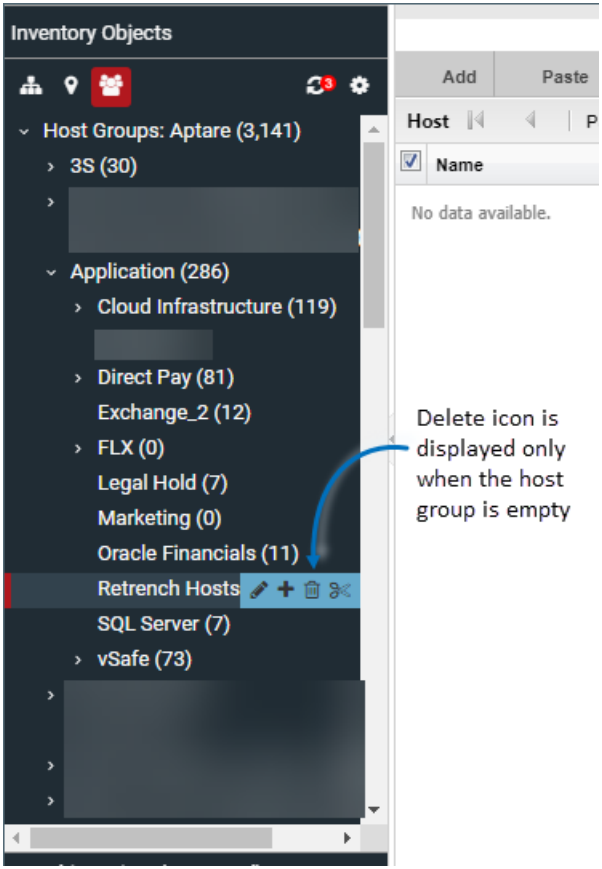
   See “Use host groups to organize your data” on page 73.

3 Search for or navigate to the Host Group to move.

See “Navigating with search” on page 35.

4 Delete the contents of each sub-group first. If there are sub-groups, you must delete the hosts from those groups first and then delete the group. Hosts can be deleted in bulk or individually.

Note: Hosts and backup servers must belong to at least one host group.



Move Host Groups

If you need to reorganize your host groups, you can move them to preserve the sub-group structures.

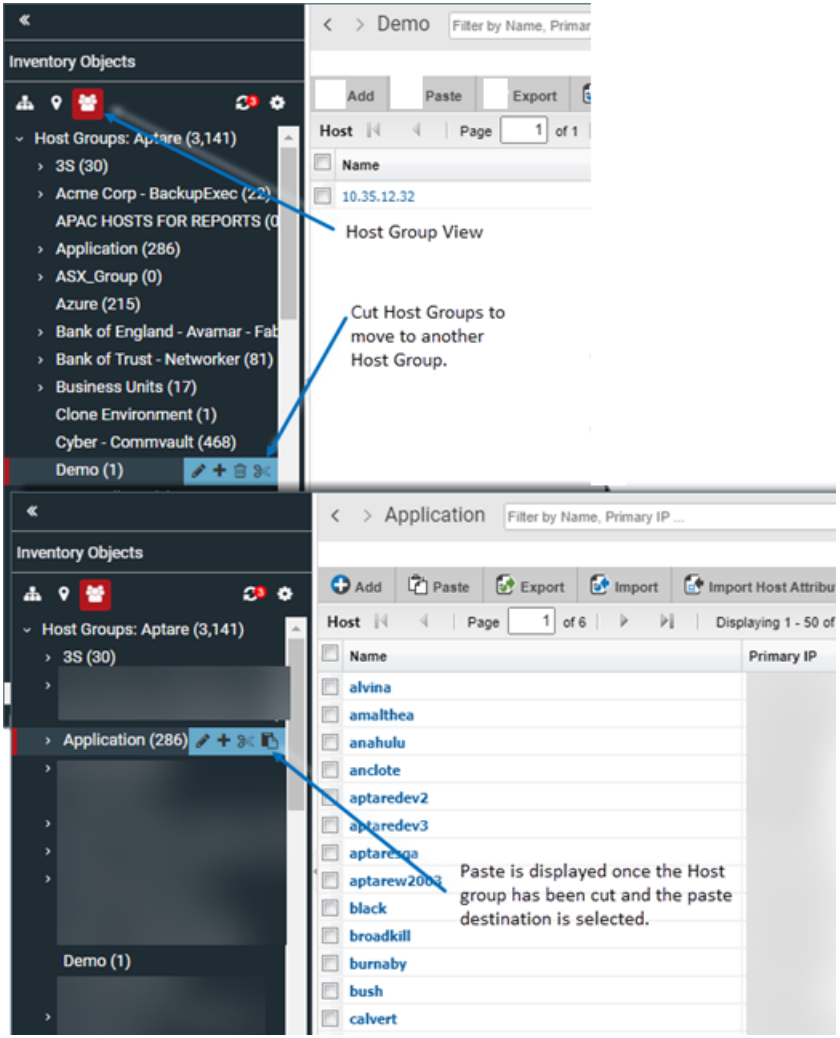
To move a host group

1. Navigate to the **Inventory.**
2. Change the view to Host Groups.

   See “Use host groups to organize your data” on page 73.
3. Search for or navigate to the Host Group to move.

   See “Navigating with search” on page 35.
4. Select the Host Group to move.

5 Click the **Cut** icon inline.



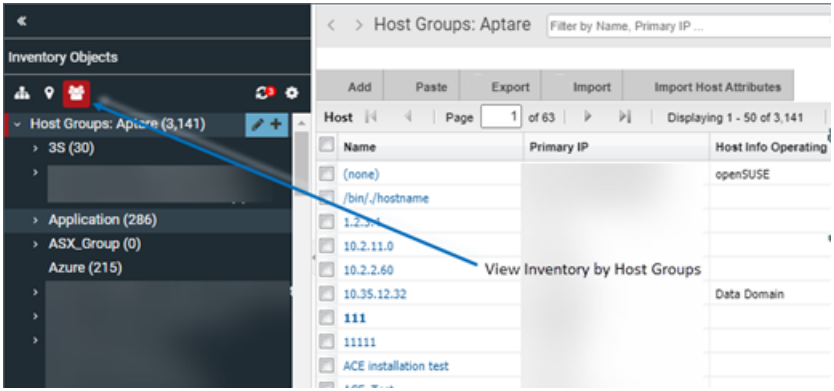
6 Select the destination group. The **Paste** icon is displayed inline.

7 Click **Paste.**

Find a Host Group ID

To identify the unique identifier associated with a host group, take the following steps in the Portal.

1. Navigate to the **Inventory.**
2. Click the **Host Groups** icon to switch the Inventory view.

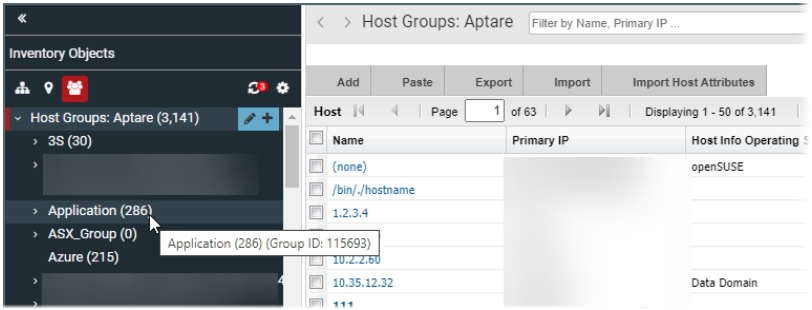


3. Verify the Host Group column is displayed on the grid, and optionally, use Advanced Filtering to locate the Host Group.

Note: The **Host Group** column, displayed in the Inventory for the Host Group management view, has sorting disabled to improve portal performance.

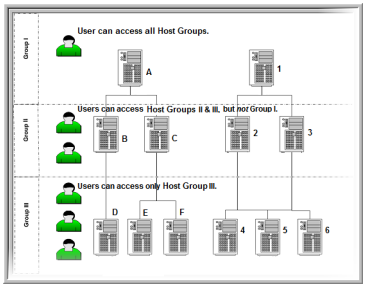
---

4. Hover your mouse over the Host Group folder in which your hosts reside. The Group ID will display in a tooltip.



Assign a User Home Host Group

When you assign a user to a host group, the user can only gain visibility to those hosts at their assigned level and below. This assigned level is called a home host group. The user will not see any host groups above their defined home group.



To assign a user to a host group

1. Choose **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Select the user name to which you want to assign a host group, and click **Edit**. The window displays the user's profile.
3. From the **Home Host group** list, select the host group to which you want the user to belong and click **OK.**

Manage Backup Servers

Backup servers are classified as a host. Although they are displayed as separate object in the Inventory, they are classified as special type of host. This means, all host report and management actions are valid for backup servers.

NetBackup Primary Servers

NetBackup Primary Servers have special considerations for setup and use.

- See "Configure the Time Zone for a NetBackup Primary Server" on page 517.
- See "Default Host Group Hierarchy for Veritas NetBackup" on page 518.

Configure the Time Zone for a NetBackup Primary Server

Whenever the Time Zone is modified, the system marks the Data Collector so that the updates are pushed to the Data Collector server. If the time zone is not explicitly configured for a NetBackup Primary Server, the system defaults to the time zone of the Data Collector server.

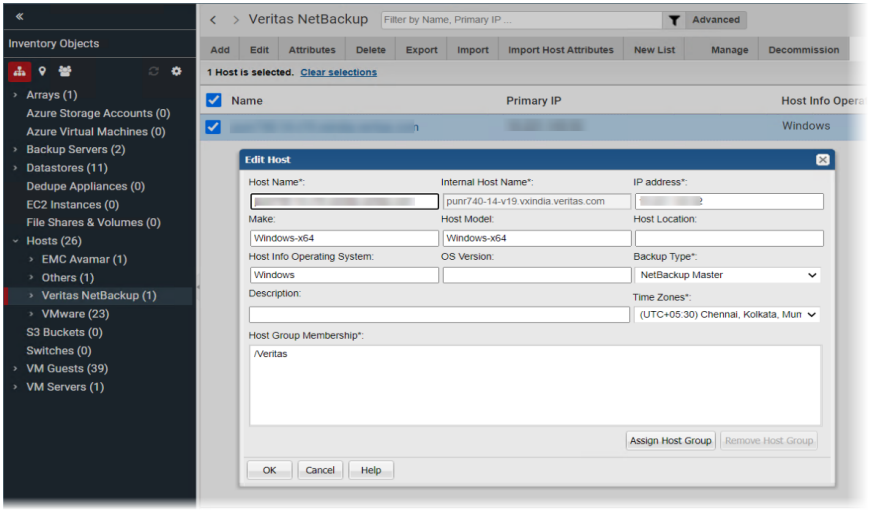
In reports, the date and time displayed for a backup transaction represents the date and time when the event actually happened.

Note: Currently, the Time Zone setting is available only for a host that is configured as a NetBackup Primary Server.

1. Search for the NetBackup Primary Server.

   See "Navigating with search" on page 35.

2. Select the specific server, when located, to edit the details.
3. Click **Edit.** The **Edit Host** dialog is displayed. The fields displayed are dynamic and Time Zones are displayed when the **Backup Type**, Primary Server is selected.



Default Host Group Hierarchy for Veritas NetBackup

During the initial installation, the Data Collector searches for all backup servers and clients, and assigns them to their respective default host groups. The result is a default hierarchy that consists of the following host groups:

- Primary Servers. If a client is being backed up by a backup server (or primary server), the Data Collector assigns the client to this host group. This group organizes clients under host groups based on the backup server name.
- Clients Not In Policy. If a client does not have a backup policy, the Data Collector assigns the client to this host group. This applicable to Veritas NetBackup only.
- Inactive Policy Clients. After the initial installation, the Data Collector may create this host group for clients that have inactive policies. For more information, See "Identifying Inactive Clients (For Veritas NetBackup)" on page 519.

This applicable to Veritas NetBackup only.

About the Primary Servers Group

The Primary Servers group represents the physical pool of backup servers. As you create additional host groups, you'll add backup servers/clients from this Primary Servers group into your new logical groups. This Primary Servers group is created automatically and updated each day.

Note: Do not create new host groups within this default primary group, as it is a system-maintained group, populated by Data Collectors. It is a best practice to retain this primary pool of servers. To set up logical groupings for reports, create additional host groups outside of this default group and add clients from the Primary Servers group.

Identifying Inactive Clients (For Veritas NetBackup)

Note: Beginning with release version 10.1, the host groups mentioned in this section are no longer automatically created. For legacy systems, if these host groups are no longer desired, a serverGroupCleanup utility can be used to re-locate the clients and then delete the system-generated group. See the *System Administrator Guide*.

There are two types of inactive clients, and in both cases these clients are not being backed up, maybe intentionally or unintentionally:

- Clients Not in Policy are not part of a policy at the time that a new Data Collector was added. Often these clients are resolved at the time of installation.
- Inactive Policy Clients were once part of a policy, but are not anymore.

To identify inactive clients

1. Navigate to **Inventory**.
2. Activate the Host Group view.
3. Select the **Veritas NetBackup** host group.

4 Select the **Clients Not in Policy** host group. If you notice that there are clients in this host group, either assign them to a policy--if they should be backed up--or delete them. In some cases, you might determine that a client need not be backed up and should be removed. In this case, refer to Deleting Hosts.

See "Deleting Hosts" on page 528.

5 Select the **Inactive Policy Clients** host group. If you notice that there are many clients in the **Inactive Policy Clients** host group and all these clients are supposed to be backed up by the same backup server, you probably accidently disabled the policy and need to enable it. In some cases, you might determine that a client does not need to be backed up, and should be removed. In this case, refer to Deleting Hosts.

See "Deleting Hosts" on page 528.

6 Give the Data Collector time to refresh its list of clients. Every night the Data Collector queries your backup servers for a list of clients and assigns clients to their appropriate host groups.

Managing hosts and backup servers

The **Inventory** enables you to quickly assess the status of your system's virtual and physical hosts by evaluating collected or imported details of all hosts from across your organization. It does not include discovered, but not collected from hosts.

The **Inventory List** view enables you to manage individual Hosts and groups of Hosts. Note, all operations are privilege based. These operations are also available from the **Host Group Overview** dashboard in the Inventory. The following management operations are only available for Hosts and Backup Servers:

- See "Adding and Editing Hosts and Backup Servers " on page 521.
-
- See "Decommission/Recommission Hosts and Backup Servers" on page 526.

The following operations are available for all objects, including hosts:

- See "Navigating with search" on page 35.
- See "Assign attributes in the inventory list view" on page 86.
- See "Export objects from the inventory list view" on page 93.
- See "Delete objects using the inventory list view" on page 94.

Adding and Editing Hosts and Backup Servers

Navigate to the Inventory and organize your hierarchy to display Hosts or Backup Servers.

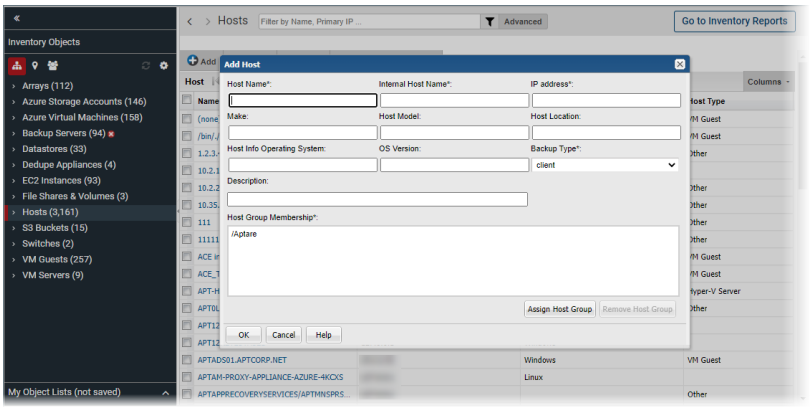
See “Hierarchy toolbar to organize your data” on page 63.

As a part of creation, you must assign Hosts/Backup Servers to a Host Group. After adding hosts or backup servers, the Refresh icon will display a badge to indicate the database has been updated with new objects. Click Refresh to display the additions in their appropriate categories in the Inventory view.

To add a host/backup server

1 Toggle to the List view with Hosts or Backup Servers displayed. The Hierarchy panel can be arranged in any configuration, but you must select Hosts/Backup Servers, to add a Host or Backup Server.

Note: When adding an EMC Data Domain Server, in the Inventory select **Hosts**, not Backup Servers.



2 Click **Add**. The **Add Host** dialog is displayed.

3 Specify the required information.

- The Host Name is the name that will be displayed in reports. This is a required field.
- The Internal Host Name must be an exact match for the name of the host as it is recognized by the product from which it was collected. This is a required field.
- IP address of the host. This is a required field.

- The Make, Host Model, Host Location, Host Info Operating System, and OS Version, if known.
- The **Backup Type** that you select from the drop-down list should characterize a specific host. Typically, this backup type is required when configuring a data collection policy. If the host is not a Backup server, select **Other**. This is a required field.

To edit host details

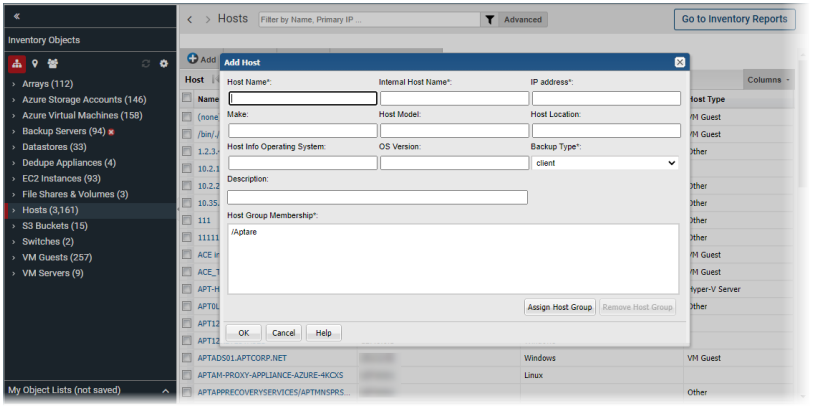
Navigate to the **Inventory** and organize your hierarchy to display Hosts.

See "Hierarchy toolbar to organize your data" on page 63.

As a part of editing, you can assign and remove hosts from Host Groups.

See "Host Group Membership" on page 507.

1. Toggle to the List view with Hosts displayed. The Hierarchy panel can be arranged in any configuration, but you must select Hosts/Backup Server to add a Host or Backup Server.
2. Select Hosts or a host group. The individual Hosts are displayed in the view pane.
3. Select the Host to edit. The **Edit** button is displayed once you select the host.
4. Click **Edit**. The **Edit Host** dialog is displayed.



5. Edit the details as required.
   - The Host Name is the name that will be displayed in reports. This is a required field.

- The Internal Host Name must be an exact match for the name of the host as it is recognized by the product from which it was collected. This is a required field.
- IP address of the host. This is a required field.
- The Make, Host Model, Host Location, Host Info Operating System, and OS Version, if known.
- The **Backup Type**that you select from the drop-down list should characterize a specific host. Typically, this backup type is required when configuring a data collection policy. If the host is not a Backup server, select **Other**. This is a required field.

Import Hosts

If your hosts are not part of a data collection process, you can still manually add them using the Import function. Import adds the data to the Portal database and displays the information in the Inventory. If a host already exists in your system, Import will update the host's details.

Before importing, create a comma-separated values (CSV) file of host data using the format specifications. The CSV file must be of UTF-8 type.

Note: This utility is designed to create and update hosts. It does not load or maintain relationships between hosts or host groups. See *Load relationships between hosts and host groups* section in the *IT Analytics Administrators Guide* for details on loading host relationships.

CSV Format Specifications (Host Import)

Create a comma-separated file that contains the following fields for a host, in the order listed in the following table. Note that each field in the CSV must have an entry, even if it is a null entry within the commas. Field values cannot contain embedded commas.

Name	Type	Value
/path_to_host_group	String	The path to the host group must be the full path from the root host group. The host group corresponds to your domain. Null values are not allowed. **Note:** If you are using the product in a multi-tenancy environment, you may not know the full path to the root host group. For security purposes, if there are multiple domains and host groups, you will only know the path within your host group. Only an administrator with a Super User role can provide the full path.
internal_name	String	The host name, typically how it is known in your data center, is a required field with a maximum of 128 characters. Null values are not allowed.
external_name	String	The host name, as you want it displayed in the Inventory, is a required field with a maximum of 128 characters. Null values are not allowed.
description	String	The host description, as it will be displayed in the Inventory, has a maximum of 256 characters.
location	String	Location of the host, a maximum of 64 characters.
ip_address	String	Host IP address, a maximum of 40 characters.
make	String	Make of the host, a maximum of 64 characters. Example: **Dell PowerEdge**
model	String	Model of the host, a maximum of 64 characters. Example: **4300**

Name	Type	Value
os_version	String	Operating system version of the host, a maximum of 128 characters. Example: **RHEL 2.1**
os_platform	String	Operating system platform of the host, a maximum of 128 characters. Example: **Linux**

EXAMPLE: hostImport.csv

```
/HDS/Finance Assets,Finance NBU Primary,Finance Backup,Main
backup server for the finance group,HQ New York,10.10.10.10,Dell,
64-bit, 10.0.10240,Microsoft Windows
```

Import Notes

- If the host already exists in the specified host group, the details are updated.

To import hosts from a CSV file

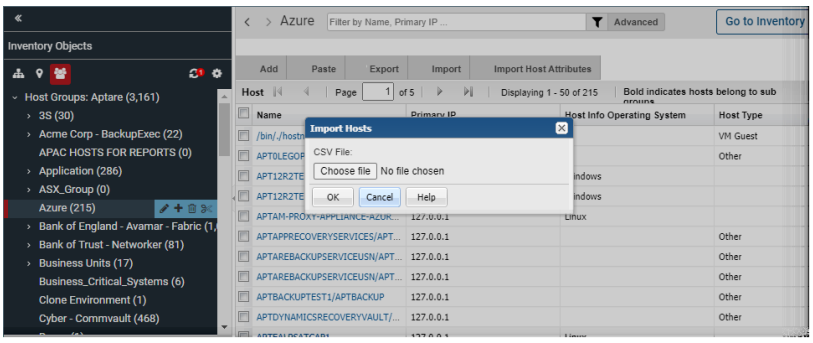
1. Prepare the CSV according to format specifications. The CSV file must be of UTF-8 type.

   See the section called “CSV Format Specifications (Host Import)” on page 523.

2. Select **Inventory**.
3. Navigate to Hosts and select a host group. This can be anywhere in the hierarchy structure that displays hosts.
4. Click **Import**.

Note: The Import function is controlled by privilege settings in the user profile.

See “Assigning user privileges” on page 578.



5 Click **Choose File** to navigate to the CSV to import.

6 Click OK. You must refresh to see the new hosts in the grid.

Decommission/Recommission Hosts and Backup Servers

Note: Prior to release version 10, the procedure for decommissioning a backup primary server directed you to change its host type to client. This previous method is no longer supported and, in fact, may result in data reporting issues. Use the procedure described in this section to decommission a host, a primary server, or any other server.

Under the following circumstances, you will need to decommission a host or primary server from the database, while still retaining its collected historical information for auditing purposes.

- Any Host: Frequently, large organizations decommission hosts and re-use the host IP addresses and/or host names. In this case, the host must be flagged in the database so that historical data remains, but the newly collected host with the same IP address and/or host name can be included in data collection and reporting.
- Backup Servers: When you decommission a backup primary server from your enterprise, you must let the Portal know that it is no longer available for data collection. After decommissioning a backup server, the active policy clients that are no longer being backed up will not count against your backup license count.

When you **decommission** a host or server, the following actions are taken to ensure that historical data remains intact:

- Renames the host in the database, according to the following format: **<host name>-decommissioned-<date>**
- Logs the decommissioning actions, including the user and date, in a database audit table. The table, apt_object_action_audit, can be accessed in the database using SQL*Plus or SQL Developer.

Note: Post decommissioning, when data collection encounters the re-purposed IP address or host name, a new host is created in the database.

When you **recommission** a host, the following actions are taken:

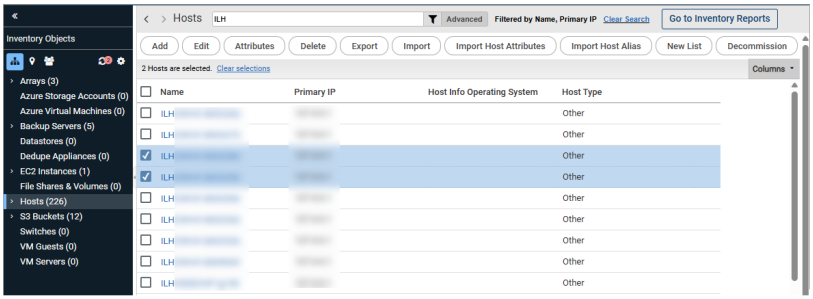
- Renames the host to its original name, discarding the labels that were appended when the host was decommissioned.

- Logs the recommissioning actions, including the user and date, in a database audit table.

Note: Recommissioning a host may not be successful, as data collection may have collected another host that has been brought online with the same name or IP address.

To decommission a host or a backup server

1 Select **Inventory**.



2 Click **Go to Inventory List**.

3 Enter your search criteria to find your target set of hosts or backup servers.

See “Filter the inventory list view” on page 85.

4 Select the items to manage. You can select all the items on a page by clicking the checkbox on the top of the management page. Note, only all the items displayed on a single page are selected.

5 Click **Decommission**. A confirmation dialog is displayed.

6 When you **decommission** a host or server, the following actions are taken to ensure that historical data remains intact:

- Renames the host in the database, according to the following format: **<host name>-decommissioned-<date>**
- Logs the decommissioning actions, including the user and date, in a database audit table.

Note: When data collection encounters the re-purposed IP address or host name, a new host is created in the database.

To recommission a host or a backup server

1. Select **Inventory**.
2. Enter your search criteria to find your target set of hosts or backup servers.

   See "Navigating with search" on page 35.
3. Select the items to manage. You can select all the items on a page by clicking the checkbox on the top of the management page. Note, only all the items displayed on a single page are selected.
4. Click **Recommission**. A confirmation dialog is displayed.
5. When you **recommission** a host, the following actions are taken:
   - Renames the host to its original name, discarding the labels that were appended when the host was decommissioned.
   - Logs the recommissioning actions, including the user and date, in a database audit table.

Note: Recommissioning a host may not be successful, as data collection may have collected another host that has been brought online with the same name or IP address.

Deleting Hosts

You may want to delete hosts from the reporting database for the following reasons:

- You do not want the hosts to appear in your reports.
- The host is inactive and should no longer be backed up.

Note: Warning: If you choose to **Delete a Host**, you are irretrievably deleting hosts from host groups and the reporting database. All related host historical data will also be permanently deleted from the database and unavailable in all reports. A pop-up window warns you of this action to prevent inadvertent deletions.

Note: When you delete a host, it removes everything related to the host except the VM Server. To remove the related VM Server, you just explicitly delete it in the Inventory. This prevents servers from being orphaned in the database.

In most cases, you should remove a host from a group, thereby un-linking it from its relationship with other hosts in that group. Also, you can easily move or add a host to other groups.

See “Decommission/Recommission Hosts and Backup Servers” on page 526.

See “Delete objects using the inventory list view” on page 94.

Manage host aliases

Managing host aliases can become cumbersome if there are multiple or duplicate aliases assigned to a host. This management option allows bulk import of host aliases through a CSV file with a check to ensure only unique aliases are imported and assigned to the host. You can also manually add unique host aliases to individual hosts present on the portal.

See “Bulk import host aliases ” on page 529.

See “View and manage aliases for individual host” on page 530.

Bulk import host aliases

Ensure you create a CSV file with the structure specified below and follow the considerations before you perform the actual bulk import.

Structure of CSV file for bulk alias import

The CSV file to be used for bulk alias import must have the following structure:

```
<domain>, <hostname>, <alias_hostname>
```

For example:

```
Enterprise, AsiaPacific, 123.123.123.123
```

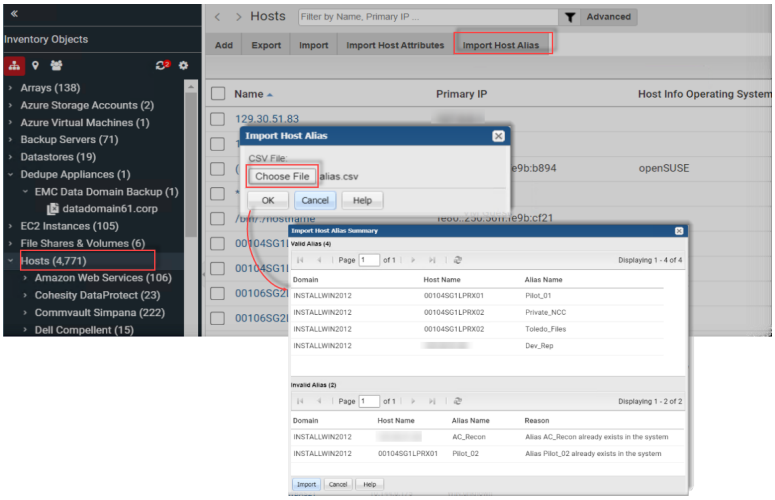
Considerations to import host aliases in bulk

A bulk import of host aliases succeeds provided the following requirements are met:

- Domain mentioned in the CSV must be added on the IT Analytics Portal.
- The hosts for which the aliases are being assigned through the CSV file must be present on the portal.
- Alias must be unique and unlike the host name. If the alias is duplicate, it is excluded from the import.
- Alias names must not contain special characters.

To import host aliases in bulk:

1. Prepare the CSV according to format specifications suggested above. The CSV file must be of UTF-8 type.
2. Select **Inventory**.
3. Navigate to **Hosts** and click **Import Host Alias**.
4. Click **Choose File** and add the CSV file. The file is analyzed and the list of valid and invalid aliases is displayed. You can hover your mouse pointer over the invalid alias names to see the reason for rejection in its tool tip.



5. Click **Import** to add the valid aliases.

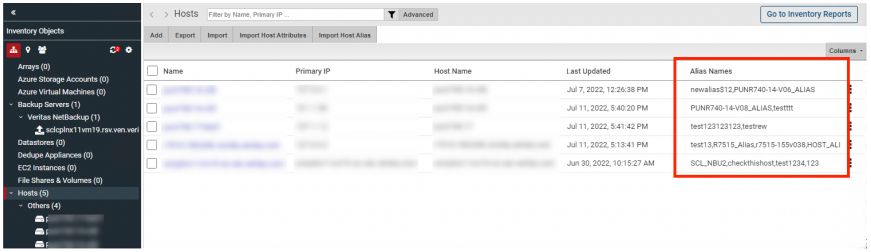
   A success message is displayed to confirm the alias import.

View and manage aliases for individual host

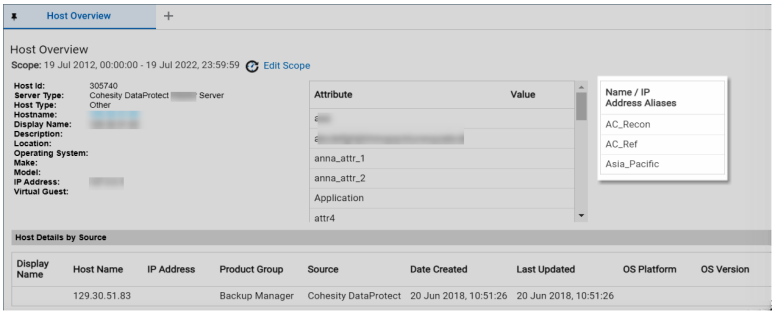
This section describes how you can view, add, edit, and delete aliases of an individual host.

View aliases assigned to a host

Host aliases can be viewed under the **Alias Names** column in the Inventory List View for all the Hosts. If the **Alias Names** column is not visible in the view, add it from the **Columns** list present on the view. You can also hover your mouse pointer in the **Alias Names** column to see the alias list in a tool tip.



Alternatively, you can also view the host aliases in the Host Overview report of the associated host. The Host Overview report opens when you click the hostname on the Inventory List view.



Add aliases to a host

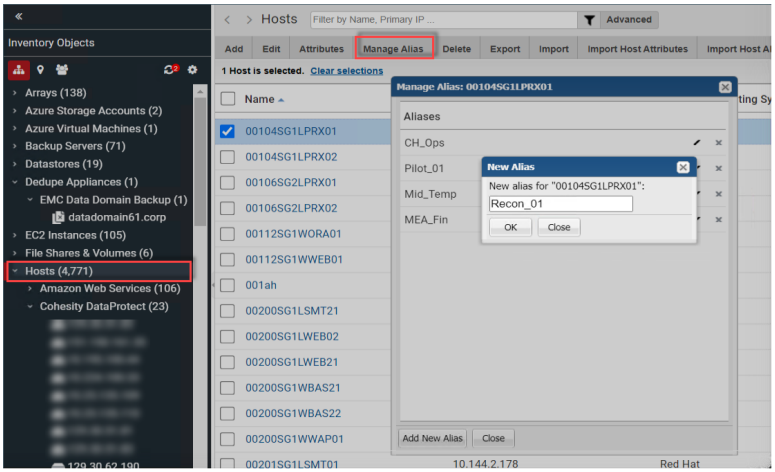
The steps below help you to add multiple aliases to a single host manually.

To add one or more aliases to a single host:

1. Open the Host Inventory view on the IT Analytics Portal.
2. Select the host for which you want to assign aliases and click **Manage Alias**.
3. On the **Manage Alias** pop-up, click **Add New Alias**.

4 Enter the new alias name and click **OK**. The name must be unique and unlike the host name. Also, avoid using special characters.

The new alias gets added to the **Manage Alias** pop-up .



5 Click **Close** on the **Manage Alias** pop-up. The new alias gets assigned to the respective host.

Edit a host alias

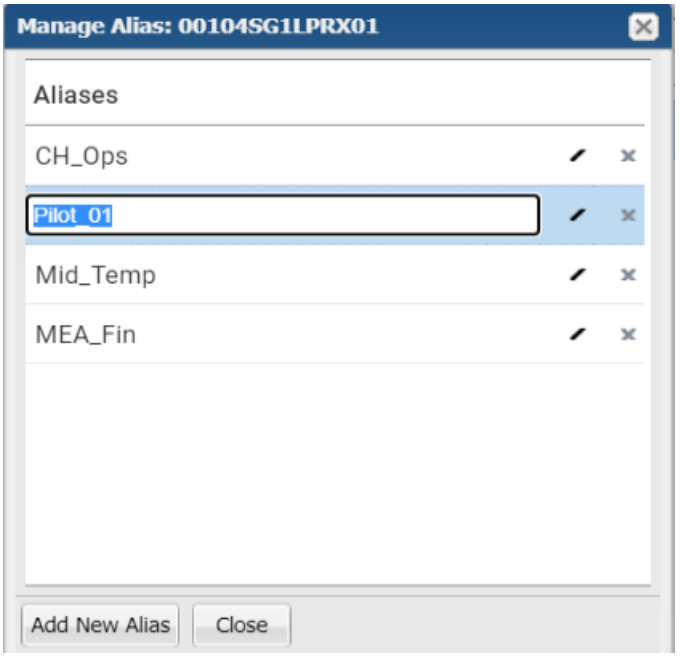
To edit a host alias:

1 Open the Host Inventory view on the IT Analytics Portal.

2 Select the host for which you want to edit the alias and click **Manage Alias**.

3 On the **Manage Alias** pop-up, click the edit icon against the alias name. The text entry becomes editable.

4 Add the required changes and click **Close** to exit the **Manage Alias** pop-up.



The changed alias is displayed under the**Alias Names** column

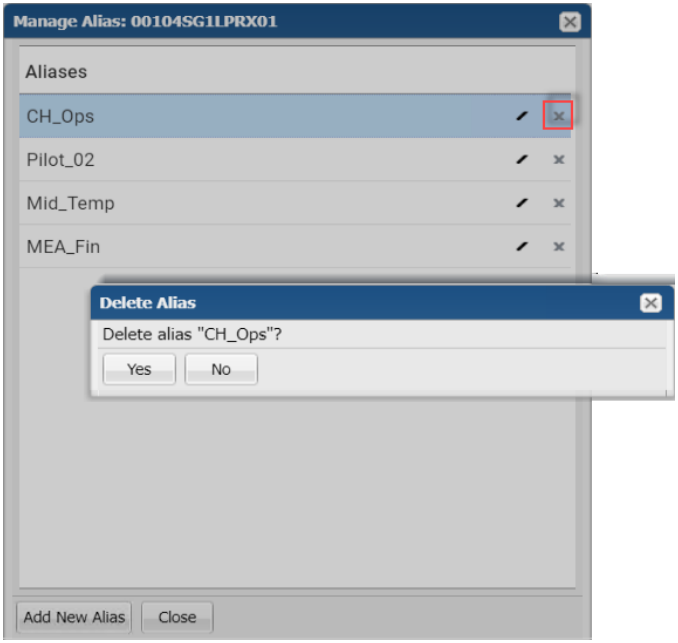
Delete a host alias

To delete a host alias:

1 Open the Host Inventory view on the IT Analytics Portal.

2 Select the host for which you want to edit the alias and click **Manage Alias**.

3 On the **Manage Alias** pop-up, click the cross icon against the alias name.



4 Click **Yes** on the confirmation message. The alias is deleted from the the **Manage Alias** pop-up.

5 Click **Close** to exit the **Manage Alias** pop-up and save changes.

The alias list is updated under the**Alias Names** column.

Chapter 13

Manage attributes and objects

This chapter includes the following topics:

- About attributes
- Attribute example
- Host groups vs attributes
- Bulk load utilities and attributes
- System attributes
- Attribute inheritance
- Manage attributes
- Set attributes on hosts
- Add attributes
- Edit or rename attributes
- Enable granular user access using authorization attributes
- Delete attributes
- Attribute naming rules
- Examples of attributes and values
- Understanding the attribute management view
- Attribute management view (Top-Level domain)

- Attribute management view (Sub-Domain)
- Export/Import dynamic templates with custom attributes
- About object maintenance
- Product specific objects for reporting
- Maintaining objects
- Assign attributes using object maintenance
- Search for objects by type in object maintenance
- Customize library objects
- Permanently remove devices, libraries, and drives

About attributes

Note: Starting with release version 10.0, attributes are no longer specific to a single object type. Attributes now apply to all object types, such as hosts, arrays, and switches.

- See " Manage attributes " on page 540.
- Attributes enable you to define a distinct data set based on a specific characteristic. Attributes represent logical relationships between objects and their relevant characteristics. This categorization mechanism aids reporting on clearly defined sets of data. Typically, you will create attributes to identify an additional data set for a report's scope or to use in the **Inventory** view to categorize data.
  See "Attribute example" on page 537.
- A default set of system attributes is included in the Portal software.
  See "System attributes" on page 538.
- Not all attributes can be modified. For example, System Attribute details (domain, name, and description) cannot be modified, but the values assigned to them can be revised. The Domain can also restrict who can modify attributes. **In the Portal, attributes that can be fully modified are displayed in bold**.
- Attributes directly relate to report scope. Report scope refers to the data selections that you make to include and exclude in a report. You can set the scope using several characteristics, including host groups, hosts, arrays, and attributes.

- See “Attribute inheritance” on page 539.
  See “Attribute naming rules” on page 550.

Attribute example

Attributes can be configured for objects to define a distinct data set, based on a particular characteristic. For example, you might set up a “maintenance_contract” attribute that you can associate with the hosts for which you have service coverage. With that set, you can generate a report to list all hosts with a maintenance contract by dragging the maintenance_contract attribute into the report scope.

Similarly, you might create an attribute to organize arrays by geographical location. In this case, for example, the report scope could include EMC Symmetrix and NetApp arrays and a Location attribute with a value set to London and another Location attribute with the value set to Edinburgh. In this example, four sets of data would be included in the report’s scope.

When you create an attribute, that attribute is automatically available for all object types, simplifying administration.

Note: Report templates created with the Dynamic Template Designer have access to only Array, Host, Library, and Drive objects and attributes.

Host groups vs attributes

Both **host groups** and **attributes** provide mechanisms for organizing data to facilitate reporting and inventory browsing. Use the following list to determine the best approach for your Portal environment.

When to use attributes (Preferred)

- Use attributes as the preferred organization mechanism because they span most object types and are not limited to hosts and host groups.
- Assign attribute values to a diverse set of objects (such as hosts, arrays, and switches) for easy Inventory access.
- Enable further filtering of host groups by assigning attributes to select hosts, such as backup clients.

When to use host groups

- Enforce security controls, limiting a user’s access to a subset of the hosts in the database. A user has an assigned home host group, which identifies the hierarchy of hosts a user can access.

- Group hosts that naturally belong together and typically will be reported on as a group.

Bulk load utilities and attributes

Bulk load utilities--that is, the scripts that import large sets of data into the database from comma-separated values (CSV) files--process attributes in the following manner:

If an attribute already exists for the attribute name that is being loaded, the utility uses the existing attribute and creates the corresponding attribute for the other object types. For example, if an array attribute named Criticality exists for an array, when the utility attempts to add an attribute with the same name, it will simply create the attribute for the remaining object types.

System attributes

A default set of system attributes is included in the Portal. These attributes cannot be deleted, however, you must create the list of values (LOV) for your enterprise.

See “ Edit or rename attributes ” on page 544.

See “Understand report data caching” on page 839.

When you assign a list of values to an attribute, the values are available to all objects. Typically, a list of values will be suitable for all object types (hosts, arrays, etc.). There may be cases where you require a list of values for one type of object and a separate list of values for another type of object. For these cases, you would need to create a user-define attribute with a different name.

The system attributes shipped with the product include:

- Application - Enables grouping by Application for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Business_Unit - Enables grouping by Business Unit for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Data_Center - Enables grouping by Data Center for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Department - Enables grouping by Department for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.

- Environment - Enables grouping by Environment for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Location - Enables grouping by Location for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Organization - Enables grouping by Organization for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Org_level1 - Enables grouping by Org_level1 for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Org_level2 - Enables grouping by Org_level2 for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute
- Org_level3 - Enables grouping by Org_level3 for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Org_level4 - Enables grouping by Org_level4 for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Owner - Enables grouping by Owner for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.
- Region - Enables grouping by Region for certain inventory configurations and report scope selectors, for all objects. User-defined values must be assigned to this attribute.

Attribute inheritance

Attribute inheritance is relevant primarily in multi-tenancy environments where domains are used to partition the database to maintain security controls. In this configuration, a hierarchical structure provides a parent-child relationship that controls access to data and also a structure for inheriting configurations from parents.

See “Understand report data caching” on page 839.

Manage attributes

You can add any number of attributes, and they can represent any characteristic. Ensure that you provide a descriptive name for the attribute and for the attribute values so that all users understand the intent of the attribute. The **Inventory** view uses attributes as a way for you to organize your view.

To manage attributes

1. Select **Admin > Advanced > Attributes**.
2. See “Add attributes ” on page 541.
3. In the Attribute Management page, several functions are available, depending on the domain ownership: Add, Edit, Delete.
   - See “Add attributes ” on page 541.
   - See “ Edit or rename attributes ” on page 544.
   - See “Delete attributes” on page 549.

Set attributes on hosts

To further define a report’s scope and search criteria, you can define custom attributes with discrete values. Attributes are characteristics of your hosts. You can add any number of attributes, and they can represent any characteristic. These attributes provide an additional method for including and excluding data in a report.

Attributes provide a way of defining a set of objects that share a certain characteristic. Attributes represent logical relationships between objects and their relevant characteristics. Typically, you’ll set up attributes to aid in defining the scope of a report. For example, you might set up a “maintenance contract” attribute that you can associate with the hosts for which you have service coverage. Or, you might create an attribute to organize hosts by geographical location so that the administrators responsible for the hosts at each corporate location can generate reports for their specific sites.

You can select hosts in bulk and assign or modify attributes associated with them. Use Search and Advanced Filters to create a results set with the hosts you’d like to modify. For example, you can create a search query to find all Windows 2008 R2 systems as reported by NetBackup. When the search results are displayed, you can quickly apply a custom attribute such as patch applied.

Add attributes

When you add an attribute, a validation process ensures that the attribute does not exist in the domain's single hierarchy path. Duplicate attribute names are allowed only in sibling domain hierarchies.

See "Attribute inheritance" on page 539.

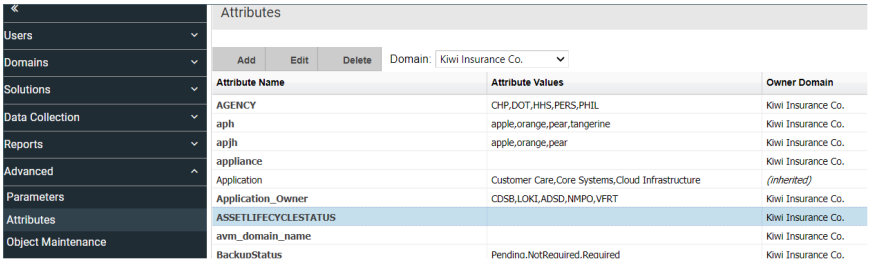
See "Attribute naming rules" on page 550.

See " About attributes " on page 536.

See "Understanding the attribute management view" on page 552.

To add attributes

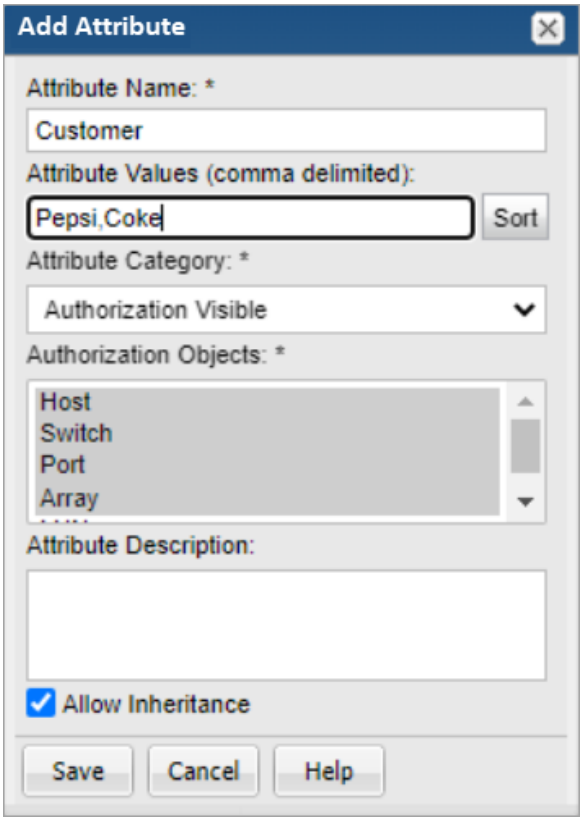
1 Select **Admin > Advanced > Attributes.**



2 From the drop-down list, select the domain to which you want to add the attribute. Your domain is displayed by default. Typically, only one domain is available and this domain selection is required only for multi-tenancy environments, such as Managed Services Partners (MSPs). When you add or delete attributes, you do so globally for your domain and all child domains.

See "Understand report data caching" on page 839.

3 Click **Add**. The **Add Attribute** dialog is displayed.



Note: Set `portal.supportAuthzAttributes=true` in the **Customs Parameter** of **System Configuration** to view **Attribute Category** drop down and **Authorization Objects** list box.

---

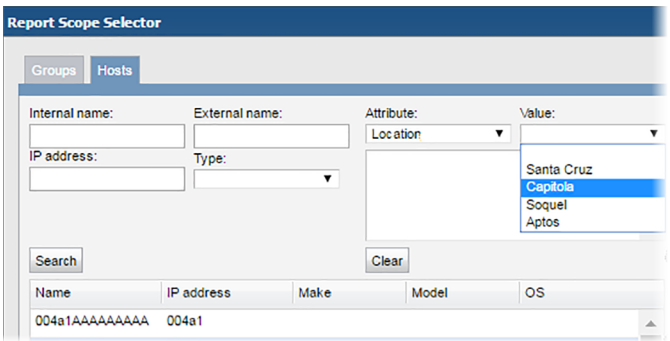
4 Enter a name.

See "Attribute naming rules" on page 550.

5 Enter a comma-separated list of values for the Attribute. Values must be unique and are case-insensitive. The order in which you enter the values, is the order they are listed for selection in the Portal.

For example, if you enter the values Santa Cruz, Capitola, Soquel and Aptos for the **Location** attribute, that is how they are listed for selection when you assign them in the **Report Scope**. This unsorted order is determined when you create the values. The following screen shot shows an example of the unsorted values.

You can choose to have the values sorted. By sorting, the values will be displayed alphabetically in the Portal.



6 Click **Sort** to alphabetically sort the values when displayed in the Portal. Attributes and values are displayed in:

- The Inventory.
  See “Set attribute values in the inventory list view” on page 88.
- Report Scope.
  See “Configure report scope with attributes” on page 116.
- Object Maintenance.
  See “Assign attributes using object maintenance” on page 557.

7 Specify the **Attribute Category** based on the description below. This option is visible provided Attribute Authorization privilege is enabled for you.

- **Authorization Visible**: Keeps the attribute visible for authorization in the attribute list for objects and it is visible in the scope picker to narrow the scope of the report like other attributes.
- **Authorization Hidden**: Keeps the attribute visible only for authorization.

8 Select objects from **Authorization Objects**. The authorization attributes will be assigned to the selected objects.

9 Enter an optional description. The description will be displayed in the Dynamic Template Designer when you are creating report templates.

10 Uncheck the **Allow Inheritance** if you do not want child domains to inherit this attribute's values. The **Allow Inheritance** checkbox is checked by default to set a flag that enables child domains to inherit the attribute's values.

When inheritance is disabled, users in sub-domains will be able to see the attribute, but they will not be able to see values in the **Inventory** or in a report scope. However, an administrator at the sub-domain or parent level can override the values of the attribute, making the attribute's values available to the sub-domain.

See "Override inherited attribute values" on page 837.

11 Click **OK** in the **Add Attributes** dialog. The new attribute is displayed.

Edit or rename attributes

See "Attribute naming rules" on page 550.

See " About attributes " on page 536.

See "Understanding the attribute management view" on page 552.

To edit attributes

Not all attribute details can be modified. For example, System Attribute details cannot be modified, but the list of values (LOV) assigned to them can be changed. The Domain can also restrict who can modify attributes. **Attributes that can be fully modified are displayed in the Portal in bold**.

If the user-defined attribute was created within your Domain, the attribute name, values, and description can all be modified. Most environments have only a single Domain. However, for multi-tenancy environments, where a hierarchy of Domains is used to partition data, you are not permitted to modify attribute names in child Domains. You can modify the inheritance flag and the list of values of an attribute that was inherited from a parent.

See "Attribute naming rules" on page 550.

See "Attribute inheritance" on page 539.

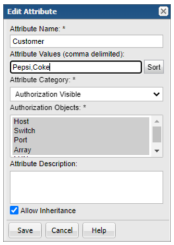
Note that when you modify the name, a validation check is made to ensure that the name you enter does not conflict with an existing attribute in your Domain hierarchy.

1 Select **Admin > Advanced > Attributes**.

2 (Multi-tenancy/multi-domain environments) From the drop-down list, select the Domain in which the attribute resides.

3 Select the attribute. Bold attribute names indicate that the attribute can be fully edited.

4 Click **Edit**.



5 Modify the name and/or description.

6 Modify the **Attribute Values**. These must be comma-separated and unique. They are case-insensitive. The order in which you enter the values, is the order they are listed for selection in the portal.

- For example, if you enter the values Santa Cruz, Capitola, Soquel and Aptos for the **Location** attribute, that is how they are listed for selection when you assign them in the **Report Scope**. This natural order is determined when you create the values. You can choose to have the values alphabetically sorted.

Note: Attribute values can also updated using the **Import Host Attribute** function available from the **Inventory**.

See “Import host attribute values” on page 90.

---

7 Click **Sort** to alphabetically sort the values when displayed in the Portal. Attributes and values are displayed in:

- The Inventory.
  See “Set attribute values in the inventory list view” on page 88.
- Report Scope.
  See “Configure report scope with attributes” on page 116.
- Object Maintenance.
  See “Assign attributes using object maintenance” on page 557.

Note: For System Attributes, only the list of values and the **Allow Inheritance** flag can be modified. For attributes inherited from a parent, you can modify only the list of values and the inheritance flag. You are not permitted to modify attribute names in child Domains.

See "Attribute inheritance" on page 539.

See "System attributes" on page 538.

8 Specify the **Attribute Category** based on the description below. This option is visible provided Attribute Authorization privilege is enabled for you.

- **Authorization Visible**: Ensures the attribute is visible for authorization in the attribute list for objects and it is visible in the scope picker to narrow the scope of the report like other attributes.
- **Authorization Hidden**: Ensures the attribute visible only for authorization in the attribute list for objects.

9 Select objects from **Authorization Objects**. The authorization attributes will be assigned to the selected objects.

10 In the **Edit Attribute** dialog, click **Save**.

Enable granular user access using authorization attributes

Authorization attributes provide a more granular control in providing access to various inventory objects for a single user or a user group. While enabling the access to specific objects such as backup servers, switches, and ports in the inventory, you can restrict the access to only specific servers, switches, or ports nested within the object categories. Thus, the inventory view for a user or a user group can be controlled from allowing access to an entire object to permitting access to only a specific set of entities nested under specific objects.

The basic requirements that enable the use of authorization attributes are as follows:

- The portal.supportAuthzAttributes attribute must be set to true. This step is essential to enable all the menus and UI elements related to authorization attributes on the portal UI. See the section called "Set portal.supportAuthzAttributes to true" on page 547.
- Authorization attributes must be already created on the IT Analytics Portal. See "Add attributes " on page 541.
- Assign authorization attributes to the portal users or user groups. See the section called "Assign authorization attributes to users" on page 547.

Set portal.supportAuthzAttributes to true

As a prerequisite you must set the **portal.supportAuthzAttributes** to **True**.

To set the attribute value:

1. On the IT Analytics Portal, go to **Admin** tab > **Advanced** > **System Configuration** and click the **Custom Parameters** tab.
2. Select **portal.supportAuthzAttributes** and click **Edit**.
3. Set its value to **True** and save the attribute.

This configuration is essential to enable the tab from where admin users can assign the authorization parameters to various users or user groups..

Assign authorization attributes to users

This procedure assigns the authorization attributes to the users. This equips the users with a more granular control over the objects.

Note: Ensure your objects are already assigned with the attributes that you plan to associate with the users in this procedure.

To assign authorization attributes to users:

1. On the IT Analytics Portal, go to **Admin** tab > **Users** > **Users and Privileges**.
2. Select a user from the user list and click the **Authorization Attributes** tab.
3. On the **Authorization Attributes** window, specify the attribute values based on the descriptions below.
   - **All Values**: Select to allow the user to see all the data objects assigned with any of the attribute value.
     For example, if an attribute Fruit has values Apple, Banana, and Orange, the user can see the data objects assigned with any attribute values.
   - **Unassigned Values**: Select to allow the user to see all the data objects without any assigned attributes or values.
   - **Values**: Enter specific attribute values separated by commas. The user can view only those data objects assigned with the specified values. If your

entries fail to match the attribute values, you will see an error while saving the changes.

4 Click **OK** to save the changes.

After assigning the authorization attributes, the user is able to view the objects in the Inventory view after logging on to the IT Analytics Portal. Ask the user to restart the portal in case the user had already logged onto the portal before you made the changes.

Note: You can follow the same procedure mentioned above to assign authorization attributes to user groups.

Note: If a user is authorized for the attributes using individual access permissions **OR** group access permission, the user is authorized to the attributes in any of the cases.

Use case of authorization attributes

The below use case can help you understand how the selection of **All Values** and **Unassigned Values** checkboxes can impact the object visibility of the users.

For example, consider two authorization attributes Customer and Site with values specified below.

Table 13-1 Authorization attributes

Authorization attribute name	Values
Customer	Coke, Pepsi
Site	SFO, NY

The inventory ports on the IT Analytics Portal are assigned the attributes and their values as below.

Table 13-2 Attribute assignment

Port Name	Customer attribute value assigned	Site attribute value assigned
Port1	Pepsi	
Port2		SFO
Post3	Pepsi	SFO

Table 13-2 Attribute assignment *(continued)*

Port Name	Customer attribute value assigned	Site attribute value assigned
Port4	Coke	
Port5	Coke	NY
Port6		NY
Port7		

The port visibility of various users assigned with the Customer and Site attributes and depending on the selection of **All Values** and **Unassigned Values** checkboxes is explained below.

Table 13-3 Port (object) visibility for users assigned with authorization attribute

User	Customer			Site			Visible Ports
	All Values	All Unassigned	Values	All Values	All Unassigned	Values	
User A	Yes			Yes			Ports 3 and 5
User B	Yes	Yes		Yes	Yes		All ports
User C		Yes			Yes		Port 7
User D			Pepsi			SFO	Port 3
User E			Pepsi	Yes			Port 3
User F			Pepsi				None
User G			Pepsi		Yes		Port 1
User H			Pepsi			NY	None
User I			Coke			NY	Port 5

Similar behavior is observed if you assign authorization attributes and their values to user groups instead of individual users.

Delete attributes

See “Attribute naming rules” on page 550.

See " About attributes " on page 536.

See "Understanding the attribute management view" on page 552.

Note: When you delete an attribute, it will continue to be available as a column in the Inventory view. Refresh the Inventory to access the most up-to-date list of attributes.

To delete attributes

1. Select **Admin > Advanced > Attributes**.
2. (Multi-tenancy/multi-domain environments) From the drop-down list, select the **Domain** from which you want to remove the attribute.
3. Select the attribute. Click **Delete**.
   - If the attribute is in use by a Dynamic Template, you will be prompted to confirm that you really want to delete the attribute. If you choose to delete an attribute that is in use, the template will no longer work as designed.
   - See "System attributes" on page 538.
     cannot be deleted.
   - Attributes inherited from a parent can only be deleted by an administrator of the Domain where the attribute was created.

Attribute naming rules

Adhere to the following rules when creating attribute names. Attributes are validated against these rules so that there are no conflicts in the database, such as duplicates or the use of Oracle reserved words.

- Limit the length to 30 characters.
- Begin the name with an alphabetic character.
- Use only ASCII, alpha, numeric, or underscore characters in the name. Spaces and special characters other than underscores are not allowed in attribute names, although they are allowed in the list of values (LOV) for an attribute.
- Names are not case-sensitive.
- Do not use Oracle reserved words. See http://docs.oracle.com/cd/E15817_01/appdev.111/b31231/appb.htm. To list the Oracle reserved words, use this SQLPlus query at the command line:

```
SQL> SELECT * from v$reserved_words;
```

- Attribute names within a domain hierarchy must be unique.

Examples of attributes and values

See “Attribute naming rules” on page 550.

See “ About attributes ” on page 536.

See “Understanding the attribute management view” on page 552.

Table 13-4 Table 1 Example Attributes and Attribute Values

Attribute Name	Possible Values	Purpose
Application	SAP Exchange	Data based on the software application running on the host.
Asset_Tag	0001234 0001235	For asset management purposes, perhaps you want to report on backup servers by asset tag.
Backup_Server	BackupServer1 BackupServer2	You will certainly want to report on backup servers/clients based on the backup server that backs up the backup server/clients' user data. Backup Server is the most common attribute, which is why the Portal creates a default group to represent this characteristic.
Business_Unit	Marketing Accounting	Backup servers/hosts often contain backups of data owned by users from specific business groups (for example, Marketing).
CPU	Opteron UltraSPARC	If you need to know how your backups are performing on your backup servers with specific CPUs, simply run reports based on this attribute.
Location	Americas Asia	If you are responsible for hosts in a region, you can select a scope for your region. It may make sense to set up host groups by geographical location or, as an alternative, create an attribute to group hosts by location.
Country_Code	004 248	You can be very specific about the location of hosts that you have spread throughout the world.

Table 13-4 Table 1 Example Attributes and Attribute Values *(continued)*

Attribute Name	Possible Values	Purpose
Host_Type	Production Test Server	This attribute can represent production vs. test machines. Data on production systems is critical to your business. Test data is important, too, but you might want to know how data is being produced on your production systems particularly.
SysAdmin	Alix Emily	Hosts are managed by this person.
OS	Linux Windows Mac	If you need to roll out patches for a particular operating system, you can quickly determine when the user data on those hosts will be backed up. Your values can be general or specific (for example, Windows 10).

Understanding the attribute management view

The Attribute Management page displays Attributes that are available for a specific domain. Multiple domains are available in multi-tenancy environments.

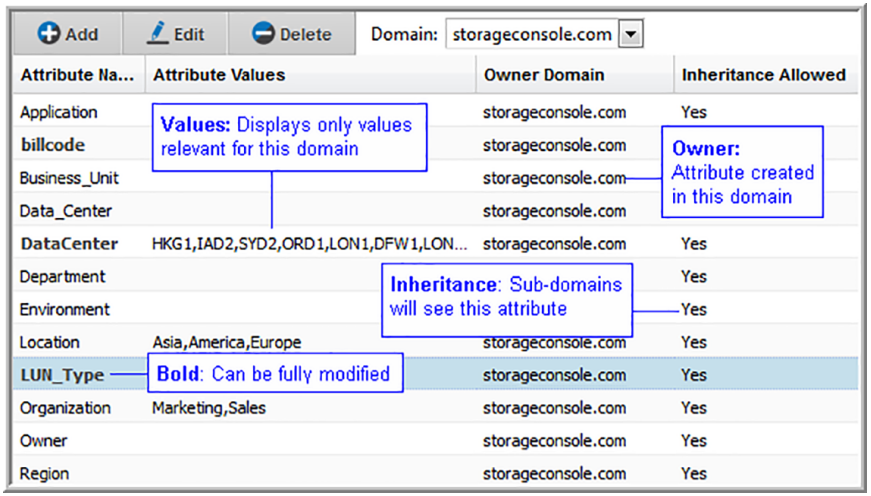
See “Attribute inheritance” on page 539.

Attribute management view (Top-Level domain)

The top-level domain has visibility into all the attributes that are available for the selected domain and for all sub-domains.

Top-level domains have an Attribute Management view with the following characteristics.

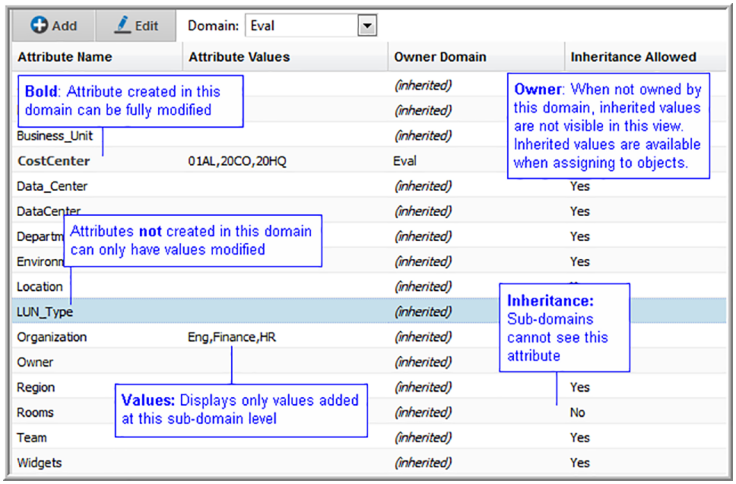
- Attributes owned by the domain: Both attribute names and values can be modified. These attribute names are shown in bold.
- Values: Only values relevant to the selected domain are displayed in the view.
- Inheritance: If allowed, sub-domains will have access to this attribute and its values.



Attribute management view (Sub-Domain)

Domains below the top-level domain have an Attribute Management view with the following characteristics.

- Attributes owned by the domain: Both attribute names and values can be modified. These attribute names are shown in bold.
- Values: Only values relevant to the selected domain are displayed in the view.
- Inherited attributes (not owned by the domain): Inherited values are not visible in this view, however, values that have been added for this domain are displayed. All values (both inherited and domain-specific) will be available when assigning attribute values to objects.
- Inheritance: If allowed, sub-domains will have access to this attribute and its values.



Export/Import dynamic templates with custom attributes

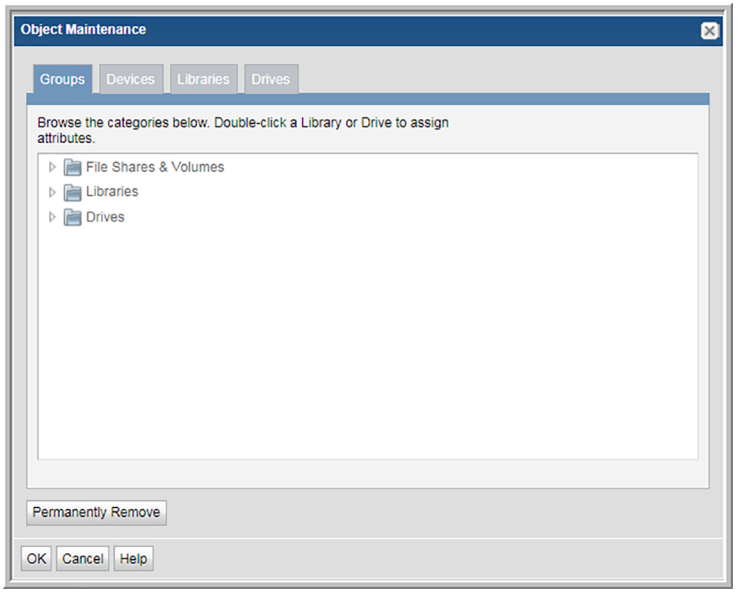
System attributes can be found in any release version 10.00 (or higher) environment. This set of attributes is common to the platform. User-defined attributes, however, can be unique to a portal. Therefore, when a custom report template, developed with the Dynamic Template Designer, embeds attributes in the logic, the distribution of the template needs to take attributes into account.

- See “System attributes” on page 538.
- See “Export a dynamic template that contains attributes” on page 293.
- See “Import a dynamic template that contains attributes” on page 293.

About object maintenance

The database contains a variety of objects on which reports can be generated. These objects can also have attributes associated with them to further refine the scope of a report. The Object Maintenance tool displays object types in your environment under a Groups tree structure. Each tab in the Object Maintenance window enables object searching by type.

To access the Object Maintenance dialog, select **Admin > Advanced > Object Maintenance**.



See “ Maintaining objects ” on page 556.

Some objects are maintained through the Inventory view. From the Inventory view you can delete instances of objects, and add and edit attributes for:

- Hosts
- Backup Servers
- Arrays
- Switches
- VM Guests
- VM Servers
- Dedupe Appliances

See “Manage objects in the inventory list view” on page 85.

See the *Licensing Guide* for details on removing objects from the license count.

Product specific objects for reporting

The objects that are available for a report’s scope are dependent on the product/feature on which you are reporting, as listed in the following table.

See “Search for objects by type in object maintenance” on page 560.

Note that the objects available for the report are dependent on the underlying product on which you are reporting. For example, certain Capacity Manager reports may report only on arrays and not hosts.

Product	Available Objects
Capacity Manager	■ Hosts ■ Arrays
Virtualization Manager	■ Hosts
Fabric Manager	■ SAN Fabrics/Switches ■ Hosts ■ Arrays
File Analytics	■ Devices (Shares and Volumes)
Backup Manager	■ Hosts ■ Libraries ■ Drives
Replication Manager	■ Hosts ■ Arrays

Figure 13.2 Product-Specific Objects

Maintaining objects

The Object Maintenance window enables several object maintenance functions:

- See “Assign attributes using object maintenance” on page 557.
  - Objects can have attributes associated with them to enable a specific selection of objects in a report’s scope. Attributes can also be used to organize objects in the **Inventory** view.
- See “Search for objects by type in object maintenance” on page 560.
  - Libraries, Drives, or Devices --have their own search windows.
- See “Customize library objects” on page 561.
  - Modify the name displayed for Tape Libraries.
- See “Permanently remove devices, libraries, and drives” on page 561.
  - Permanently deletes the object from the database. All historical information will be deleted from the database.

Some objects are maintained through the **Inventory** view.

See “Manage objects in the inventory list view” on page 85.

See “About object maintenance” on page 554.

See “Product specific objects for reporting” on page 555.

Assign attributes using object maintenance

Attributes enable you to define a distinct data set based on a specific characteristic. They give you the ability to refine a report’s scope. Attributes can also be used to categorize objects within the **Inventory** view. The object type determines where attributes can be assigned.

Object Type	Product Area for Management
Arrays	Inventory
Backup Servers	Inventory
Dedupe Appliances	Inventory
Devices (shares and volumes)	Object Maintenance for searching only; attributes cannot be assigned
Drives	Object Maintenance
Hosts	Inventory
Libraries	Object Maintenance
Switches	Inventory
VM Guests	Inventory
VM Servers	Inventory

Prerequisite: Prior to assigning an attribute to an object, you must create a list of values for the attribute. To configure attributes that can be assigned to objects, refer to managing attributes.

See “ Manage attributes ” on page 540.

To assign attributes in Object Maintenance

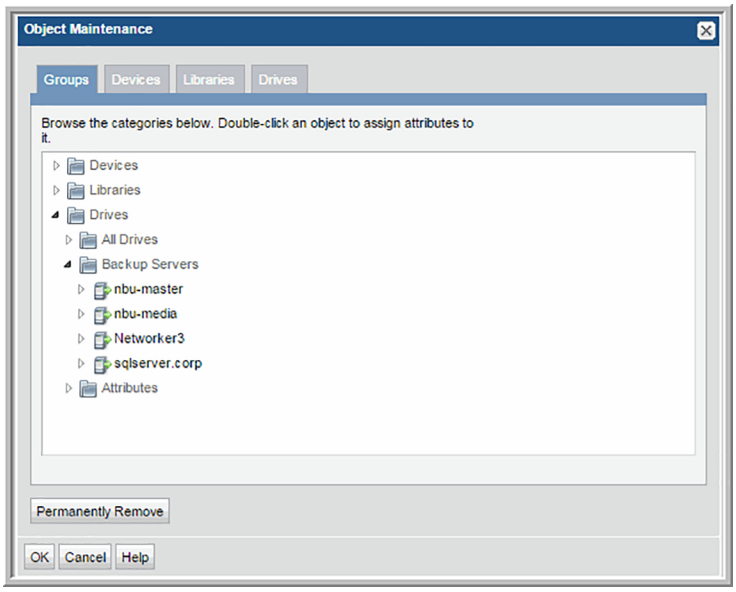
Within Object Maintenance, the following object types can have attributes associated with them: **Libraries** and **Drives**. Assigning attributes to Devices (Shares and Volumes) is **not** supported.

See “Manage objects in the inventory list view” on page 85.

1 Select **Admin > Advanced > Object Maintenance**.

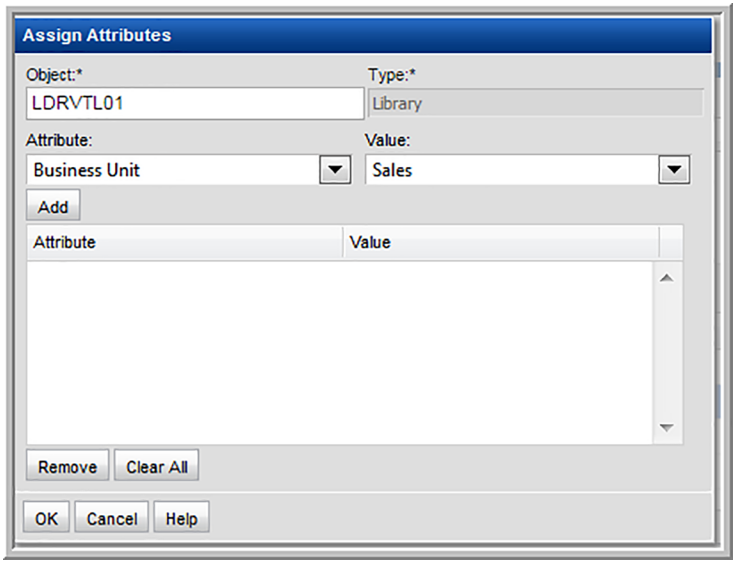
2 Find the object that you want to modify using one of the following methods:

- Expand the Groups tree to find a specific object.



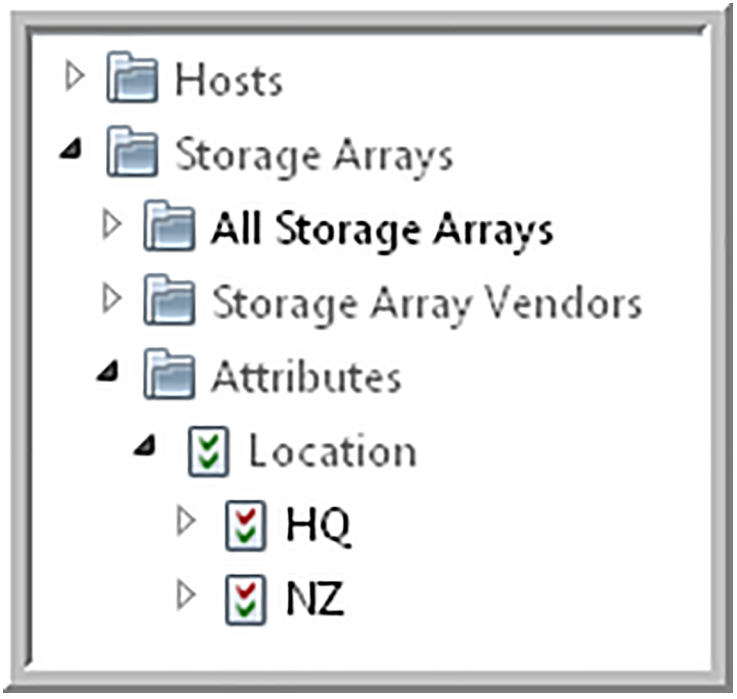
- Click a tab-- Libraries or Drives--to use the search window for that object type.

3 Double-click an object to launch the **Assign Attributes** window.



4. Select the Attribute and associated Value from the drop-down lists and click **Add**.
5. Click **OK** to save this configuration to the database.

   Now, when you generate a report for this object, a list of attributes will be available for you to generate a report based on that attribute.

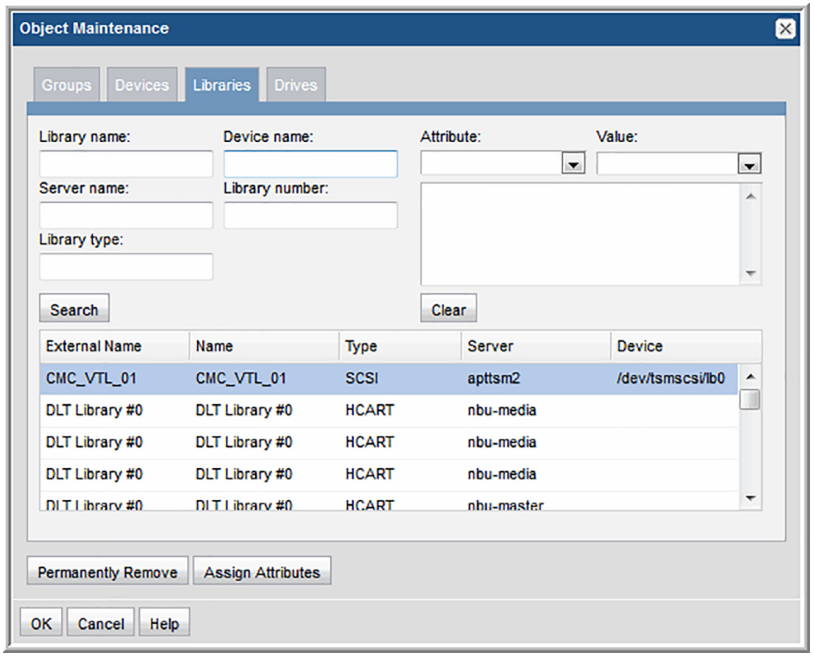


Search for objects by type in object maintenance

Object Maintenance manages the following object types: devices, libraries and drives. All other object types are managed in the **Inventory** view.

1. Select **Admin > Advanced > Object Maintenance**.
2. Click a tab-- Libraries, Drives, or Devices--to use the search window for that object type.

Example Search Window:



The asterisk (*) wildcard is supported for searches.

Customize library objects

In some environments, it is desirable to customize Tape Library names that are displayed in reports.

1. Select **Admin > Advanced > Object Maintenance**.
2. Click the **Libraries** tab and click **Search**.
3. Double-click a Library in the list.
4. Modify the name shown in the **Object** field and click **OK** in the **Assign Attributes** window.

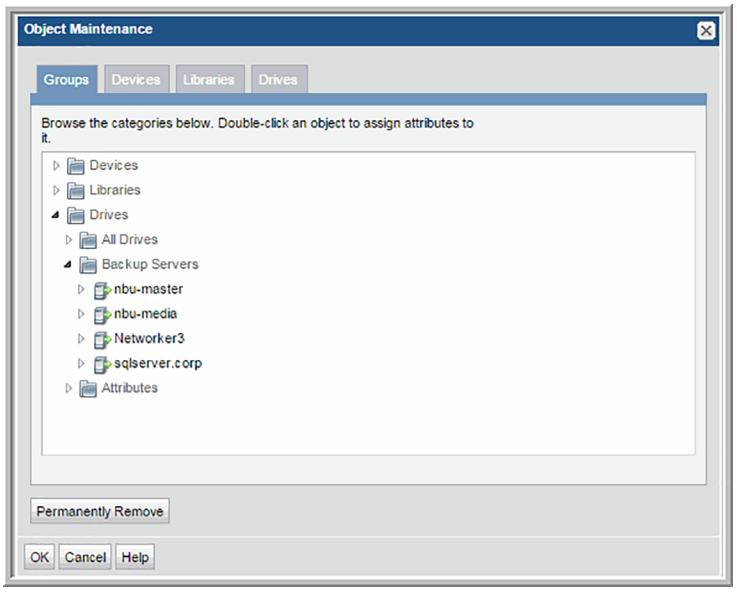
Permanently remove devices, libraries, and drives

Object Maintenance manages the following object types: devices, libraries and drives. All other object types are managed in the **Inventory** view.

Currently, the objects that can be removed through the Object Maintenance tool are:

- Devices
- Libraries
- Drives

1. Select **Admin > Advanced > Object Maintenance**.
2. Find the object that you want to modify using one of the following methods:
   - Expand the Groups tree to find a specific object.



   - Click a tab-- Devices, Libraries, or Drives--to use the search window for that object type.
3. Select an object in the list and click **Permanently Remove** in the Object Maintenance window. You are prompted to verify this deletion.

Note: CAUTION: When you permanently remove an object, all historical information is also deleted from the database.

Chapter 14

Provide Portal access and user privileges

This chapter includes the following topics:

- Providing user access to the portal
- Creating portal user accounts
- About user types
- Selecting user groups for users
- Working with user groups
- Creating user groups
- Assigning users to user groups
- About user privileges
- Assigning user privileges
- Setting user group privileges
- Enabling new product report templates
- Importing reports
- Enabling cloud privileges
- Granting access to template designers
- Enabling all privileges in a category
- Managing users and user groups

- Searching for users and user groups
- Editing user accounts
- Impersonating user accounts
- Setting / Resetting passwords
- Account lockout
- Managing user group home pages (Administrator)
- Removing portal users
- Deactivating user accounts

Providing user access to the portal

To provide users access to the Portal, perform the following sequence of steps:

1. Set up host groups.

   See "Add Host Groups" on page 504.

2. Understand the different user types.

   See "About user types" on page 565.

3. Create a portal account for each portal user.

   See "Creating portal user accounts" on page 565.

4. Create user groups.

   See "Creating user groups" on page 570.

5. Provide each user a welcome email, outlining the user's username and default password. Include instructions about how the user can log on to the portal. Mention that the portal prompts the user to change the default password.

   See "Assigning user privileges" on page 578.

Note: For Managed Services Partners and Administrators: Multiple browser windows on the same computer share a session data, when different account IDs are logged into the Portal from separate browser windows. For Managed Services Partners and Administrators, who may be setting up environments for clients, multiple logins may cause unpredictable results in generated reports and other operations because of the shared session data. You can workaround this issue by using multiple browser types. For example, if testing accounts, login with one account to Google Chrome and another with Mozilla Firefox.

Creating portal user accounts

This procedure assumes that you identified the host group to which the user must belong, and that the host group or subgroup is available. Using host groups (or a home group), dictates how wide a net a user can cast when generating report data. When you create a user account, you specify a home host group for that user. If not, go to

See "Plan your host group hierarchy" on page 502.

About user types

There are three portal user types:

User Type	Rights
Administrator	Manage user accounts and set up host groups at or below the Administrator's assigned group. An Administrator can create both End User and Administrator accounts, but only within the Administrator's home group. In an MSP (Managed Services Provider) environment, each client has Administrator accounts that have access only to the client's domain and only the host groups within that domain. **Note:** Portal upgrades will automatically enable privileges for newly added reports and the display of the **Inventory** view including all objects, for all Administrators. Refer to the release notes for the list of reports and features introduced in a specific product release.

User Type	Rights
Super User	A Super User's privileges cannot be revised by any other user. Everything an Administrator can do in addition to the following: ■ Access the entire Portal host group hierarchy from top to bottom regardless of the user's group assignment. See "Plan your host group hierarchy" on page 502. ■ Manage Oracle table space. ■ Define and manage server backup cycles. ■ Create both End User and Administrator accounts for any group within the host group hierarchy. ■ Access all default and user-generated reports. ■ View **New** and **Updated** badging on system report templates when they are made available. See "About badging" on page 102. ■ Impersonate a user profile. See "Impersonating user accounts" on page 595.
End User	■ Those features for which privileges have been granted by the Administrator. The end user will only be able to utilize those features at or below their assigned home group. ■ An End User can create only End User accounts within the user's own home group (domain).

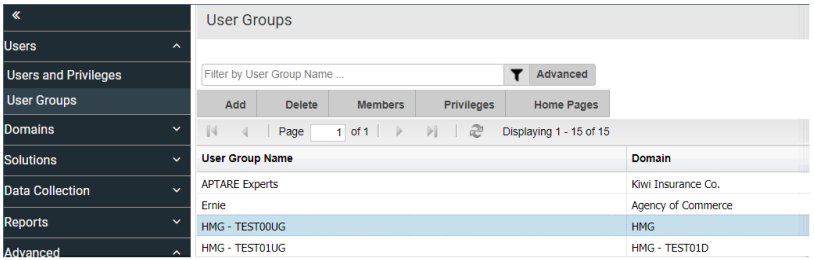
To create a user account

When you create a user account, you add user details and a password. Users can modify their passwords and some of their profile information, however, they cannot modify their access privileges. Access privileges and group membership is setup through a separate operation.

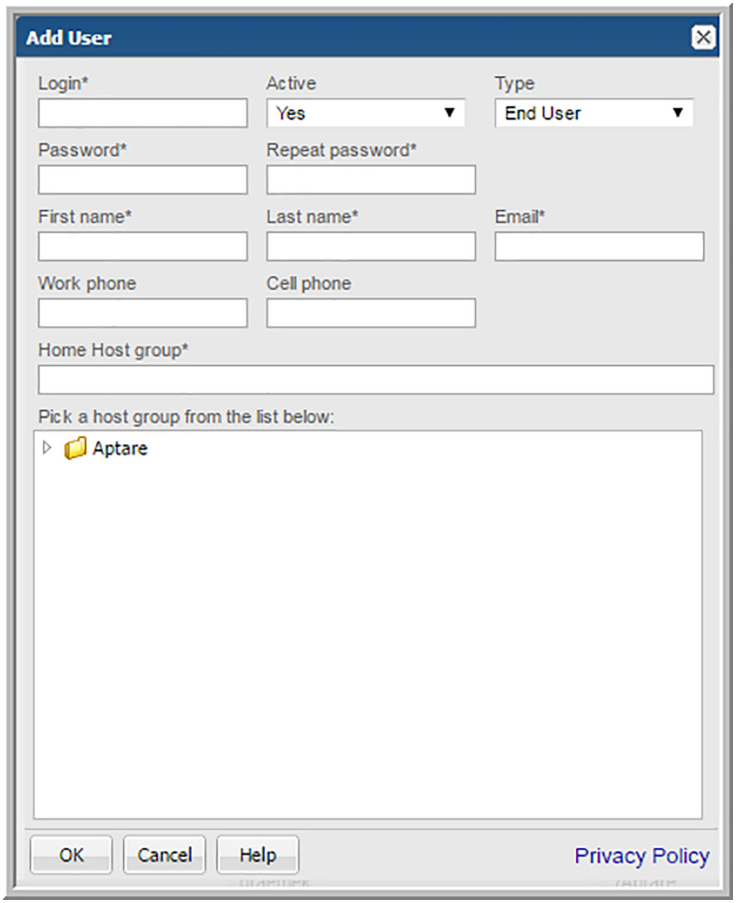
Preferences for locale, number and date formatting are set by the user using the user account menu.

See "Manage your profile and set a language preference" on page 772.

1 Select **Admin > Users > Users and Privileges**. The window displays all Portal users.



2 Click **Add** to create a new user. The **Add User** dialog is displayed.



3 Complete the fields.

- Required fields are denoted with an asterisk. **First name** and **Last name** fields are limited to 64 characters.
- The **Login** ID should be in an email format, as required by LDAP. This field is limited to 128 characters.
- From the selector, select the User Type. For a definition of the privileges associated with a user type,
  See "About user types" on page 565.
- Select a Home Host group. A home host group is a host group to which a user belongs. This limits a user's access to data. A user can access any host groups that are lower in the hierarchy from the home host group.

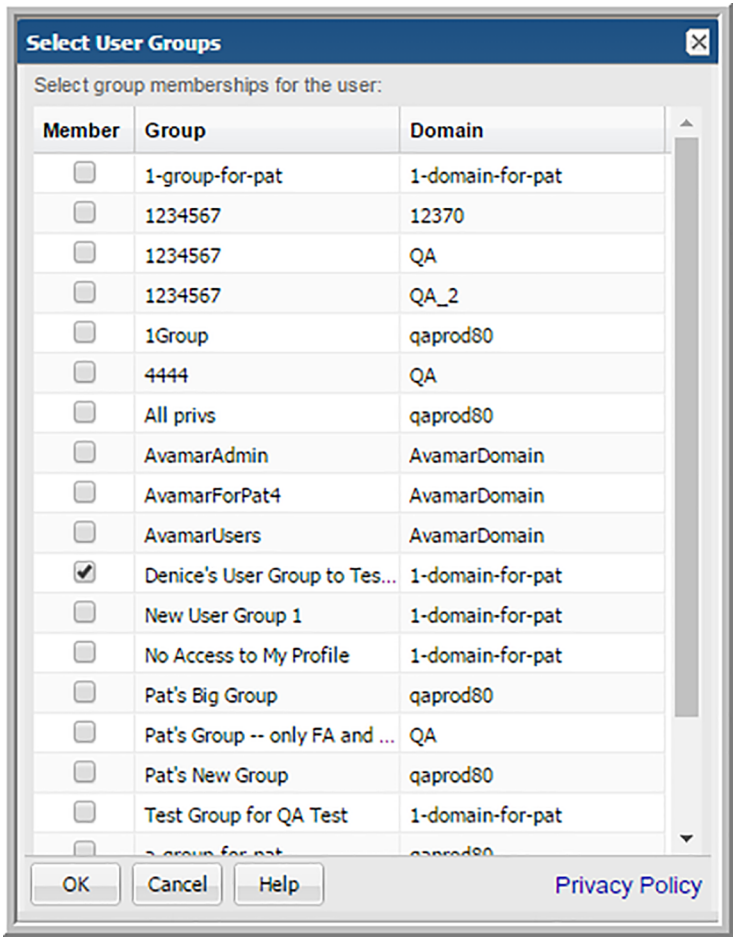
4 Click **OK** to create the user.

Selecting user groups for users

User groups provide an efficient way of managing many users at once. You can assign privileges to a group and they are propagated to the users in that group.

See "Working with user groups" on page 569.

1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Select a user.
3. Click **Groups**. The **Select User Groups** dialog is displayed.



4. Click the groups to assign membership to the selected user.

Working with user groups

User groups provide an efficient way of managing many users at once.You can assign privileges to a group and they are propagated to the users in that group. For example:

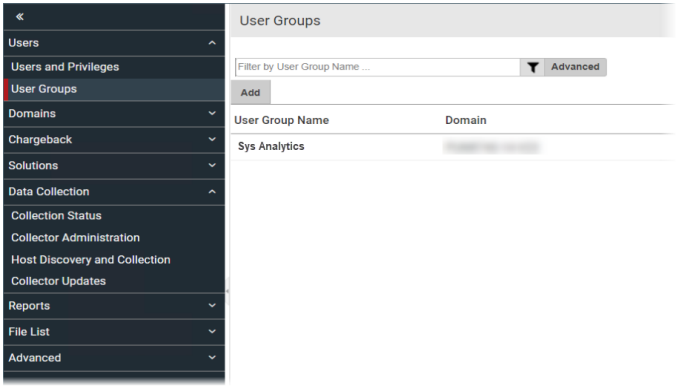
- Enable access to specific reports
- Enable access to functional areas
- Assign a set of home pages to launch at login

Privileges also can be assigned to individual users. A user's privileges are determined from both the user and the group settings. For more information see:

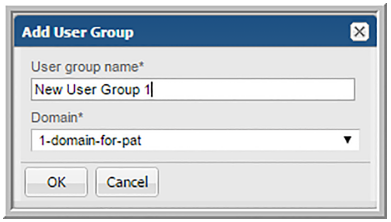
- See "Creating user groups" on page 570.
- See "Assigning users to user groups" on page 571.
- See "Setting user group privileges" on page 581.

Creating user groups

1. Select **Admin > Users > User Groups**.



2. Click **Add**. The **Add User Groups** dialog is displayed.



3. Enter the name of the group.
4. Select the domain in the Domain drop-down list and click **OK**. To learn about domains, go to

---

Note: The user privileges are governed by the product license associated with the domain to which the user group is assigned.

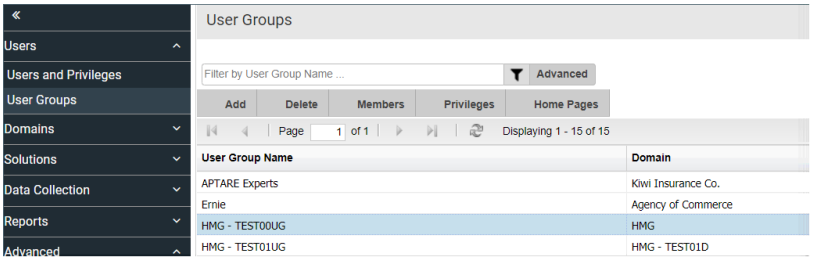
See "Add/Configure a domain" on page 766.

Assigning users to user groups

Both users and user groups must exist to complete the next set of steps.

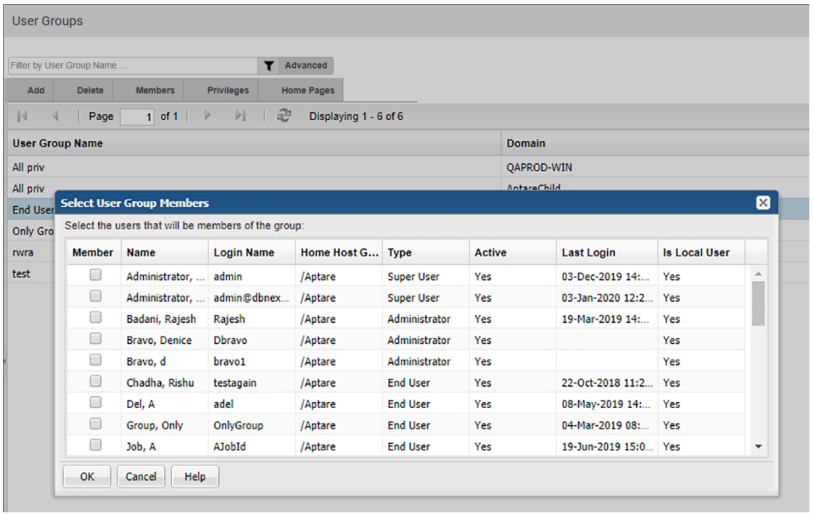
To assign users to user groups

1. Select **Admin > Users > User Groups**. The window displays all user groups.
2. Select the user group to which you want to add a user.



3 Click **Members**. A list of users is displayed. This dialog lists details for each member including if the user was created in the local Portal or if they were added through an Active Directory (AD) system.

4 Select the check box for the user that you want to add to the user group, then click **OK**.



Note: A user with Administrative privileges is only permitted to add a user to a group for which that Administrative user is also a member. Therefore, when you create a group, immediately add the Administrative user to the group to enable management of that group.

About user privileges

Privileges can be set for both user groups and individual users.

Note: The user privileges are governed by the product license associated with the domain to which the user is assigned.

- See “About user types” on page 565.
- See “Assigning user privileges” on page 578.
- See “Managing users and user groups” on page 593.

Privileges are organized by areas in the portal: **Inventory, Reports, Alerts, Admin**. Various categories of privileges enable you to tailor user access capabilities:

Table 14-1 Table 1 User Privileges

Privilege category	Privilege description
Inventory	
	Enables or restricts the ability to access to the **Inventory** tab. **Note:** Reports and templates available through the **Inventory** tab are individually controlled by **Report** specific privileges.
Enable Access to Objects	Enables or restricts the ability to access to Inventory objects.
Arrays	Enables or restricts the access to the Array object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately. **Note:** User access to LUNs can be configured separately irrespective of access granted to the Arrays.
Azure Storage Account	Enables or restricts the access to the Azure Storage Accounts object. This includes the ability to export.
Azure Virtual Machine	Enables or restricts the access to the Azure Virtual Machines object. This includes the ability to export.
Backup Servers	Enables or restricts the access to the Backup Server object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately.
Datastores	Enables or restricts the access to the Datastore object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately.
Dedupe Appliances	Enables or restricts the access to the Dedupe Appliance object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately.
File Shares & Volumes	Select to provide access to File Share & Volumes on the Inventory page. This includes the ability to export.
EC2 Instances	Select to provide access to Amazon Web Services (AWS) EC2 Instances on the Inventory page. This includes the ability to export.

Table 14-1 Table 1 User Privileges *(continued)*

Privilege category	Privilege description
Hosts	Enables or restricts the access to the Host object in the Inventory. This includes the ability to decommission, recommission, export and import objects. Assigning attributes and permanent deletion are permissioned separately.
S3 Buckets	Select to provide access to Amazon Web Services (AWS) S3 Buckets on the Inventory page. This includes the ability to export.
Switches	Enables or restricts the access to the Switch object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately. **Note:** User access to Ports can be configured separately irrespective of access granted to the Switches.
VM Guests	Enables or restricts the access to the VM Guest object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately.
VM Servers	Enables or restricts the access to the VM Server object in the Inventory. This includes the ability to export. Assigning attributes and permanent deletion are permissioned separately.
Manage Objects	Enables or restricts the ability to manage to Inventory objects.
All Objects: Permanently Delete	Enables or restricts the ability to permanently delete objects from the Inventory view and the database.
All Objects: Assign Attribute Values	Enables or restricts the ability to assign attributes and change an object's attribute values within the Inventory. **Note:** These privileges are specific to the **Inventory** tab.
Hosts: Import from CSV	Enables the user to use a CSV file to add/update hosts from an external source and ultimately manage them from the **Inventory**.
Hosts: Add/Edit/Move	Enables or restricts access to add and modify Hosts. This includes moving (cutting and pasting) hosts from one Host Group to another.
Hosts: Recommission and Decommission	Enables or restricts he ability to recommission and decommission hosts and primary servers. These are audited actions and the historical data is preserved in the database. Even with this privilege, a password is required for the action.

Table 14-1 Table 1 User Privileges *(continued)*

Privilege category	Privilege description
Host Groups: Add/Edit/Move	Enables or restricts access to add and modify Host Groups. This includes moving (cutting and pasting) one Host Group to another, allowing you to create sub groups.
Reports	
	Enables or restricts the ability to access to the **Reports** tab. Access to reports within a product/group--for example, System Administration or Management Reports. Access can be on a per-report basis or for an entire report menu group. New reports must be explicitly enabled for users. Access is automatically granted for Super Users and Admins. **Note:** Portal upgrades will automatically enable privileges for newly added reports and the display of the **Inventory** view along with all objects for all Administrators. Refer to the release notes for the list of reports and features introduced in a specific product release.
Enable Access to Cloud Reports	Enables access to new and updated reports delivered through the Cloud. **Note:** The Cloud section is only displayed on the **Reports** tab when this privilege is enabled.
Dynamic Template Designer	Enables or restricts access to the Dynamic Template Designer - a tool that does not require Structured Query Language (SQL) knowledge to create custom report templates.
SQL Template Designer	Enables or restricts access to the SQL Template Designer - a tool to create custom report templates using SQL.
Alerts	
Alert Configuration	
Configure Thresholds & Alerts	Enables or restricts access to Alert Configuration.
Designers	
Admin	
Users	

Table 14-1 Table 1 User Privileges *(continued)*

Privilege category	Privilege description
Users and Privileges	Enables or restricts access to view, search and modify user details. This includes editing group memberships, privileges and password resets.
User Groups	Enables or restricts access to view, search and modify user group details. This includes adding, editing user membership, and privileges. User groups provide an efficient way of managing many users at once.
Domains	
Domains	Enables or restricts access to Domain administration. This includes the ability to add, edit and delete domains.
Chargeback	
Backup	Select to provide access to Backup Billing and Usage Policy administration functions including adding, editing and deleting.
Capacity	Select to provide access to Capacity Billing and Usage Policy administration functions including adding, editing and deleting.
SAN Fabric	Select to provide access to SAN Fabric Billing and Usage Policy administration functions including adding, editing and deleting.
Solutions	
Storage Optimization	Select to provide access to Storage Optimization and corresponding management functions.
Risk Mitigation	Select to provide access to Risk Mitigation and corresponding management functions.
Data Collection	
Collection Status	Select to provide access to monitor the health and status of Data Collectors.
Collector Administration	Select to provide access to create and configure Collectors and Collector policies. Stop, start, edit, and policy deletion is also permitted. On-Demand collection runs are also available for supported vendors.
Host Discovery and Collection	Select to provide the ability to configure Host Resources Data Collection and manage the Host Discovery and Collection.

Table 14-1 Table 1 User Privileges *(continued)*

Privilege category	Privilege description
Collector Updates	Select to enable or restrict the view of the current status for data collectors. This area also provides access to the update the Upgrade Manager and aptare.jars.
Reports	
Thresholds	Select to provide access to add, edit and delete Threshold Policies. Threshold Policies enable you to establish Low, Warning, and Critical levels from which to manage the state of your capacity utilization.
NetBackup Discovery	Select to provide access to the management and configuration of host discovery policies. The Discovery module is specific to NetBackup
Backup SLA	Select to provide access to add, edit and delete Backup Service Level Agreements (SLA) Group polices. When you establish an SLA, performance is tracked against objectives. The SLA policies are where objectives are set.
Primary Schedules	Select to provide access to creating and editing primary schedules. Primary schedules can be defined at a global level and then applied to specific saved reports, causing the reports to be emailed or exported on a regular basis.
Backup Windows	Select to provide access to add, edit and delete custom backup windows.
Method Designer	Enables or restricts the access to the Method Designer, an advanced feature that requires experience in SQL (Structured Query Language) query development. Methods can be used only in report templates that have been created using the Dynamic Template Designer.
Email/Export on Schedule	Select to provide the ability to email or export user reports based on an establish schedule.
Share Reports	Select to provide the ability to share custom reports and templates with users and user groups.
Create Ticket	Select to provide the ability to create tickets associated with backup jobs.
File List	
File List Export	Enables or restricts access to the File List Exporter - a utility for extracting the File Analytics collected metadata into a comma-separated values (.csv) file.

Table 14-1 Table 1 User Privileges *(continued)*

Privilege category	Privilege description
Advanced	
Parameters	Select to provide the ability to add, edit, delete and download advanced parameters. Advanced parameters are a mechanism for customizing internal settings with the help of Support.
Attributes	Select to provide the ability to create, edit and delete attributes and modify their values. Attributes define a distinct data set based on a specific characteristic.
Authorization Attributes	Select to provide the ability to: ■ Create an authorization attribute ■ Assign authorization attributes to a user ■ Assign authorization attributes to data objects
Object Maintenance	Select to provide the ability to manage Devices, Libraries and Drives. Permanently delete and the assignment of attribute values are permitted operations.
Support Tools	Select to provide the ability to download data collector logs, portal logs and configuration files for a specified time period.
System Notifications	Select to provide the ability to view and manage system level notifications. Notifications alert the user to any issues that require attention.

Portal upgrades automatically enable privileges for newly added reports and the display of the **Inventory** view including all objects, for all Administrators. Refer to the *IT AnalyticsRelease Notes* for the list of reports and features introduced in a specific product release.

Assigning user privileges

You can assign users privileges, such as enabling access to specific reports and authorizing administrative configuration features.

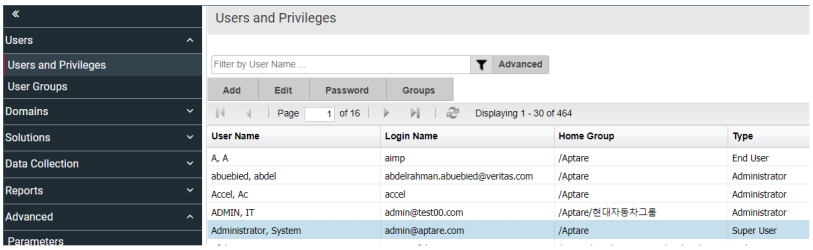
- See “About user privileges” on page 572.
- See “Setting user group privileges” on page 581.
- See “Enabling new product report templates” on page 586.
- See “Enabling cloud privileges” on page 590.

- See “Granting access to template designers” on page 591.
- See “Enabling all privileges in a category” on page 592.

To assign privileges to an individual user

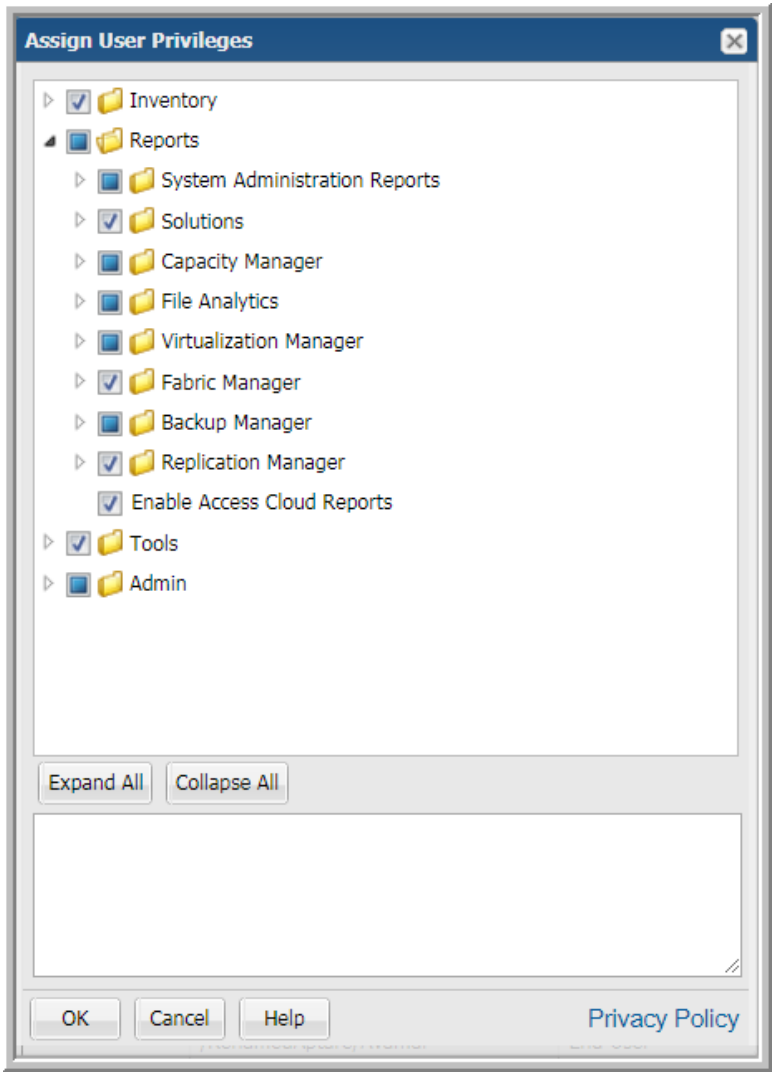
1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Search for a user if required.

   See “Searching for users and user groups” on page 593.



3. Select the user name to which you want to assign privileges.

4 Click **Privileges**. The **Assign User Privileges** window is displayed.



5 Click the arrow beside the appropriate privilege folder to expand it. Or, click the **Expand All** button at the bottom of the window.

See “About user privileges” on page 572.

6 Individually select the check boxes that correspond to the required privileges or select all entries in a folder by following the procedure described in

See “Enabling all privileges in a category” on page 592.

A blue checkbox indicates some privileges are selected within the category.

7 Click **OK**.

Setting user group privileges

Once you set up a User Group, you can configure the access privileges for the members of that group. For a list of privileges that can be configured refer to the following.

See “About user privileges” on page 572.

Each category folder can be expanded to select a subset of the access category. Expand the folder and select the privileges to be assigned to the selected User Group. A blue checkbox indicates some privileges are selected within the category.

To select an entire category, such as all reports, follow the procedure described in the following section.

See “Enabling all privileges in a category” on page 592.

To assign privileges to a User Group

1 Select **Admin > Users > User Groups**

2 Search for a user group if required.

See “Searching for users and user groups” on page 593.

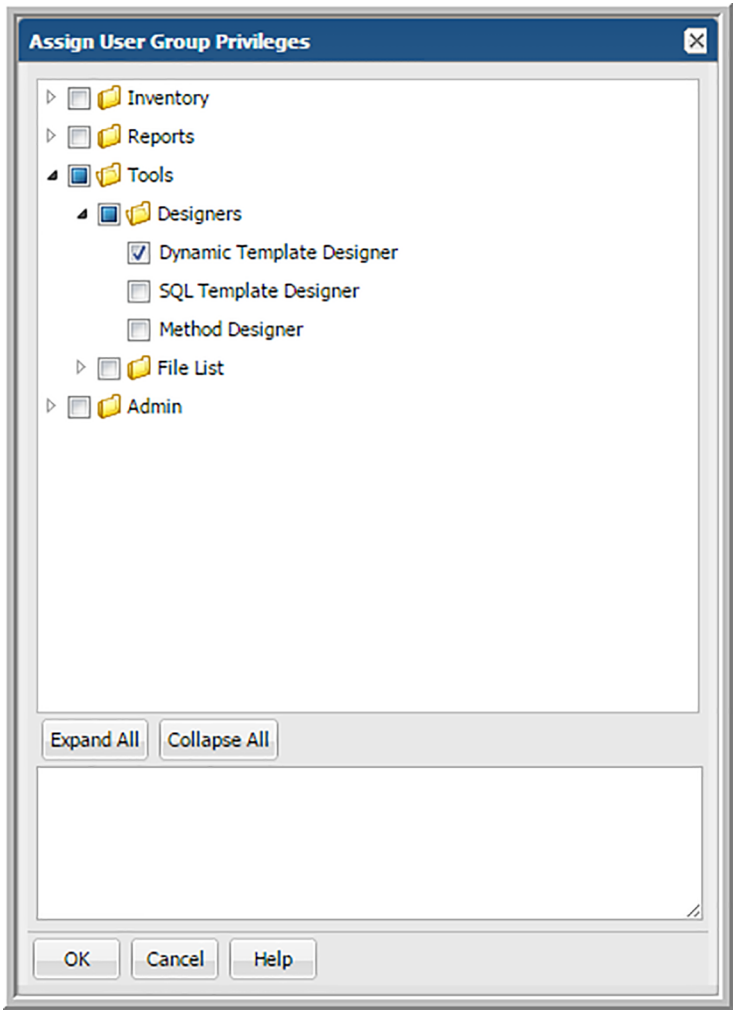
3 Select a **User Group Name**.

4 Click **Privileges**. Click the arrow beside the appropriate privilege folder to expand it. Or, click the **Expand All** button at the bottom of the window.

See “About user privileges” on page 572.

5 Individually select the check boxes that correspond to the required privileges or select all entries in a folder by following the procedure described in the following section.

See "Enabling all privileges in a category" on page 592.



6 Click **OK**.

Enabling new product report templates

When new system report templates are introduced, existing user accounts need to be modified to gain access to the new templates. For Super Admin, notification badges are displayed in the **Reports** tab. These badges quickly identify new and updated system report templates. Badges are automatically displayed on their product folder indicating when a template change has been introduced. This badge serves as a notification. The Super Admin must use privileges to enable each product report template for users or user groups.

See "About badging" on page 102.

When they are introduced, new system report templates, delivered automatically by IT Analytics appear in the correct folders in the privileges section. All reports are displayed, regardless of the products you have licensed and installed.

Note: Portal upgrades will automatically enable privileges for newly added reports and certain features/functions, for all Administrators. Refer to the release notes for the list of reports and features introduced in a specific product release.

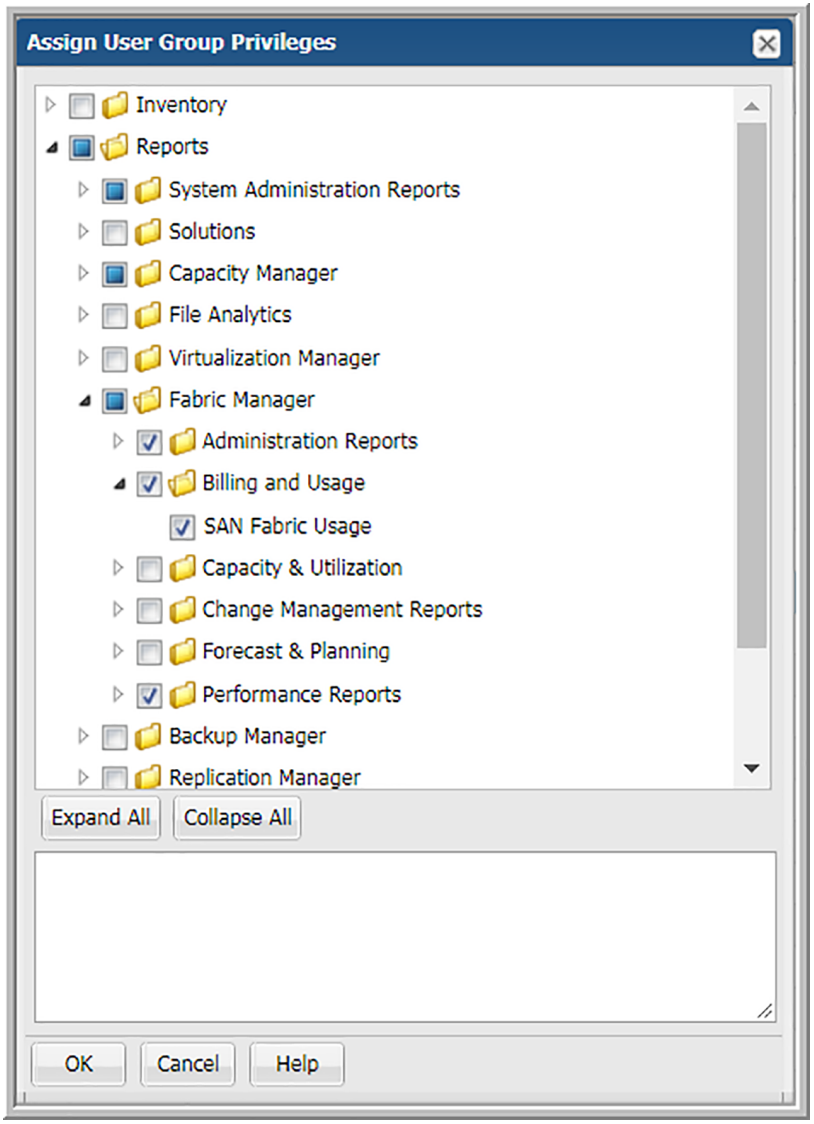
RECOMMENDATION: The User Group feature lends itself to quickly enabling access to report templates that become available.

To grant access to report templates

1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Search for a user if required.

   See "Searching for users and user groups" on page 593.
3. Select a user name and click **Privileges**.

4 Expand **Reports**.



5 Click the reports to grant access. Select the folder to grant access to an area of reporting. Expand the folder to select individual reports.

This can be done by individual user or by user group. User groups are the most efficient way to grant access to any report template, feature or enhancement.

Importing reports

IT Analytics allows an administrator user to assign **Import Reports** privileges to the users on cloud portal.

Note: By default, all the users are assigned with the *Import Reports* privileges.

---

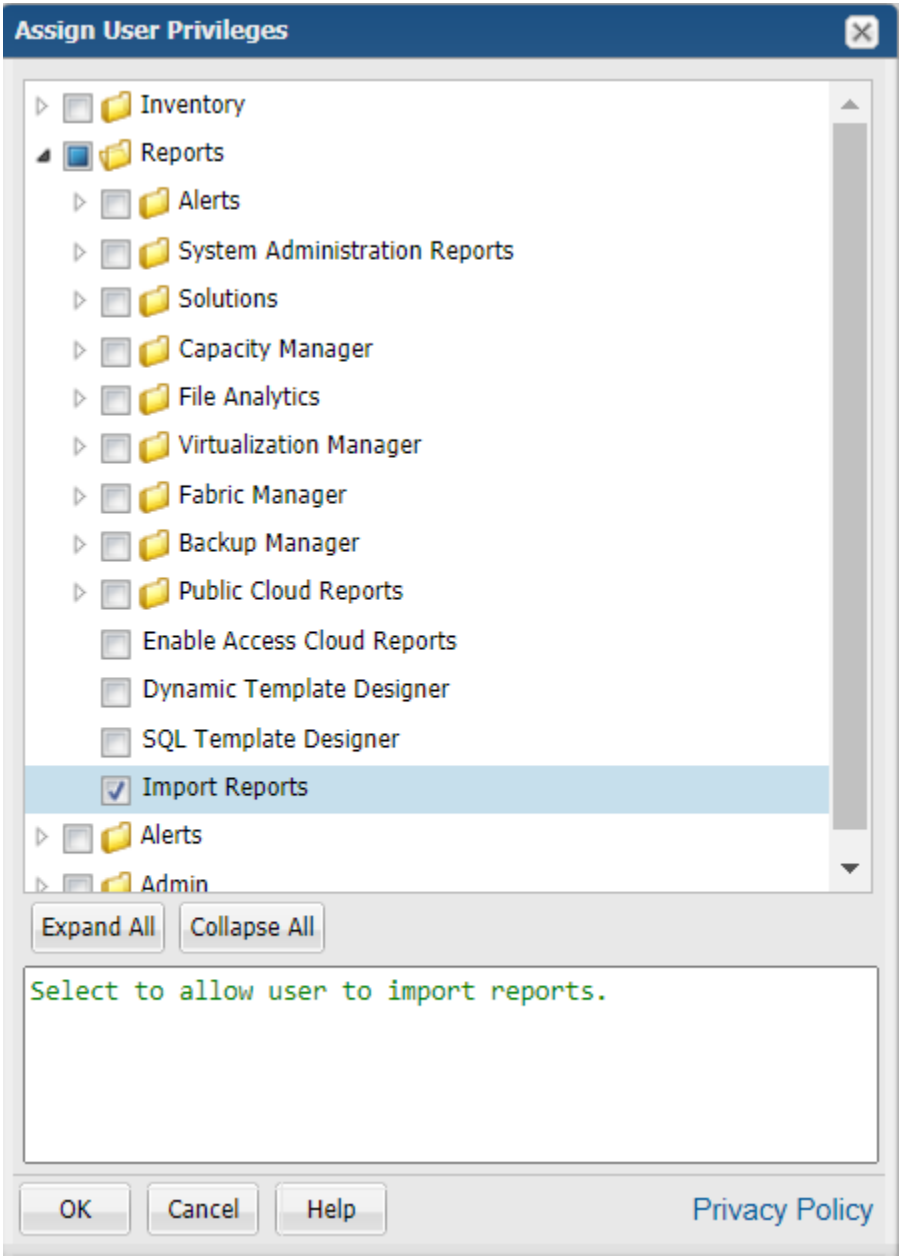
1. Navigate to **Admin > Users > Users and Privileges**. *User and Privileges* page is displayed.

   Note: Search for a user in **Filter by User Name** field, if required.

   Note: See “Searching for users and user groups” on page 593.

   ---

2. Select the user and then click **Privileges**. **Assign User Privileges** dialog box is displayed.
3. Click **Reports** toggle button to expand reports folder.
4. Click **Imports Reports**.


Assign User Privileges
Inventory
Reports
Alerts
System Administration Reports
Solutions
Capacity Manager
File Analytics
Virtualization Manager
Fabric Manager
Backup Manager
Public Cloud Reports
Enable Access Cloud Reports
Dynamic Template Designer
SQL Template Designer
Import Reports
Alerts
Admin
Expand All
Collapse All
Select to allow user to import reports.
OK
Cancel
Help
Privacy Policy

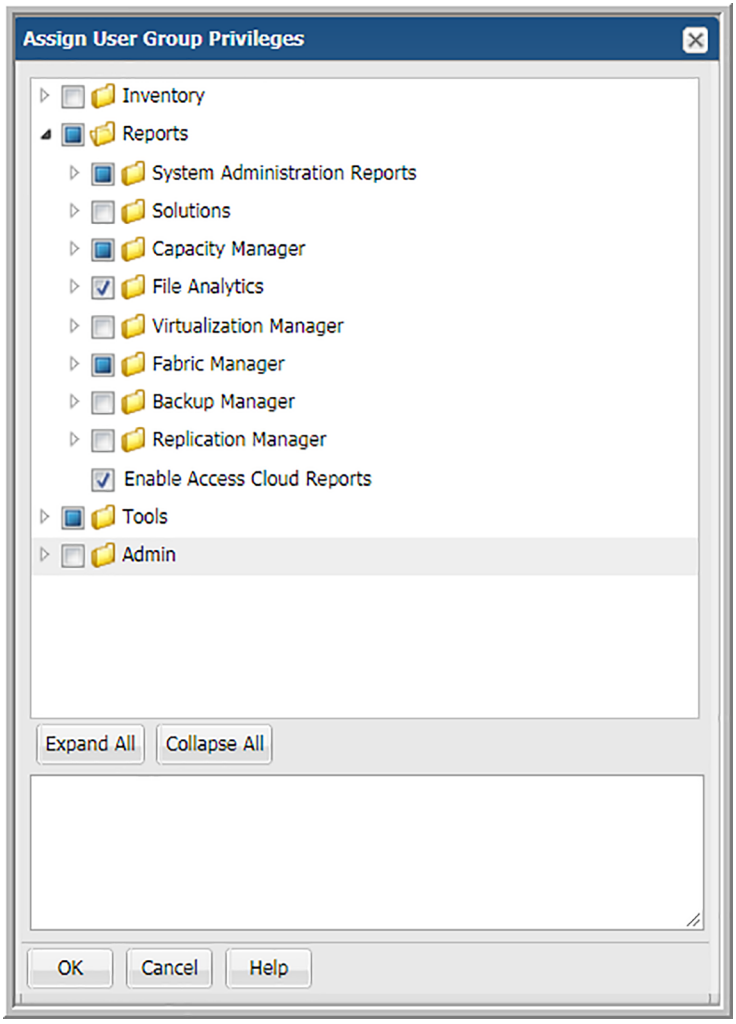
5. Click **OK**.

Enabling cloud privileges

IT Analytics periodically publishes new user report templates and makes them available through the Cloud section on the **Reports** tab. You must enable the privilege for users to see the Cloud section on the **Reports** tab, view and have access to the report templates.

1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Search for a user if required.

   See “Searching for users and user groups” on page 593.
3. Select a user name and click **Privileges**.
4. Select **Reports> Enable Access Cloud Reports**.



5. Click **OK**.

Granting access to template designers

Several tools give you the capability to create report templates to satisfy your organization's reporting requirements.

- Dynamic Template Designer: A tool to create custom reports by dragging and dropping database components, which are then used to dynamically generate the Structured Query Language (SQL) query "behind the scenes" of the Portal.

See “Dynamic Template Designer Overview” on page 220.

- SQL Template Designer: An interface to enable custom report creation using SQL skills.
- Method Designer: A tool to create a method, which enables special processing to be incorporated into a report template. A method can take parameters, perform an action, such as a complex calculation, and then return a value. This advanced feature requires experience in SQL query development. Methods can be used only in report templates that have been created using the Dynamic Template Designer.

1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Search for a user if required.

   See “Searching for users and user groups” on page 593.
3. Select a user name.
4. Click **Privileges**.
5. Select **Reports** and choose your template designer.
6. Click **OK**.

Enabling all privileges in a category

For User and User Group administration, you can configure privileges for individual functions, or you can configure access for an entire category of privileges.

Note: The following procedure assumes that you already have created the user or user group. This procedure also can be used when creating a new user or group.

To enable all privileges within a category

1 Select **Admin > Users > Users and Privileges**. The window displays all Portal users.

2 Search for a user if required.

See “Searching for users and user groups” on page 593.

3 Select a user name.

4 Click **Privileges**.

5 Click the category folder--**Reports, Tools**, or **Admin**--to select it. You can expand folders and then select sub-folders to enable all functions within a particular sub-folder. If the checkbox beside a folder is blue, this indicates some, but not all privileges have been selected.

6 Click **OK**.

Managing users and user groups

- See “Searching for users and user groups” on page 593.
- See “Editing user accounts” on page 593.
- See “Impersonating user accounts” on page 595.
- See “Setting / Resetting passwords” on page 597.
- See “Account lockout” on page 598.
- See “Managing user group home pages (Administrator)” on page 599.
- See “Removing portal users” on page 604.
- See “Deactivating user accounts” on page 604.

Searching for users and user groups

Use Search to locate Users and User Groups.

See “Navigating with search” on page 35.

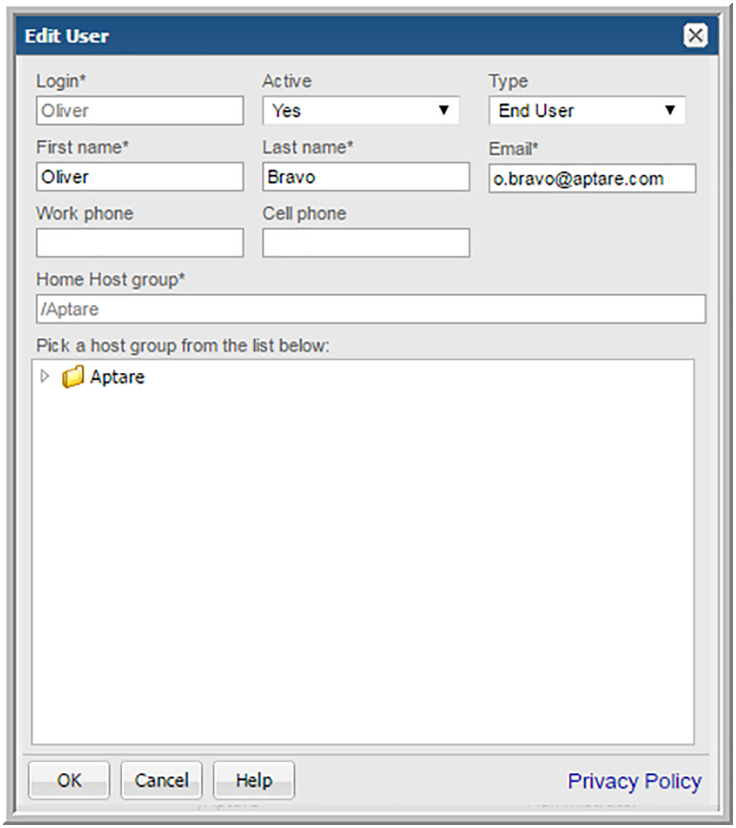
Editing user accounts

When you create a user account, you create the user details, access privileges, group membership and a password. Users can modify their passwords and most of their profile information, however, they cannot modify their access privileges. Preferences for locale, number and date formatting are set by the user using the user account menu.

See “Manage your profile and set a language preference” on page 772.

1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Search for a user if required.

   See “Searching for users and user groups” on page 593.
3. Select a user and click **Edit**. The **Edit User** dialog is displayed.



4. Revise the fields.
   - Required fields are denoted with an asterisk. First name and last name fields are limited to 64 characters.
   - The Login ID should be in an email format, as required by LDAP. This field is limited to 128 characters.
   - From the selector, select the User Type. For a definition of the privileges associated with a user type,
     See "About user types" on page 565.
   - Select a Home Host group. A home host group is a host group to which a user belongs. This limits a user's access to data. A user can access any host groups that are lower in the hierarchy from the home host group.
5. Click **OK** to edit the user.

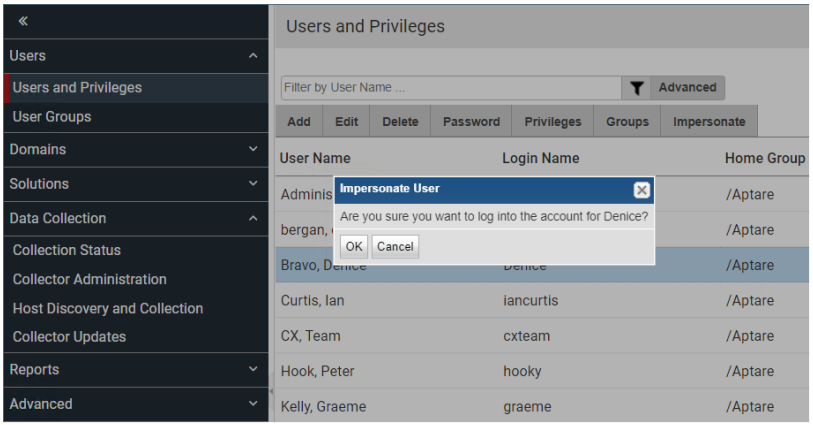
Impersonating user accounts

A user, assigned the Super User role can impersonate another user. The Impersonate feature enables a Super User to easily log into a separate user session to for example, diagnose issues, verify permissions, and manage a user's scheduled reports. This is especially useful in an MSP environment. This information is captured in an audit.log file. The Impersonate feature is also useful when employees leave an organization and access is required to export custom reports from an account.

To impersonate a user account

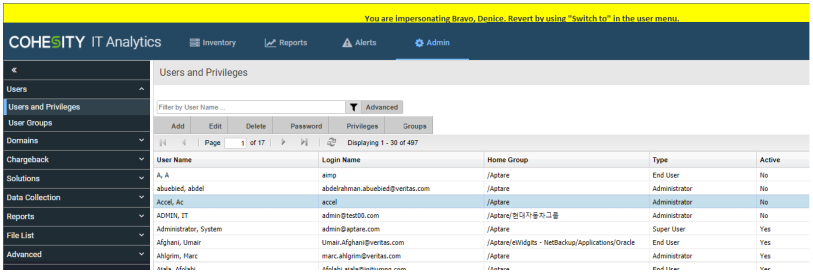
1. Navigate to **Admin > Users > Users and Privileges**.
2. Select a **User Name**.

3 Click **Impersonate**. The **Impersonate User** message prompts you for a confirmation.



4 Click **OK**. You are redirected to the last open session of the selected user. The privileges associated with the user account are retained, so you are logged in as that user.

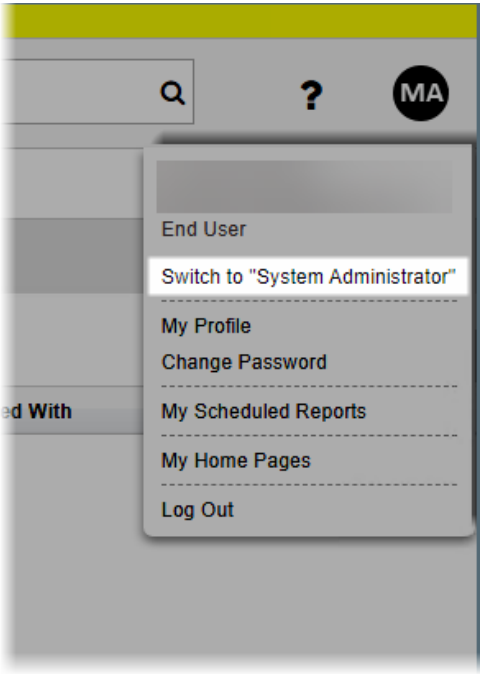
A yellow ribbon is displayed at the top indicating the user you are impersonating.



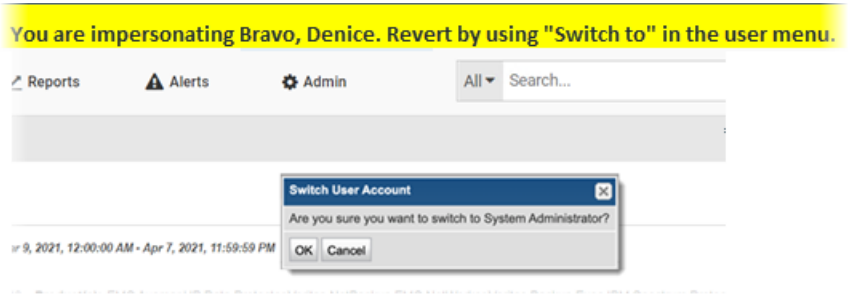
To switch back to the super user account

Once logged into the new session, you are logged out of your own session. You must explicitly switch back into your own user account.

1 Click the user menu and click **Switch to System Administrator**.



2 Click **OK** on the confirmation message to switch back to the super user account.



Setting / Resetting passwords

Once you create a user, you need to assign an initial default password, for example, employee ID. The user can change this password from within the Portal.

1. Select **Admin > Users > Users and Privileges**. The window displays all Portal users.
2. Search for a user if required.

   See “Searching for users and user groups” on page 593.
3. Select the user.
4. Click **Password** and enter the password.

Password notes

Passwords can include the equal sign (=) character.

Passwords are stored in the database using the following encryption algorithm:

- Algorithm: AES/CBC/PKCS5Padding
- Strength: 128-bit

Account lockout

A user is allowed five attempts to login. If the user is unable to remember the password after five attempts, the account is locked. This value can be overridden by a Super User by adjusting the setting under **Advanced<System Configuration< Portal**.

An administrator must reactivate the account before the user can login.

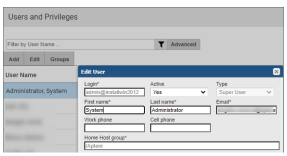
To unlock the account

1 Click **Admin>Users>Users and Privileges**.

2 Search for a user if required.

See “Searching for users and user groups” on page 593.

3 Select the user name and click **Edit**.



4 Change the **Active** field from **No** to **Yes**.

5 Click **OK**.

Managing user group home pages (Administrator)

Home pages are associated with user account and display when a user logs into the Portal. Each user can assign a different report (or set of reports) as their home pages and they are launched automatically each time they log in.

This functionality is also available for User Groups. System Administrators can designate a set of reports as home pages and assign them to User Groups. Each member of the User Group receives the designated home page set and they automatically launch each time the user logs in. These assigned reports are displayed on the user's My Home Page Administration dialog along with any Home Page reports selected by the user.

When users belong to multiple User Groups with Home Pages assigned, all reports are combined, sorted and added to the member's list of Home Pages.

My Home pages and user group home pages (Administrator)

When Home Pages are assigned through User Groups:

- Members can reorder the display order.
- Members cannot delete the reports. Members can always delete the reports they individually selected as Home Pages.
- Reports are automatically added for each User Group member when updated by an Administrator.
- Reports are automatically removed for each User Group member when deleted by an Administrator.
- Reports are removed for each User Group member when the User Group is deleted.
- Reports are automatically added to the top of the user's Home Page list.
- New reports are flagged with an asterisk.
- Total number of reports automatically launched is configurable by an Administrator.
  See "Manage My Home Pages" on page 211.

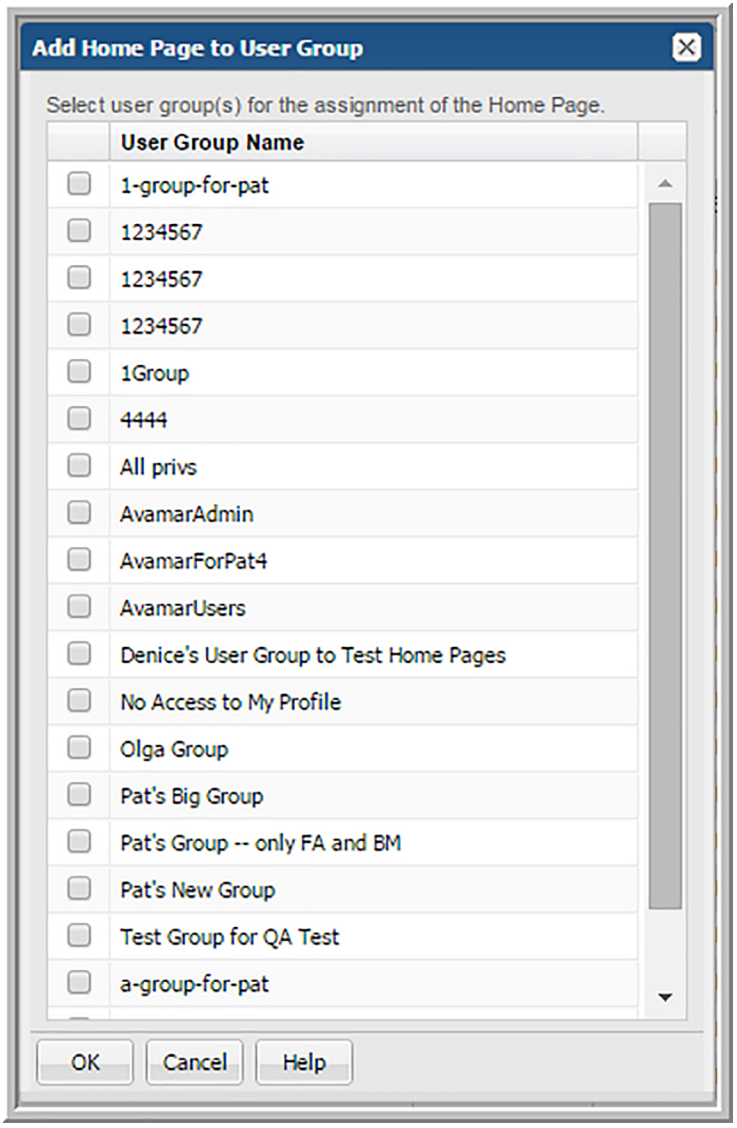
To designate reports or dashboards as User Group home page(s)

You can designate one or multiple reports as home pages for a User Group. You can add reports as home pages, however, only the first five reports listed on the logged in user's **My Home Page Administration** dialog will automatically launch at login.

When reports are added using User Groups, the first five are added to the top of the display list. However, users manage the ultimate home page display order regardless of the assignment origin: user-selected or assigned through User Groups.

1 Generate a saved report or dashboard.

2 Right-click and select **Add to User Group Home Pages**. The **Add Home Page to User Group** dialog is displayed.



3 Select the User Group(s) for the assignment.

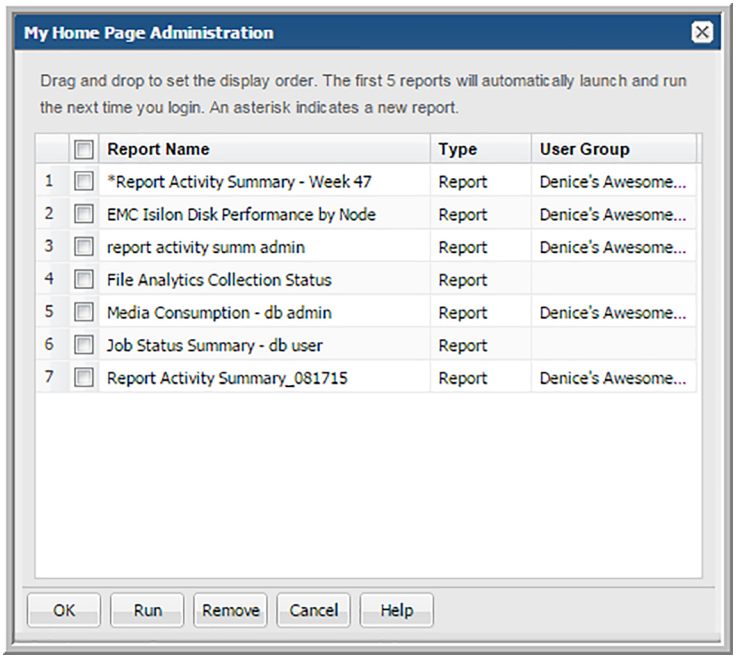
4 Click **OK**.

To add new reports to the home page set

- New Home Pages added to a User Group with an existing list are automatically display at the top of the user's Home Page list. By default, five reports can be added to the top of the list. This value is configurable. Refer to the following.

  See "Manage My Home Pages" on page 211.

  These new Home Pages display an asterisk (*) beside them on the user's My Home Page Administration dialog to indicate they've been added through the User Group. Once the report is run, the asterisk is removed.



To remove a report or dashboard from User Group home pages

If all Home Pages are removed, the Inventory view is displayed (if the privilege is selected for a user). If privileges are not set, the Reports view is displayed. User Group assigned home pages are automatically removed from the individual users list when the Administrator removes them from the group.

Note: Individual users cannot remove home pages assigned through the User Groups. They must be removed from the User Group.

1 Select **Admin > Users > User Groups**.

2 Select the **User Group Name**.

3 Click **Home Pages**. The **User Group Home Page Administration** dialog is displayed.

4 Select the reports (Home Pages) to remove.

5 Click **Remove**. The reports are automatically removed from the User Group's members' **My Home Page Administration** dialog. The report will be removed from the member's list the next time they log in.

Removing portal users

Delete a user's account if that user leaves the company, or if that user no longer needs access.

To remove a Portal user

1 Select **Admin > Users > Users and Privileges**. The window displays all Portal users.

2 Search for a user if required.

See "Searching for users and user groups" on page 593.

3 Click the user name to select the user account to delete.

4 Click **Delete**, then confirm your changes.

Deactivating user accounts

Consider deactivating a user's account if that user does not intend to use the Portal for an extended period of time, such as in the event of an employee's leave of absence. If you want to permanently suspend the user account, refer to the following.

See "Removing portal users" on page 604.

To deactivate a user account

1 Select **Admin > Users > Users and Privileges**. The window displays all Portal users.

2 Search for a user if required.

See "Searching for users and user groups" on page 593.

3 Select the user to deactivate, and click **Edit**. The user's profile is displayed on the **Details** tab.

4 Select **No** from the **Active** drop-down list, and click **OK**.

Chapter 15

Configure primary schedules and backup windows

This chapter includes the following topics:

- About primary schedules
- Configure primary schedules
- About custom backup windows
- Custom backup window requirements
- Configure custom backup windows

About primary schedules

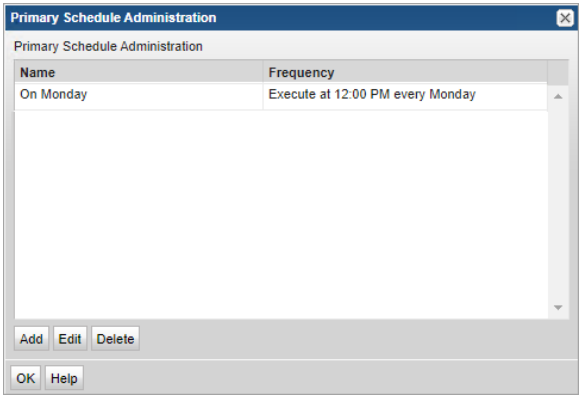
Primary schedules can be defined at a global level and then applied to specific saved reports, causing the reports to be emailed or exported on a regular basis. The advantage of a primary schedule is that any changes to a schedule automatically affect its associated reports. So, for example, a particular department may need a set of monthly reports for capacity planning. Once these reports have been configured to run according to a primary schedule, they easily can be generated on a more frequent schedule, such as once a week. Rather than changing each report, you need only change a single primary schedule configuration.

See “Configure primary schedules” on page 606.

Configure primary schedules

To add/edit primary schedules for report emails and exports:

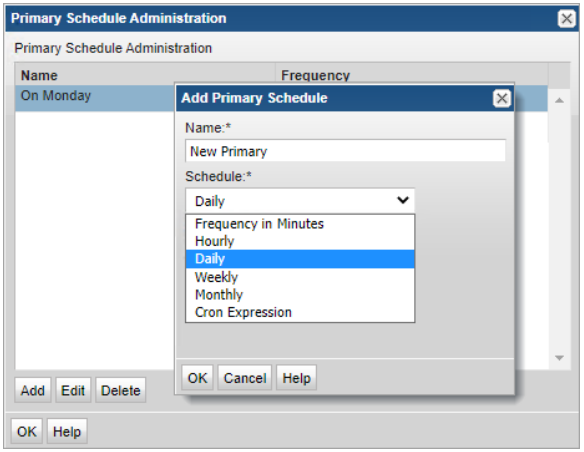
1 Select **Admin > Reports > Primary Schedules**.



2 Click **Add/Edit** in the **Primary Schedule Administration** window.

For a description of how Primary Schedules can be implemented, refer to the following.

See “About primary schedules” on page 605.



3 Create a name and select the scheduling elements: Frequency in Minutes, Hourly, Daily, Weekly, Monthly or a Cron Expression.

4 Select a time.

5 Click **OK**.

About custom backup windows

Often an enterprise measures success by Service Level Agreements (SLA) that have been defined for backups. Typically, backups begin at the end of the business day, but they do not finish before the end of the day--thereby skewing the success statistics for the day. To more accurately reflect backup SLA metrics, you can re-define a day, for example, to mean 4:00 p.m. today until 4:00 p.m. tomorrow.

Custom Backup Windows can be defined at a global level. A Backup Window can then be applied to the following reports to arrange the data based on the selected custom backup window:

- Job Status Summary
- Mission Control
- Data Protection Dashboard

Note: Several automated reports and custom report templates have Backup Windows already defined. The customization of Backup Windows, as discussed herein, only applies to the above list of reports.

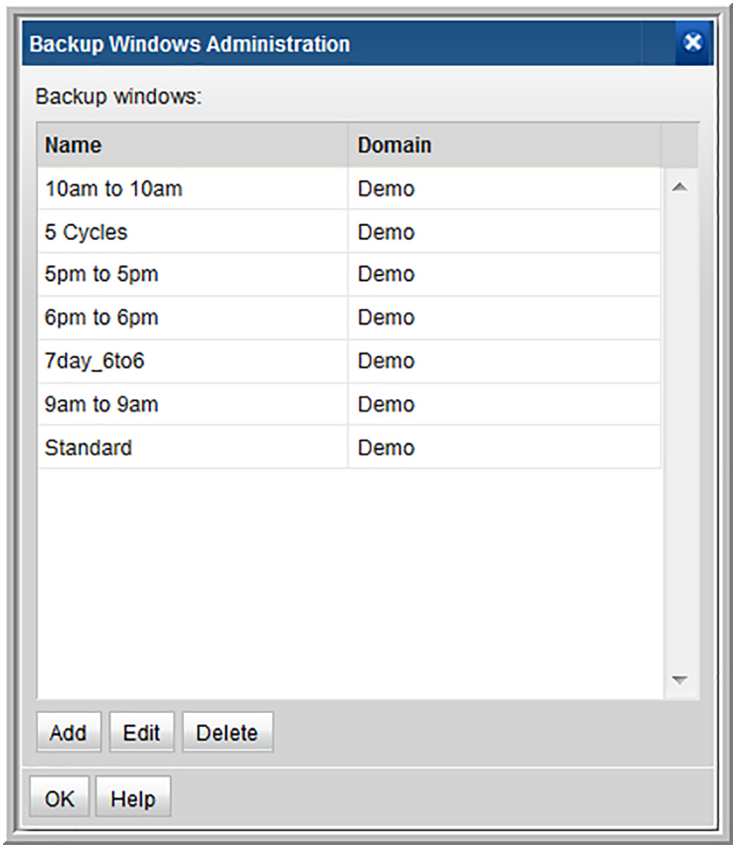
Custom backup window requirements

- The total of all the hours defined in a custom backup window must equal 168--a full week.
- A week is defined as starting at 00:00 Sunday, unless you redefine it with a Custom Backup Window.

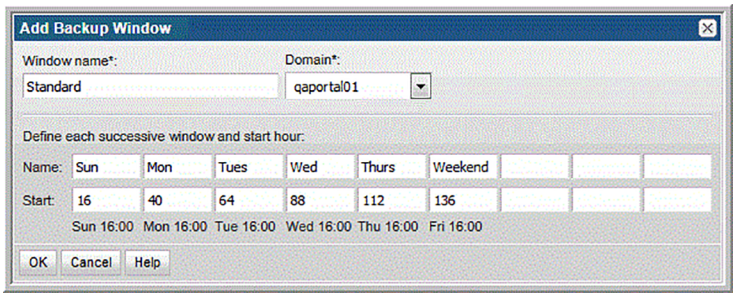
Configure custom backup windows

To create/edit a custom backup window

1 Select **Admin > Reports > Backup Windows**.



2 Click **Add/Edit**. To delete a Backup Window, click a window name to select it, and then click **Delete**.



3 Enter a **Window name** and select a **Domain** from the drop-down list.

4 Specify each day or segment of the backup window by entering a Name for the day/segment and also the Start Hour. The Hour must be specified as an integer. The ending hour of each segment is computed automatically when the next start hour is entered. See the example screen in step 2 and use the following guidelines to define the window.

- When you create a custom backup window, you must account for a full week--that is, no more than 168 hours without any gaps or overlaps. Typically, you'll define each day in the backup window and perhaps create a weekend time frame to account for the hours in which no backup jobs run.
- 0 = Saturday at midnight (1 second after Saturday 23:59:59)
- The week does not have to begin with 0 and the week cannot exceed 168 hours.

Chapter 16

Add, edit, and move policies

This chapter includes the following topics:

- Add/Edit a threshold policy
- Add/Edit a capacity billing and usage policy
- Capacity Chargeback policy types
- Add/Edit a backup SLA policy
- Add/Edit a SAN Fabric Chargeback policy
- IT Analytics Billing and Chargeback policies
- Add/Edit a Billing and Chargeback policy

Add/Edit a threshold policy

Note: For Backup Billing and Chargeback Policies,

See “Add/Edit a Billing and Chargeback policy” on page 623.

Threshold Policies enable you to establish Low, Warning, and Critical levels from which to manage the state of your capacity utilization. Various reports will represent the threshold crossings with highlighted fields to help you identify issues.

See “Threshold Example” on page 611.

You can set a variety of thresholds:

- File System
- Host Groups

- Host
- LUN
- Database
- NetApp Aggregate

To add/edit a threshold

1. Select **Admin > Reports > Threshold**
2. Click **Add** to add a new policy or select an existing policy and click **Edit**.
3. When you click **Add**, the **Add Threshold Policy** window launches.
4. Select a type from the **Threshold type** drop-down.
5. Select an instance, displayed in the window with the **Parent** column header.

   Use the search feature to find a Parent instance.
6. Specify a **Low**, **Warning**, and **Critical** threshold and click **OK**.

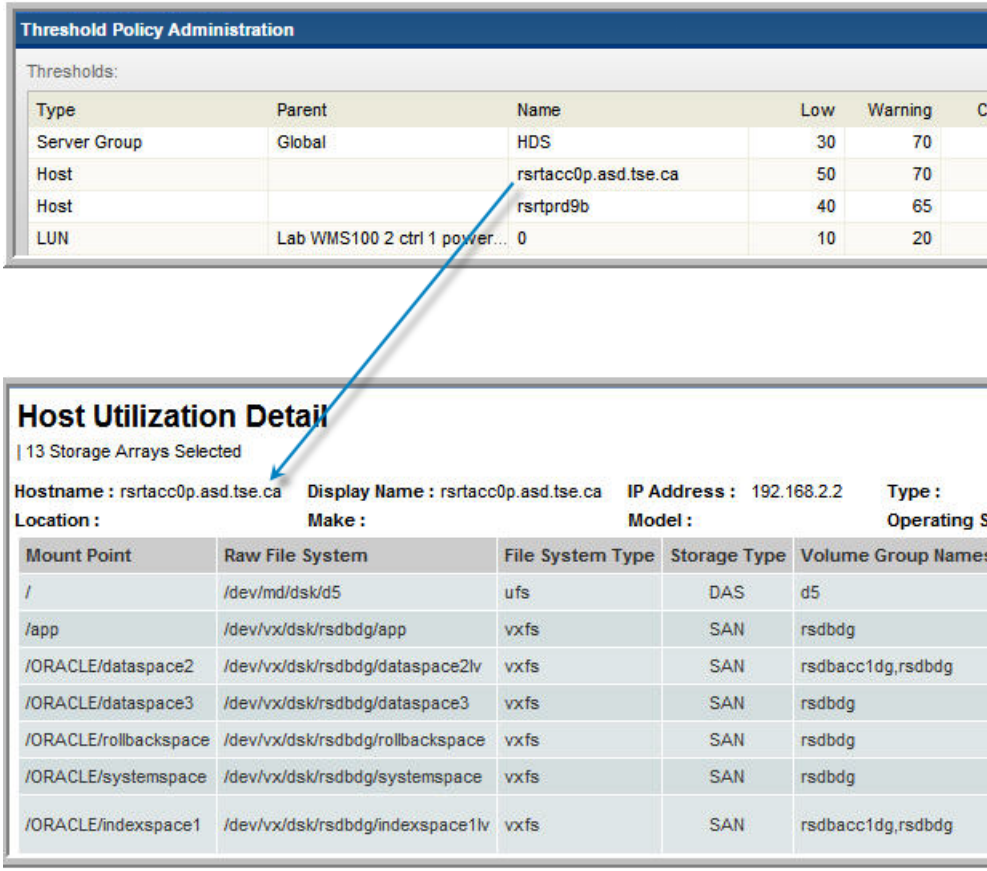
   Note: All three threshold levels are required.

   As a guideline, the low threshold is used to indicate over-provisioned capacity, typically anything that will not be filled up more than about 30%. The Warning and Critical thresholds represent the point at which you are at risk of running out of space, typically 70% and 90% respectively.

Threshold Example

The threshold value represents the percentage used. So, if you set a Critical threshold of 80, that means that you want reports to indicate when there is 20% or less space available.

The following example shows how configured thresholds affect the values shown in reports.



Add/Edit a capacity billing and usage policy

Capacity Billing and Usage Policies for chargeback can be configured to allocate costs for storage array usage. Costs can be associated with a variety of storage elements to enable a detailed approach to usage costs. The values that you configure in the Capacity Billing and Usage Policies for chargeback are used by the following reports:

- Aggregated Chargeback
- Chargeback Array Capacity
- Chargeback By Host
- Chargeback Policy Capacity

Note: Capacity Chargebacks can be configured for block storage only; file-based storage is not supported for Array Capacity Chargeback.

To add/edit a capacity chargeback policy

1. Select **Admin > Chargeback > Capacity**
2. Click **Add** to add a new policy or select an existing policy and click **Edit**.
3. Configure the following elements:

   Note: An asterisk (*) denotes a required entry.

Domain	Select a domain. Typically, only one domain is listed, unless the portal is maintained by a Managed Services Provider (MSP).
Name	Enter a name for the policy.
Priority	■ Priority - 1 is the highest priority ■ If a LUN subscribes to two chargeback policies with the same priority, two LUNs will be listed in the capacity chargeback reports.
Cost/GB	Enter the cost associated with this capacity billing and usage policy.

4. Click **Add** to select a Policy Type.
5. Continue the configuration process, using the details provided in the following.

   See “Capacity Chargeback policy types” on page 613.

Capacity Chargeback policy types

One or more policy types comprise a Capacity Chargeback Policy, as described in this section. Policy Types can be combined to provide “or” logic. If two or more of the same policy types are configured in a policy, the conditions will be “ORed” when the Portal evaluates the policy during report generation. Conditions within the same policy type are joined by OR, while AND logic is used between policy types.

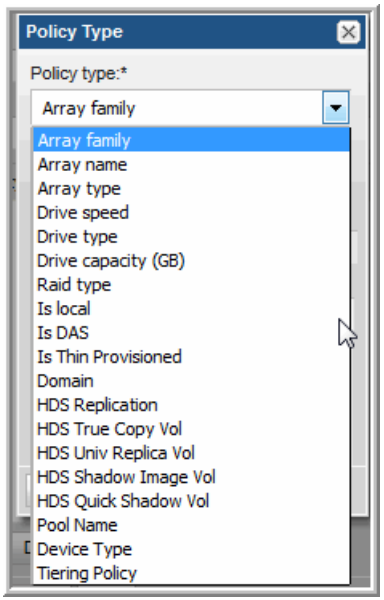
See “Example of Capacity Billing and Usage Policy Logic” on page 617.

Note: Capacity Chargebacks can be configured for block storage only; file-based storage is not supported for Array Capacity Chargeback.

The following procedure is the second part of configuring a Capacity Chargeback Policy.

See "Add/Edit a capacity billing and usage policy" on page 612.

1. Select **Admin > Chargeback > Capacity**
2. Click **Add**.
3. In the **Policy Type** window, select a **Policy type** and one or more **Policy values** from the drop-down list.



- Policy Types include: Array Family, Array Name, Array Type, Drive Speed, Drive Type, Drive Capacity (GB), RAID Type, Is Local, Is DAS, Is Thin Provisioned, Domain, HDS Replication, HDS True Copy Vol, HDS Univ Replica Vol, HDS Shadow Image Vol, HDS Quick Shadow Vol, Pool Name, Device Type, and Tiering Policy.
  See "Drive Types for Capacity Chargeback" on page 618.
- For EMC Arrays, Device Type should be used instead of RAID Type for LUN mapping.
- Capacity Chargeback Policies can be configured for thin-provisioned storage pools for: HP USP and USP-V Dynamic Pools, CLARiiON and Symmetrix Thin Pools, and IBM XIV. To configure this in a Capacity Chargeback Policy, select **Pool Name**.
- Capacity chargeback policies support virtualized storage. Supported virtualized storage systems include Hitachi Virtual Storage Platform (VSP),

Hitachi NAS, EMC VPLEX, IBM SVC, NetApp 7-Mode, and NetApp Cluster Mode. The chargeback policy (drive speed, drive type, drive capacity, and RAID type) uses characteristics of the back-end arrays' LUNs, where the physical disks for the LUNs actually reside, to automatically categorize the LUN into the correct storage tier.

- These HDS Capacity Chargeback Policies--HDS True Copy Vol, HDS Univ Replica Vol, HDS Shadow Image Vol, HDS Quick Shadow Vol--augment the HDS Replication policy, enabling flexibility for policies to cover various combinations of configurations.

Note: Policy types are applicable to specific arrays and therefore, not all policy types will be listed in your Portal. For example, the Tiering Policy type is available only for the following arrays: Hitachi NAS (HNAS), EMC Symmetrix Fast, and Compellent.

---

4. Click **OK** to save the configuration.

   The values that appear in the Policy Value drop-down list are derived from your existing database. However, in anticipation of new storage resources for which data will be collected, you can enter values.

Array Family	Select an Array Family from your environment, or enter a name.
Array Name	Select an Array Name from your environment, or enter a name.
Array Type	Select an Array Type from your environment, or enter a name.
Device Type	Select a type; for example, 2-Way BC Mir (Meta Head, Non-Exclusive Access) or RAID-5 (non-Exclusive Access).
Domain	Capacity Billing and Usage Policies are domain-specific. Therefore, this policy type would rarely be used. The intent of this policy type is to provide the capability for a "catch-all" rule to cover all objects in a domain.

Drive Capacity (GB)	Select a Drive Capacity for drives in your environment, or enter a size. Drive Capacity is treated as a range of values--that is, the value entered here plus or minus the value that is configured as a system parameter. The default value set in the System Configuration is 10 GB. **Advanced Configuration**: To modify the drive capacity range, enter a value here and in **System Configuration (Admin>Advanced>System Configuration Data Collection** tab), provide a value that indicates a range; for example, 320 plus or minus 50. Search the online documentation for System Configuration from the Portal.
Drive Speed	Select a Drive Speed for drives in your environment, or enter a speed.
Drive Type	Select a Drive Type from your environment, or enter a type. See "Drive Types for Capacity Chargeback" on page 618. The display name is associated with an internal policy value. Several policy display names translate to the same internal value. For example, 5-SSD and 8-FMD both translate to a value of SSD. In this example, if you attempt to add both policy display names to the chargeback policy, only one will be retained since they are duplicates of the same value.
HDS Quick Shadow Vol	P-VOL, POOL, Simplex, V-VOL
HDS Replication	JNL-VOL, MF-JNL, P-VOL, POOL, S-VOL, SP-VOL, Simplex, Unknown, V-VOL.
HDS Shadow Image Vol	P-VOL, S-VOL, Simplex
HDS True Copy Vol	P-VOL, S-VOL, Simplex
HDS Univ Replica Vol	JNL-VOL, MF-JNL, P-VOL, S-VOL, Simplex
Is DAS	Direct-attached Storage: Yes or No

Is HDT	Is Hitachi Dynamic Tiering enabled: Yes or No. When Yes, this associates LUNs to HDT pools, to enable distinct and separate chargeback rates.
Is Local	This policy type applies to virtual array scenarios (for example, Hitachi virtual arrays), where back-end arrays are feeding front-end arrays: Yes or No
Is Thin Provisioned	Is thin provisioning enabled: Yes or No
Pool Name	Select a storage pool from the list.
RAID Type	Select a RAID Type from your environment, or enter a type.
Tiering Policy	This policy type is available only for the following arrays: Hitachi NAS (HNAS), EMC Symmetrix Fast, and Compellent. Select a storage tiering policy value, such as Tier 0 (3090-Cluster), from the list.

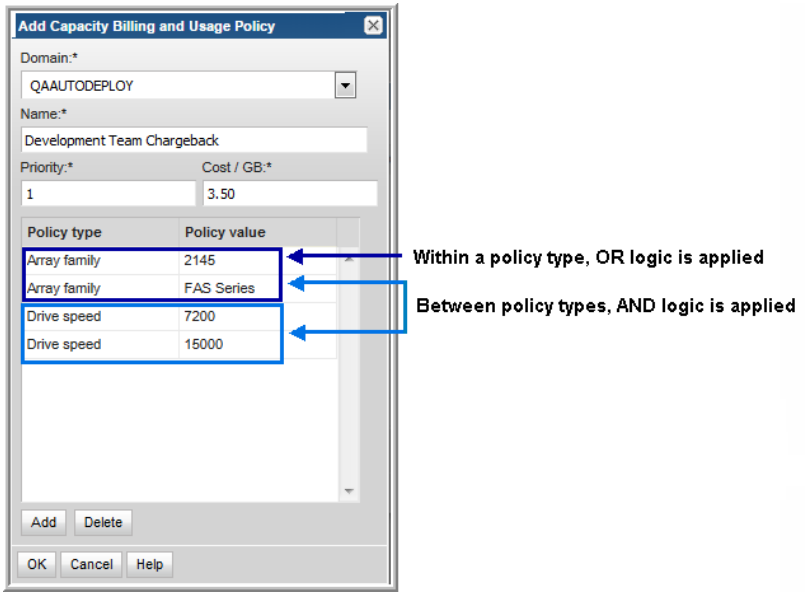
Repeat these steps to define additional Policy Types for a Capacity Chargeback Policy.

Example of Capacity Billing and Usage Policy Logic

Several policy types can be combined in a Capacity Billing and Usage Policy. Conditions within the same policy type are joined by OR, while AND logic is used between policy types.

The configuration shown in See Figure 16-1 on page 618. translates to the following logic: (Array family = 2145 OR Array family = FAS Series) AND (Drive speed = 7200 OR Drive Speed = 15000)

Figure 16-1 Example of Capacity Chargeback Policy Logic



Drive Types for Capacity Chargeback

The following table lists the drive types supported for each array vendor.

Table 16-1 Drive Types for Capacity Chargeback

Vendor and Array	Drive Type Values
Dell Compellent	FC, SAS, SSD
EMC CLARiiON	Fibre Channel, SATA, SAS, ATA, SATAII, NL SAS, FC SSD, SATA Flash, SAS Flash
EMC Isilon	SATA, SAS, SSD, Unknown
EMC Symmetrix	FC, SATA, EFD, SAS. Select **Dynamic** - for multiple drive types of Storage Resource Pool (SRP).

Table 16-1 Drive Types for Capacity Chargeback *(continued)*

Vendor and Array	Drive Type Values
HDS	■ 0:FC ■ 1:SATA ■ 2:BD ■ 4:SAS ■ 5:SSd ■ 7:SAS(SED) ■ 8:FMD ■ 9: FMC ■ Unknown FMD: Flash Module Drive SED: Self-encrypting drives SAS: Serial Attached SCSI FMC: Flash Module Compression (another name for the FMD DC2 drives)
IBM 3-4K	Fibre Channel, Serial ATA (SATA), SAS
IBM 6K	ENT, NL, FC, SSD, Volume, Unknown
NetApp	FCAL, SATA, SAS, ATA, EATA, LUN, SCSI, XATA, XSAS, FSAS, BSAS, SSD, MSATA, Unknown
NetApp Cluster-Mode	FCAL, SATA, SAS, ATA, EATA, LUN, SCSI, XATA, XSAS, FSAS, BSAS, SSD, MSATA, Unknown

Note: Beginning with release version 9.2.01P5, EMC CLARiiON drive types include more granular values for Capacity Billing and Usage Policies, as listed in the following table.

Table 16-2 Drive Types for Capacity Chargeback

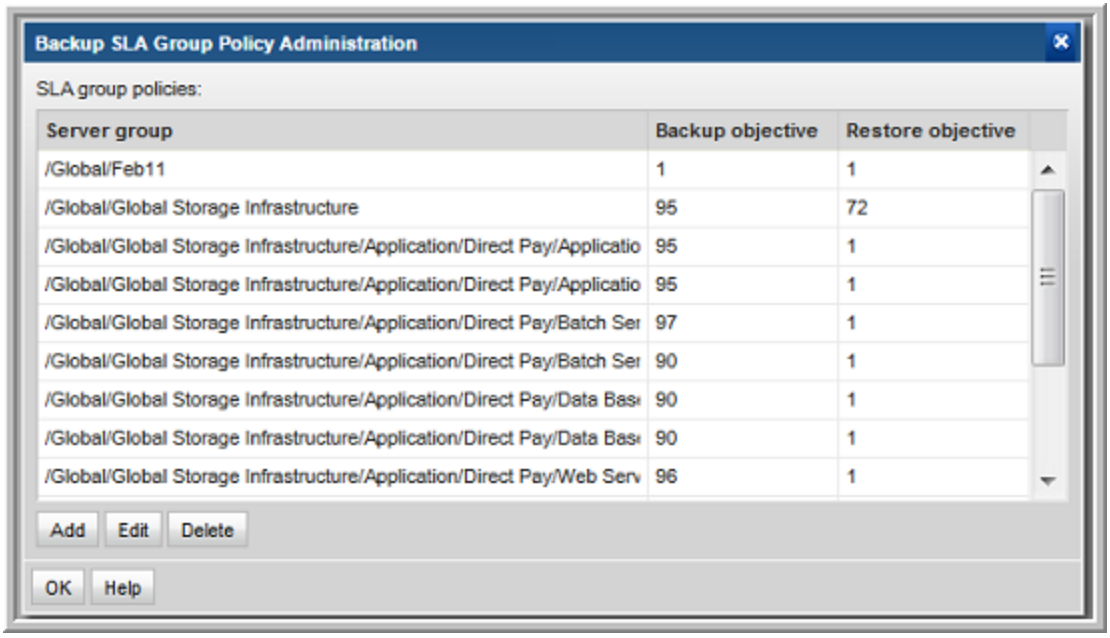
EMC CLARiiON Drive	Former Value	Current Value
FC SSD	SSD	FC SSD
NL SAS	SAS	NL SAS
SAS Flash	SSD	SAS SSD
SATA Flash	SSD	SATA SSD

Add/Edit a backup SLA policy

When you establish Service Level Agreements (SLA), you need to be able to track performance against objectives. A single service-level group policy should be configured for a user's home group. Then, in the Backup Status SLA report, you can monitor performance based on the SLA objectives.

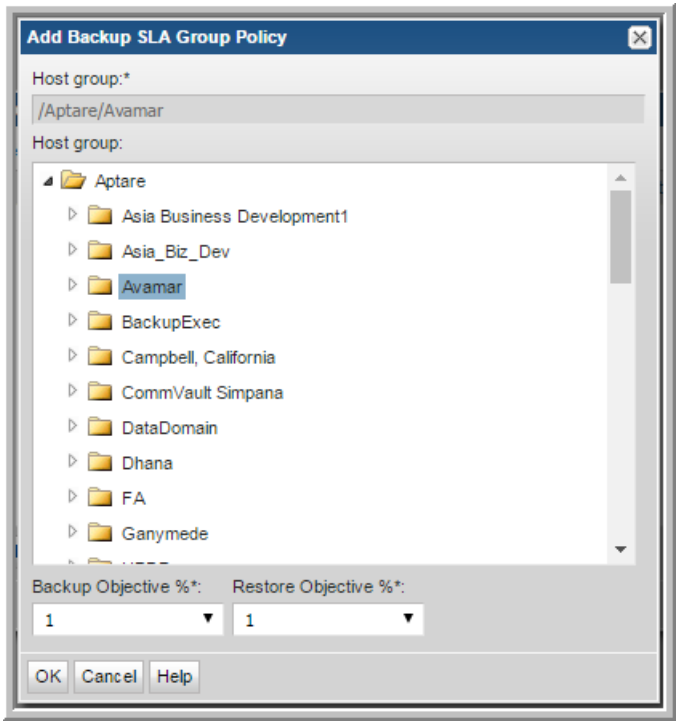
To add/edit an SLA policy

1 Select **Admin > Reports > Backup SLA**.



2 Click **Add** to add a new policy or select an existing policy and click **Edit**.

3 When you click **Add**, the **Add Backup SLA Group Policy** window launches.



4 Expand the Host group list and select a **Host group.**

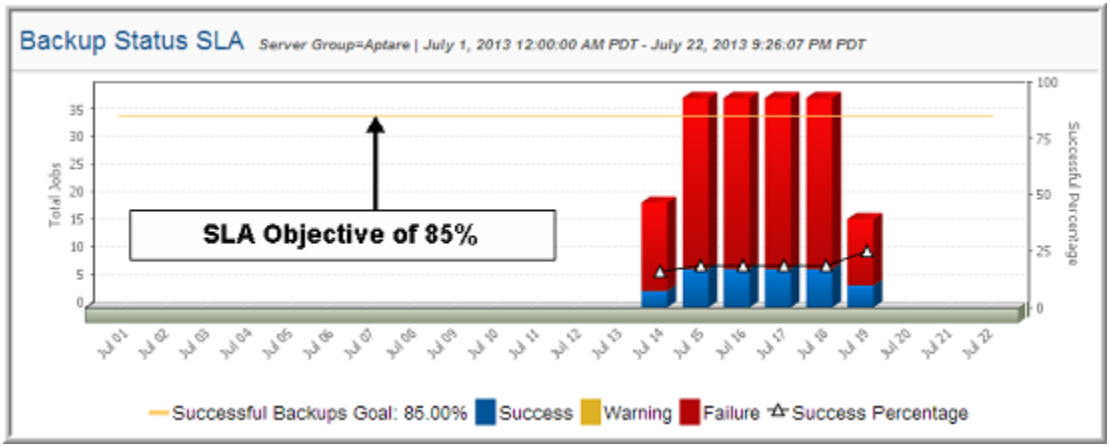
5 Select a percentage that represents the service-level expectation for successful backups from the **Backup Objective %** drop-down list.

6 Select a percentage that represents the service-level expectation for restores from the **Restore Objective** drop-down list.

7 Click **OK** to save the configuration.

Backup SLA Policy Example

The SLA percentage represents the expectation for backup success. So, if you set an objective of 85, you can monitor the success for daily backups, using the Backup Status SLA report, as shown in the following screen shot.



Add/Edit a SAN Fabric Chargeback policy

Create SAN Fabric Chargeback Policies to associate a cost with fabric and port usage. These policies are used by the SAN Fabric Usage report.

To add/edit a SAN Fabric Chargeback policy

1. Select **Admin > Chargeback > SAN Fabric**.
2. Either click **Add** or select an existing Domain and click **Edit**.
3. Click **Add** to add Policy Types and Values. Use the following table to supply values in this window.

Domain	Select a domain from the list. Typically, only one domain is listed, unless the portal is maintained by a Managed Services Provider (MSP).
Name	Enter a name for the policy.
Priority	Enter a priority, where 1 is the highest priority.
TB cost	Enter a cost per terabyte.
Port cost	Enter a cost per switch port.

Policy Type & Policy Value	Possible choices include: ■ Port speed (Gbps) - 1, 2, 4, 8, or 10 Gbps ■ Fabric name - The fabric values list is derived from the SAN Fabric data that has been collected for your environment. ■ Domain - The values list is derived from Domains that are configured in the Portal.

IT Analytics Billing and Chargeback policies

This section describes how to create Billing and Chargeback Policies. This policies are then used by the following reports, as described in the Report Reference Guide.

- Billing & Chargeback Summary
- Server Consumption Summary

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.

Add/Edit a Billing and Chargeback policy

This procedure assumes that your Host Group hierarchy represents the logical relationships--business units, customers, etc.--on which you want to report. When you set the per-GB and per-tape usage costs, you do so at the host group level. Therefore, child host group policies take precedence over parent groups. If a parent host group (Group A) has a billing and usage policy set at $0.50/Gbyte and $35.00/Tape and a child host group (Group B) has a usage policy set at $0.55/Gbyte and $40.00/Tape, the host contained in Group B would be billed at $0.55/Gbyte and $40.00/Tape rather than the billing rate of the parent group.

When IT Analytics discovers that there is no policy set for any given group, it traverses the host group hierarchy looking for a policy in a parent group until it finds a billing and usage policy to use. If the application cannot find any policy as it

searches up the host group hierarchy, a value of $0.00 will be assigned for both per-GB and per-tape usage.

To add a new billing policy for a Host Group

1. Click **Admin > Chargeback > Backup**. The **Backup Billing And Usage Policies** window is displayed.
2. Click **Add**.
3. Select a Host Group name to which the policy will apply. Note that a policy set at any particular host group will be assumed by all child host groups that do not have a policy.
4. Enter values for Cost Per GByte, Cost Per Tape, and Cost per Duplicated GB and then click **OK**.

Cost per GByte	The amount each host within the host group should be billed per GB of disk space consumed. The acceptable range is from $0.0001 - $999.9999.
Cost per tape	The amount each host within the host group should be billed per tape drive consumed. The acceptable range is from $0.0001 - $999.9999.
Cost per duplicated GB	The amount each host within the host group should be billed per duplicated consumption. The acceptable range is from $0.0001 - $999.9999.

Chapter 17

Solutions administration

This chapter includes the following topics:

- Storage Optimization solution overview
- Understand the value of Storage Optimization historical data
- Configure Storage Optimization rules
- Enable Storage Optimization rules
- Storage optimization rule prerequisites and logic
- Risk Mitigation solution overview
- Configure Risk Mitigation rules

Storage Optimization solution overview

Storage teams continually seek actionable analytics that can be used to eliminate inefficient storage utilization. Growing storage requirements and shrinking budgets demand on-going processes. The Storage Optimization solution identifies both on-prem and cloud resources that may be reclaimable or that can be optimized.

Analytics can be used to find dark storage, in an effort to contain costs and delay/avoid purchases. By examining collected historical data for multiple data points, Storage Optimization provides visibility into storage utilization over time, with true analytics that support infrastructure rightsizing. Already collected data is leveraged to deliver a detailed storage utilization perspective. For example, while collecting data from storage arrays, the system can use the LUN-to-host mapping to retrieve additional host details, such as the capacity that has been allocated to a host. Often storage allocated to a host never gets used.

The following scenarios represent a sampling of challenges faced by data center teams:

- Decommissioned hosts leave behind storage that is never de-provisioned.
- Overprovisioned hosts may offer candidates for storage reclamation.
- VMs that are not in the VM inventory consume disks in datastores.
- Temporary VM snapshots are forgotten, yet their files consume space.
- Idle and powered-off VMs continue to have storage allocated to them.
- Expensive tier-1 storage contains files that should be migrated to less expensive storage.
- LUN reclamation opportunities can be realized by identifying unallocated LUNs, allocated but undiscovered LUNs, discovered but unused LUNs, and LUNs that have had no activity in the past 30 days.
- Mergers, acquisitions, and migrations require illumination of storage usage.

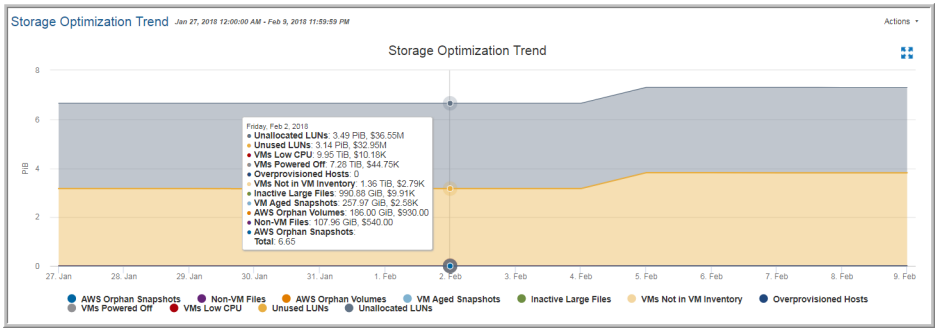
To learn how to use the Storage Optimization solution to address common use cases, explore the following:

- See "Understand the value of Storage Optimization historical data" on page 626.
- See "Configure Storage Optimization rules" on page 627.
- See "Enable Storage Optimization rules" on page 638.
- See "Storage optimization rule prerequisites and logic" on page 641.

Understand the value of Storage Optimization historical data

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.

When analyzing storage utilization data over time, interesting trends emerge, as illustrated in the following example. This chart shows that months after successful storage reclamation efforts, storage consumption gradually climbed again. While this storage growth may be valid, this trend warrants further investigation because the growth occurred in categories known for inefficient use of storage resources.



The Storage Optimization process runs in the background on a regular schedule, gleaning historical data from already collected data. This process uses the parameters in active optimization rules to define the data that it persists in database tables designed for Storage Optimization reporting. The following steps provide a high-level view of the sequence of Storage Optimization actions.

1. Configure Storage Optimization rules to specify the criteria that the process uses to filter historical data.

   See “Configure Storage Optimization rules” on page 627.

2. Activate rules relevant for your environment.

   See “Enable Storage Optimization rules” on page 638.

3. Schedule the Storage Optimization collection process.

   See “Enable Storage Optimization rules” on page 638.

Configure Storage Optimization rules

A set of storage optimization rules are provided 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. You can also associate costs with rules. Cost sources may related to Chargeback values or you can set your own custom value. For example, the Unallocated LUNs rule can be configured to exclude LUNs less than a certain size. While various use cases drive how you configure a rule, the goal is to have analytics that help you identify storage optimization candidates and trends that require attention. This on-going process should periodically assess trends and codify business practices.

See “Storage Optimization solution overview” on page 625.

See “Enable Storage Optimization rules” on page 638.

See *Report Reference Guide* for details on Storage Optimization Solution Reports.

Once configured, a scheduled process gathers historical data for these categories so that you can identify areas that require further scrutiny. Accompanying reports present data that can be monitored over time, enabling an actionable process to maintain an optimized storage environment.

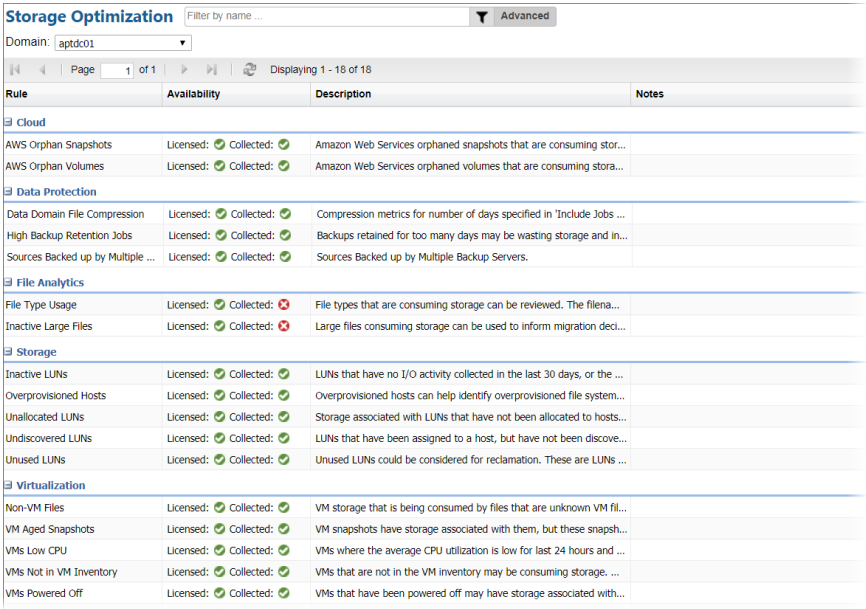
Best Practice

When configuring values for parameters, be as liberal as possible, initially. Then, over time, change parameters to produce a narrower actionable list. For optimal data comparisons, avoid frequent parameter modifications.

To edit a storage optimization rule

If optimization rules are not modified, the historical data process uses an active rule's default settings to collect the historical data.

1 Select **Admin > Solutions > Storage Optimization**.



Rule	Rules are listed within relevant categories, such as Cloud and Storage.

Availability	If a particular type of collection is not licensed or collected, storage optimization data will not be available, regardless of how a rule is configured. In some cases, a Portal may have the necessary license, but collection may not have been enabled and/or completed. ■ Cloud does not require a specific license in order to deploy a data collector. ■ Data Protection requires Protection or Complete Suite license. ■ File Analytics requires a Complete Suite license. ■ Storage requires Storage or Complete Suite license. ■ Virtualization requires Protection or Complete Suite license.
Description	The full description of the Storage Optimization rule can be viewed by placing your mouse over the description.
Notes	Enter operational notes for future reference.
Status	■ Green check mark indicates successful collection of storage optimization historical data for enabled rules. ■ Red X indicates failed historical data collection. It could be that collection is attempting to access data for a product module that is not in your Portal environment. Click the red icon to view the Database Error Aggregation report. ■ A non-colored circle indicates that the storage optimization process did not run, typically because the rule is not enabled.
State	Indicates if the rule is Enabled or Disabled.

Last Run	The date and time that the Storage Optimization process ran and evaluated the collected data against the rule's configured parameters.

2 Select a rule in the Storage Optimization grid and click **Edit**. Or, simply double-click the rule to access the edit dialog.

See "Storage optimization rule prerequisites and logic" on page 641.

Storage Optimization Rule	Description
Cloud Rules	
AWS Orphan Snapshots	Amazon Web Services orphaned snapshots that are consuming storage may be impacting costs. This rule identifies snapshots for EC2 instances that no longer exist. When an EC2 instance is deleted, its snapshots need to be reviewed and its volumes brought back into the pool of usable storage. This rule also allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For actionable information, see the Storage Optimization detail reports.
AWS Orphan Volumes	Amazon Web Services orphaned volumes that are consuming storage may be impacting costs. This rule identifies volumes for EC2 instances that no longer exist. When an EC2 instance is deleted, its snapshots need to be reviewed and its volumes brought back into the pool of usable storage. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For actionable information, see the Storage Optimization reports.
Data Protection Rules	
Sources Backed Up by Multiple Servers	Sources backed up by multiple servers may be wasting storage and increasing maintenance costs. This rule identifies sources backed up by more than one server within a specific number of days (default is 7 days). If a source is moved from one backup system to another backup system, this rule would identify those changes too. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. The cost is calculated based on the number of different servers minus 1, times the average job size in GiB. For example, if a source is backed up by 3 different servers, the average size of these jobs is 200 GiB, and $0.25 for each GiB, then the saving might be: (3-1) * 200 GiB * $0.25 = $100

Storage Optimization Rule	Description
Data Domain File Compression	Compression metrics for the last 30 days are evaluated for inefficient storage usage and compression ratios for Data Domain clients. A low compression ratio may indicate a storage optimization opportunity or it may warrant a move of low compression clients to less expensive storage. This rule also allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For additional details, see the Data Domain NetBackup File Compression Summary report.
High Backup Retention Jobs	Backups retained for too many days may be wasting storage and increasing maintenance costs. This rule identifies backups with a high number of retention days. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports.
File Analytics Rules	
File Type Usage	File types that are consuming storage can be reviewed. The filename extension identifies a file type, such as iso, log, and cab. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. This File Type rule is relevant only if File Analytics data collection is enabled. For a list of file types that are relevant for your environment, see the File Types report.
Inactive Large Files	Large files consuming storage can be considered when making tiered storage migration decisions. This rule also allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. This rule is relevant only if File Analytics data collection is enabled. For additional details, see the Largest Files report.
Storage Rules	

Storage Optimization Rule	Description
Inactive LUNs	LUNs that have no I/O activity collected in the last 30 days, or the collected value during that time period is zero, are candidates for storage reclamation. This rule explicitly examines 30 days of collected LUN performance data. This rule can filter the data by RAID type. The format for RAID type specification is specific to storage vendors. For a list of RAID types in your environment, run the LUN Utilization Summary report. This rule also allows you to set a Cost Source for displaying associated costs in the generated reports. Choose from Chargeback values or add a custom value. If you have set up Chargeback rules, those are more granular and can provide greater cost accuracy. See "Storage optimization rule prerequisites and logic" on page 641.
Overprovisioned Hosts	File system usage can help identify overprovisioned file systems and hosts. For additional information, run the Host Utilization Summary report and the Host Filesystem Reclamation Candidates cloud report. This rule also allows you to set a Cost Source for displaying associated costs in the generated reports. Choose from Chargeback values or add a custom value. If you have set up Chargeback rules, those are more granular and can provide greater cost accuracy. See "Storage optimization rule prerequisites and logic" on page 641.
Unallocated LUNs	Storage associated with LUNs that have not been allocated to hosts can be considered for storage optimization. This rule also allows you to set a Cost Source for displaying associated costs in the generated reports. Choose from Chargeback values or add a custom value. If you have set up Chargeback rules, those are more granular and can provide greater cost accuracy. This rule can filter the data by RAID type. The format for RAID type specification is specific to storage vendors. For a list of RAID types in your environment, run the LUN Utilization Summary or the Unallocated LUNs report. In addition, the Reclamation Summary report can reveal potential reclamation categories. See "Storage optimization rule prerequisites and logic" on page 641.

Storage Optimization Rule	Description
Undiscovered LUNs	LUNs that have been assigned to a host, but have not been discovered and therefore are not seen on the host side, indicate a reclamation opportunity. The storage in these orphaned LUNs is not available for mounting file systems. This rule also allows you to set a Cost Source for displaying associated costs in the generated reports. Choose from Chargeback values or add a custom value. If you have set up Chargeback rules, those are more granular and can provide greater cost accuracy. This rule can filter the data by RAID type. The format for RAID type specification is specific to storage vendors. For a list of RAID types in your environment, run the LUN Utilization Summary report. In addition, the Reclamation Summary report illustrates potential reclamation categories. See "Storage optimization rule prerequisites and logic" on page 641.
Unused LUNs	Unused LUNs could be considered for reclamation. 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. This rule also allows you to set a Cost Source for displaying associated costs in the generated reports. Choose from Chargeback values or add a custom value. If you have set up Chargeback rules, those are more granular and can provide greater cost accuracy. This rule can filter the data by RAID type. The format for RAID type specification is specific to storage vendors. For a list of RAID types in your environment, run the LUN Utilization Summary report. In addition, the Reclamation Summary report illustrates potential reclamation categories. See "Storage optimization rule prerequisites and logic" on page 641.
Virtualization Rules	
Non-VM Files	VM storage that is being consumed by files that are unknown VM file types can be considered for storage optimization. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For additional details, see the Datastore Usage Breakdown and the VM Files Summary reports. See "Storage optimization rule prerequisites and logic" on page 641.

Storage Optimization Rule	Description
VMs Aged Snapshots	VM snapshots have storage associated with them, but these snapshots have been forgotten for some time. Aged snapshots present a reclamation opportunity. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For additional details, see the VM Snapshot Summary. See "Storage optimization rule prerequisites and logic" on page 641.
VMs Low CPU	VMs where the average CPU utilization is low for last 24 hours and CPU usage is less than 5%, may provide a storage optimization opportunity. Filter your actionable list by guest size and guest state, for example, including only running large guests. This rule allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For additional details, see the VM Summary report. See "Storage optimization rule prerequisites and logic" on page 641.
VMs Not in VM Inventory	VMs that are not in the VM inventory may be consuming storage. When a VM is removed from the inventory, it doesn't necessarily mean that the associated storage is returned to the storage pool. This rule also allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For additional details, see the Datastore Usage Breakdown and the VM Files Summary reports. See "Storage optimization rule prerequisites and logic" on page 641.
VMs Powered Off	VMs that have been powered off may have storage associated with them. For additional details, see the VM Summary report. This rule also allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. See "Storage optimization rule prerequisites and logic" on page 641.

Storage Optimization Rule	Description
VMs Undiscovered Disks	VMs that have been provisioned storage, but that storage is not being used by the VMs, may indicate a storage optimization opportunity. This rule also allows you to set a custom value for Cost per GiB for the purpose of displaying costs in the generated reports. For additional details, see the Physical Disk Utilization report. See "Storage optimization rule prerequisites and logic" on page 641.

3 Click in a parameter field to view the green help text at the bottom of the dialog and then use those details to configure parameters. Each rule includes parameters unique to the data that is being evaluated.

See "Edit Storage Optimization Rule Example" on page 636.

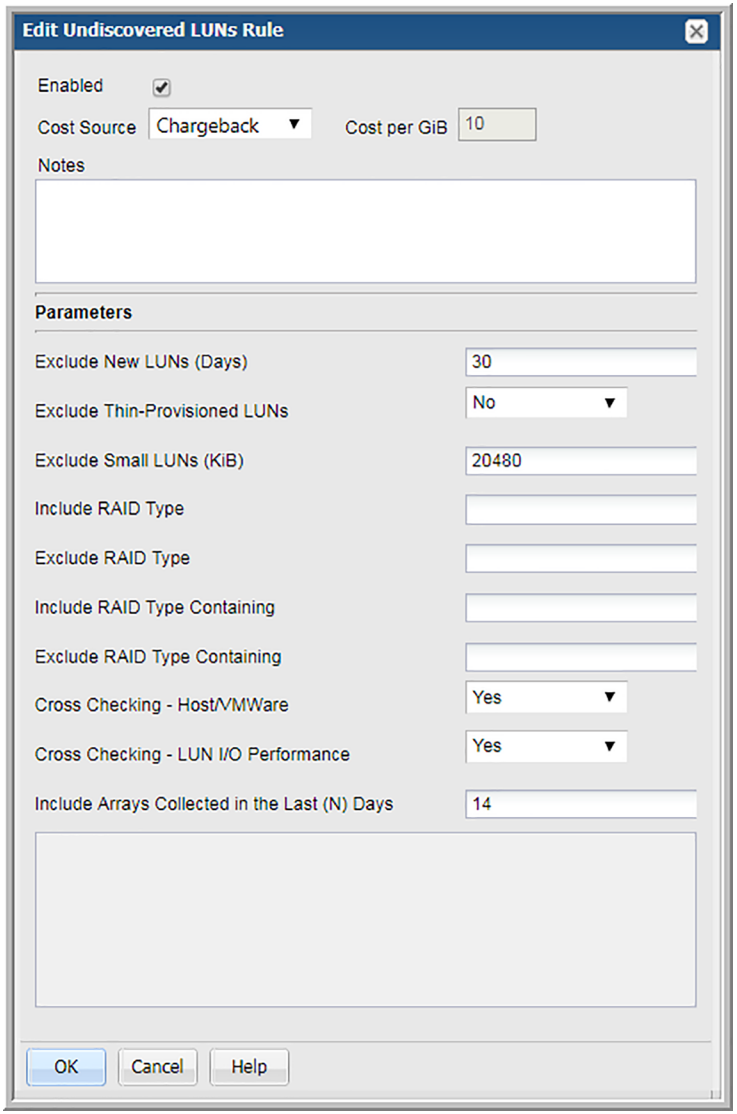
See "Common storage optimization parameters" on page 637.

Edit Storage Optimization Rule Example

Because each Storage Optimization rule has its own set of parameters, rely on the green help text in the Edit dialog to guide your configuration.

The following example illustrates the types of parameters that can be adjusted to glean an actionable list from collected data.

See "Common storage optimization parameters" on page 637.



Common storage optimization parameters

The parameters shown in the following list represent parameters that are common to multiple storage optimization rules. This list, however, does not represent a comprehensive list of storage optimization parameters. Refer to the green help text in each rule for parameter specifics.

See “Configure Storage Optimization rules” on page 627.

See "Edit Storage Optimization Rule Example" on page 636.

Table 17-1 Common storage optimization parameters

Parameter	Description and Value Examples
Cost Source	Select a cost source to enable another dimension of Storage Optimization. By selecting a cost source, you can associate a monetary value to storage inefficiencies. Choose to use the default Chargeback values or set your own cost value.
Host Type	Examples of host types include: CIFS, Windows, and Linux. For a list of host types in your environment, run the Host Summary report. A comma-separated list of case-sensitive host types is supported. A null value signifies that this parameter will be omitted from processing.
Include/Exclude File Type	Enter a comma-separated list of case-sensitive file types collected via File Analytics. Examples of file types include: zip, out, dmp, iso, gz. For a list of file types in your environment, run the File Types report. A null value signifies that this parameter will be omitted from processing.
Include/Exclude RAID Type	The format for RAID type specification is specific to storage vendors. For a list of RAID types in your environment, run the LUN Utilization Summary report. A comma-separated list of case-sensitive RAID types is supported. A null value signifies that this parameter will be omitted from processing.
Include/Exclude RAID Type Containing	Include the RAID types containing the specified characters. The format for RAID type specification is specific to storage vendors. For a list of RAID types in your environment, run the LUN Utilization Summary report. A comma-separated REGEXP expression is supported. To exclude this parameter from the rule, use a blank value.
Include Guest State	Possible VM guest states include: running, shuttingdown, resetting, standby, notRunning, and unknown. A comma-separated list of case-sensitive states is supported. A null value signifies that this parameter will be omitted from processing.

Enable Storage Optimization rules

identifies candidates for optimization or reclamation, found in common use cases, and provides rules that categorize these use cases. The optimization rules include

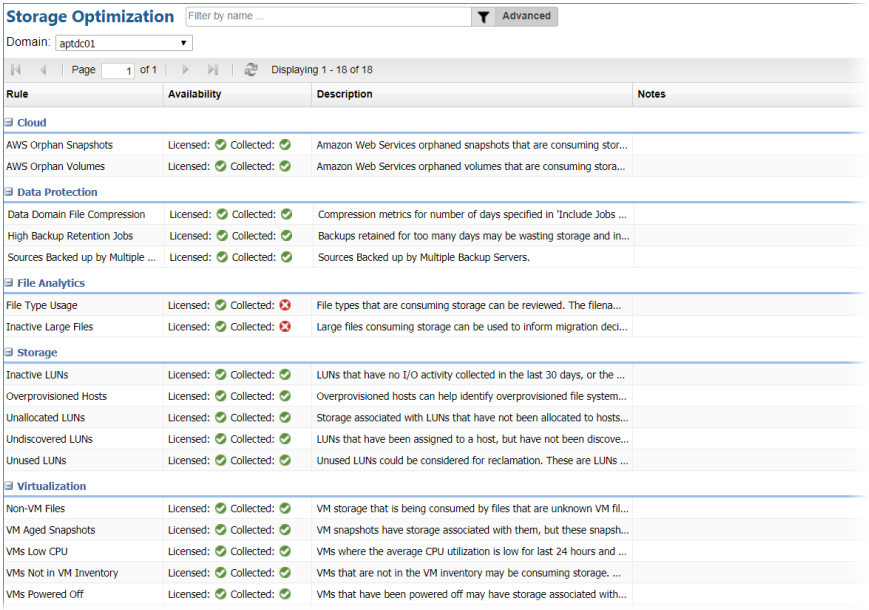
dials (parameters) that can be set to filter collected data for storage optimization reports.

- See “Understand the value of Storage Optimization historical data” on page 626.
- See “Configure Storage Optimization rules” on page 627.
- See “Storage optimization rule prerequisites and logic” on page 641.
- See *Report Reference Guide* for details on Storage Optimization Solution Reports.

To enable storage optimization rules

1 Select **Admin > Solutions > Storage Optimization**.

These rules filter already collected data to produce actionable lists to support your storage optimization goals.



Rule	Availability	Description	Notes
Cloud			
AWS Orphan Snapshots	Licensed: ✓ Collected: ✓	Amazon Web Services orphaned snapshots that are consuming stor...	
AWS Orphan Volumes	Licensed: ✓ Collected: ✓	Amazon Web Services orphaned volumes that are consuming stora...	
Data Protection			
Data Domain File Compression	Licensed: ✓ Collected: ✓	Compression metrics for number of days specified in 'Include Jobs ...	
High Backup Retention Jobs	Licensed: ✓ Collected: ✓	Backups retained for too many days may be wasting storage and in...	
Sources Backed up by Multiple ...	Licensed: ✓ Collected: ✓	Sources Backed up by Multiple Backup Servers.	
File Analytics			
File Type Usage	Licensed: ✓ Collected: ✗	File types that are consuming storage can be reviewed. The filena...	
Inactive Large Files	Licensed: ✓ Collected: ✗	Large files consuming storage can be used to inform migration deci...	
Storage			
Inactive LUNs	Licensed: ✓ Collected: ✓	LUNs that have no I/O activity collected in the last 30 days, or the ...	
Overprovisioned Hosts	Licensed: ✓ Collected: ✓	Overprovisioned hosts can help identify overprovisioned file system...	
Unallocated LUNs	Licensed: ✓ Collected: ✓	Storage associated with LUNs that have not been allocated to hosts...	
Undiscovered LUNs	Licensed: ✓ Collected: ✓	LUNs that have been assigned to a host, but have not been discove...	
Unused LUNs	Licensed: ✓ Collected: ✓	Unused LUNs could be considered for reclamation. These are LUNs ...	
Virtualization			
Non-VM Files	Licensed: ✓ Collected: ✓	VM storage that is being consumed by files that are unknown VM fil...	
VM Aged Snapshots	Licensed: ✓ Collected: ✓	VM snapshots have storage associated with them, but these snapsh...	
VMs Low CPU	Licensed: ✓ Collected: ✓	VMs where the average CPU utilization is low for last 24 hours and ...	
VMs Not in VM Inventory	Licensed: ✓ Collected: ✓	VMs that are not in the VM inventory may be consuming storage. ...	
VMs Powered Off	Licensed: ✓ Collected: ✓	VMs that have been powered off may have storage associated with...	

Best Practice

Activate rules and then periodically review results in reports to identify trends for which interventions may be required. For example, initial results may reveal candidates for reclamation. Storage reclamation, however, is not a one-time event. Over time, storage usage may creep up in unsuspecting places. Therefore, processes must be in place to monitor storage growth and also to continually identify gaps in operational processes.

2 Enable the rules applicable to your storage optimization goals by double-clicking a rule and clicking the **Enabled** checkbox to set the rule's state.

Note: Even if optimization rules are not modified, the historical data process uses an enabled rule's default settings to filter the historical data.

---

3 Modify rule parameters, if needed, using the details listed in the following.

See “Configure Storage Optimization rules” on page 627.

4 Schedule the process that collects historical data for trend reports by clicking the clock icon at the upper right of the Storage Optimization window.

Best Practice

Use a monthly schedule to capture point-in-time data that can be compared month-to-month in storage optimization trending reports.

Storage optimization rule prerequisites and logic

Simply enabling a rule does not guarantee that relevant data will be collected. The following table lists prerequisites and dependencies for successful reporting. For an overview of each rule, refer to the following.

See “Configure Storage Optimization rules” on page 627.

Table 17-2 Storage optimization rule prerequisites and logic

Storage Optimization Rule	Prerequisites	Rule Processing Logic
Cloud Rules and Logic		
AWS Orphan Snapshots	■ AWS EC2 snapshot data must be collected ■ Collected snapshots do not have EC2 instances	■ Checks if a corresponding EC2 instance is found ■ Rule parameters are used to exclude small volumes, new snapshots, and snapshots with specific states
AWS Orphan Volumes	■ AWS EC2 volumes without snapshots must be collected ■ Collected volumes do not have EC2 instances	■ Checks if a corresponding EC2 instance is found ■ Checks if a corresponding EC2 Snapshot is found ■ Rule parameters are used to exclude small volumes and new volumes so that it can ignore volumes that have been orphaned for only a few days
Data Protection Rules and Logic		

Table 17-2 Storage optimization rule prerequisites and logic *(continued)*

Storage Optimization Rule	Prerequisites	Rule Processing Logic
Data Domain Compression	■ NetBackup data must be collected ■ Data Domain Inventory and file-level compression data must be collected within the last 30 days	■ Rule parameters are used to exclude new clients, small capacity, and high compression ratios because low compression ratios can be used to identify optimization opportunities ■ **Note**: Initial backups may have only a small system effective compression ratio
High Backup Retention Jobs	■ Backup jobs must be collected. ■ Collected jobs must be backup events with a status of success or warning, and cannot be expired.	■ Rule parameters are used to exclude small jobs, and jobs with low backup retention
Sources Backed Up by Multiple Servers	■ Backup jobs must be collected.	■ Rule parameter is used to include jobs that finish within last few days.
File Analytics Rules and Logic		
Inactive Large Files	■ File Analytics data must be collected	■ Rule parameters are used to include/exclude host and file types; specify size of small files to exclude and the number of days for active files ■ User must determine if an old large file can offer reclaimable space, based on the last access date
File Type Usage	■ File Analytics data must be collected	■ Rule parameters are used to include/exclude host and file types; exclude file types with small number of files and consumed storage
Storage Rules and Logic		

Table 17-2 Storage optimization rule prerequisites and logic *(continued)*

Storage Optimization Rule	Prerequisites	Rule Processing Logic
Inactive LUNs	■ LUN data must be collected ■ LUN performance data must be collected to determine if I/O activity occurred within the days specified in the rule parameters ■ File systems discovered: mount points and file system names collected	■ Excludes internal LUNs, such as RDFs (EMC Symmetrix remote data facility) ■ Checks if the LUN has a storage path ■ Checks the LUN creation date in order to exclude new LUNs ■ Rule parameters are used to include/exclude new LUNs, thin-provisioned LUNs, small LUNs, and specified RAID types
Overprovisioned Hosts	■ Array LUN data must be collected ■ Host Inventory or VMware data must be collected. Host and file systems discovered and in use: mount points and file system names collected ■ Only SAN storage is supported	■ Checks if the LUN has a storage path ■ Checks file system activity for no used capacity growth in the last 90 days ■ Excludes internal LUNs, such as RDFs (EMC Symmetrix remote data facility) ■ Rule parameters are used to include/exclude thin-provisioned LUNs, small file systems, and low available space ■ Greater precision can be achieved if LUN performance data is collected to determine if there was any read I/O, even when usage has not grown ■ Best Practice: Use a small growth rate percentage to isolate overprovisioned hosts

Table 17-2 Storage optimization rule prerequisites and logic *(continued)*

Storage Optimization Rule	Prerequisites	Rule Processing Logic
Unallocated LUNs	■ Array LUN data must be collected	■ Excludes internal LUNs, such as RDFs (EMC Symmetrix remote data facility) ■ Checks that no storage path exists for the LUN ■ Rule parameters are used to include/exclude new LUNs, thin-provisioned LUNs, small LUNs, and specified RAID types
Undiscovered LUNs	■ Array LUN data must be collected	■ Excludes internal LUNs, such as RDFs (EMC Symmetrix remote data facility) ■ Checks if the LUN has a storage path ■ Rule parameters are used to include/exclude new LUNs, thin-provisioned LUNs, small LUNs, and specified RAID types ■ Greater precision can be achieved if VMware or Host Inventory data is collected (to verify that no hosts have been discovered) and if LUN performance data is collected

Table 17-2 Storage optimization rule prerequisites and logic *(continued)*

Storage Optimization Rule	Prerequisites	Rule Processing Logic
Unused LUNs	■ Array LUN data must be collected ■ Host Inventory or VMware data must be collected	■ Excludes internal LUNs, such as RDFs (EMC Symmetrix remote data facility) ■ Checks if the LUN has a storage path ■ Checks that no file systems have been discovered ■ Rule parameters are used to include/exclude new LUNs, thin-provisioned LUNs, small LUNs, and specified RAID types ■ Greater precision can be achieved if LUN I/O performance data is collected
Virtualization Rules and Logic		
Non-VM Files	■ VMware Inventory data must be collected ■ Datastore Scan data must be collected	■ Rule parameters are used to exclude small and/or recently modified files and specific file types
VM Aged Snapshots	■ VMware Inventory data must be collected ■ Datastore Scan data must be collected	■ Rule parameters are used to exclude small files and new snapshots (based on creation date)
VM Low CPU	■ VMware Inventory data must be collected	■ Examines CPU performance statistics for the last 24 hours ■ 0% CPU usage treated as an idle machine ■ Rule parameters are used to exclude small VM guests and include low CPU usage

Table 17-2 Storage optimization rule prerequisites and logic *(continued)*

Storage Optimization Rule	Prerequisites	Rule Processing Logic
VM Not in VM Inventory	■ VMware Inventory data must be collected ■ Datastore Scan data must be collected	■ Checks if VM files have been discovered ■ Rule parameters are used to exclude small and/or recently modified files and specific file types
VM Powered Off	■ VMware Inventory data must be collected	■ Checks if VM guest is off ■ Rule parameters are used to exclude small VM guests and VM guest recently powered off
VM Undiscovered Disks	■ VMware Inventory data must be collected	■ Checks if VM guest is on and state is running ■ Rule parameters are used to exclude small VM guests and include low virtual disk usage ■ Greater precision can be achieved if VM performance data is collected ■ Best Practice: To identify unused, wasted disk space, enter a high value for the Include Low Virtual Disk % parameter

Risk Mitigation solution overview

Data center administrators must manage infrastructure risk for both backup/recovery and storage operations to prevent downtime for critical applications. Backup and recovery plans rely on data that supports disaster recovery and compliance objectives. Storage administrators must balance storage utilization to optimize performance for critical applications. In partnership with the Storage Optimization Solution, the Risk Mitigation Solution supplies not only data, but curated analytics to enable proactive management for backup compliance and storage performance.

See “Configure Risk Mitigation rules” on page 647.

Configure Risk Mitigation rules

A set of risk mitigation rules are provided to assess areas within your enterprise that may be at risk of meeting data protection objectives. These rules include parameters that can be configured to isolate specific conditions relevant for your environment. For example, the Clients with No Recent Backups rule can be modified to specify the number of days for which no backups occurred and also to exclude retired clients. While various use cases drive how you configure a rule, the goal is to have analytics that help you identify areas at risk and trends that require attention. This on-going process should periodically assess trends and codify business practices.

See “Risk Mitigation solution overview” on page 646.

Once configured, a scheduled process gathers historical data for these categories so that you can identify areas that require further scrutiny. Accompanying reports present data that can be monitored over time, enabling an actionable process to reduce risk.

Best Practice

When configuring values for parameters, be as liberal as possible, initially. Then, over time, change parameters to produce a narrower actionable list. For optimal data comparisons, avoid frequent parameter modifications.

To edit a risk mitigation rule

If risk mitigation rules are not modified, the historical data process uses an active rule’s default settings to collect the historical data.

1 Select **Admin > Solutions > Risk Mitigation**

Rule	Rules are listed within relevant categories, such as Cloud and Storage.
Availability	If a particular type of collection is not licensed or collected, risk mitigation data will not be available, regardless of how a rule is configured. In some cases, a Portal may have the necessary license, but collection may not have been enabled and/or completed. ■ Data Protection rules require a Backup Manager license. ■ Storage requires a Capacity Manager license.
Description	The full description of the Risk Mitigation rule can be viewed by placing your mouse over the description.
Notes	Enter operational notes for future reference.

Status	■ Green check mark indicates successful collection of risk mitigation historical data for enabled rules. ■ Red X indicates failed historical data collection. It could be that collection is attempting to access data for a product module that is not in your Portal environment. Click the red icon to view the Database Error Aggregation report. ■ A non-colored circle indicates that the background process did not run, typically because the rule is not enabled.
State	Indicates if the rule is Enabled or Disabled.

Last Run	The date and time that the background process ran and evaluated the collected data against the rule's configured parameters.

2 Select a rule in the Risk Mitigation grid and click **Edit**. Or, simply double-click the rule to access the edit dialog.

Risk Mitigation Rule	Description
Data Protection Rules	
Backup Job Duration Variance	Compares a client's average job duration to help identify clients with successful backups where the current period's average job duration is more than the prior period's average job duration. Significant differences in average backup job duration may indicate a data protection issue.
Backup Job Size Variance	Compares average job size for clients with successful backups where the current period's average is lower than the prior period's average job size, which may help to identify backup issues.
Compliance RTO RPO	Considers RTO (Recovery Time Objectives) and RPO (Recovery Point Objectives) for backups by determining when/if the last full backup was performed. Then, add in the time it takes to apply any incremental backups. Assists in computing RTO (Recovery Time Objectives) and RPO (Recovery Point Objectives) for backups by determining when/if the last full backup was performed. Then, add in the time it takes to apply any incremental backups to determine if you are meeting your SLAs.
NetBackup Disk Pool Forecast	Provides NetBackup disk pool statistics for the number of weeks in the selected period are examined to forecast the date when storage will run out within the next three years. If the prediction is beyond three years, a status is returned.

Risk Mitigation Rule	Description
Source Overall Status Summary	Considers sources for which backup jobs were not successful to determine risk. This rule helps in finding such sources by providing Status Summary. Determining if source backups were successful is complicated, especially if there are multiple policies and schedules defined for that source and if there are multiple streams per backup set. Also there needs to be an established cutoff time to determine what to do if a source is still running or has not made all of its attempts. The following criteria is considered: 1. If a source fails all of its jobs it is failed. 2. If a source successfully completes all of its jobs it is successful. 3. If a source completes all of its jobs with status 1 (skipped files) it was partially successful and probably OK. 4. If a source has a mixture of successful jobs and failed jobs, it needs further examination to determine if the jobs were truly successful. Now, there is logic that can be applied to #4 in order to programmatically determine whether a source was successful or not, but that logic varies from customer to customer.
Sources Consecutive Failure	Evaluates sources where consecutive backups have failed or no backups have occurred for consecutive days. This rule examines the past 14 days of history, providing insights to possible problematic areas. **Best Practice**: Schedule this rule to run every day at the end of the backup window. This rule works with any backup product.
Sources with No Recent Backups	Reviews details of sources that have not been backed up in a defined number of days to help determine if the sources are at risk. Specify the number of days for which backups have not occurred to determine the risk.
Suspect Backups with Job Size	Backup jobs with an unexpected small size may signify issues with the policy setup.
Storage Rules	

Risk Mitigation Rule	Description
Host Multi-Pathing Exposure	Identifies hosts that are at risk because they have less than the specified number of paths. Examines LUN mappings of hosts that do not have multiple HBA ports and Array ports configured between a host and a LUN. Normally, the requirement is 2 HBA and 2 Array ports are configured, so when any HBA or Array port fails, there is another port to keep the connection between the host and the LUN.
Hot Array Ports	Identifies overactive array ports, which may indicate a risk to application performance. Array port performance data is examined to identify spikes in data transferred.
Hot LUNs by Read IO	Reveals spikes in Read I/O performance metrics, which may indicate an area of risk. This rule uses a unique yet simple algorithm to identify abnormal performance patterns.
Hot LUNs by Read Response	Reveals spikes in Read Response Time metrics, which may indicate an area of risk. This rule uses a unique yet simple algorithm to identify abnormal performance patterns.
Hot LUNs by Write IO	Reveals spikes in Write I/O activity, which may indicate an area of risk. This rule uses a unique yet simple algorithm to identify abnormal performance patterns.
Hot LUNs by Write Response	Reveals spikes in Write Response Time metrics, which may indicate an area of risk. This rule uses a unique yet simple algorithm to identify abnormal performance patterns.
Thin Pool Forecast	Uses multi-vendor and multi-metric pool capacity and forecast data to identify storage at risk.
Virtualization Rules	
VM Datastore Forecast	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.
VM Guest Disk Forecast	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.

Chapter 18

Manage and monitor data collection

This chapter includes the following topics:

- Data collection overview
- Data collection component configuration
- About data collection tasks
- Data collection installation summary
- Update the Local Hosts file for data collection
- Data Collector security and data encryption
- Manage Data Collectors and collection policies
- About validation and status
- Add/Edit Data Collectors
- Enable/Disable data collectors
- Enable and disable data collection policy schedules
- Review collectors and collection status
- Deleting a data collector
- Upgrade Data Collectors
- Monitoring data collection status
- Organize the collection status view

- Quick Filters
- View data collection status
- Troubleshoot data collection status
- Use reports to monitor data collection status

Data collection overview

The Data Collector is a centralized and remotely managed Java application responsible for interfacing with enterprise objects, such as backup servers and storage arrays, gathering information related to storage resource management.

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.

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 that are to be monitored and included in its data collection process.

Data collection component configuration

1. On the Portal Server:
   - Create a Data Collector in the Portal to enable the Portal server to receive data from the Data Collector server. In the Portal, you must first create a Data Collector and then populate it with product-specific or host data collection policies. A single Data Collector can be installed for multiple capacity, fabric, virtualization, File Analytics, and backup products. This Data Collector configuration in the Portal contains the configuration details for communicating with the corresponding Data Collector Server.
   - Add subsystem-specific Data Collector Policies. A Data Collector Policy provides the configuration details required to communicate with a subsystem to retrieve data that will be stored in the database. These details are specific to the vendor of the enterprise object from which data is collected. Policies also allow you to set the schedule for data collection. Prior to creating Data Collector Policies, a Portal Data Collector must be created.

2. On the Data Collector Server:
   - Add the Portal IP address to the Local Hosts file on the Data Collector server or on any available client with web-browsing capabilities.

     Note: Only edit the local hosts file if a DNS entry hasn't already been set up in your enterprise to resolve both http://aptareportal.yourdomain.com and http://itanalyticsagent.yourdomain.com to the Portal IP address.

   - Install the Data Collector software. This software component, installed on the Data Collector Server, interfaces with each of the supported subsystems to extract meta-data about the underlying environment. For example, backup data can include job details and tape inventory information. In the case of Capacity Manager, the Data Collector communicates with the storage arrays in your SAN (Storage Area Network) to collect meta-data.

About data collection tasks

A Data Collector regularly queries your enterprise objects for specific information, and each information type is called a collection task. Each collection task runs at specific intervals, and not all collection tasks run at the same intervals.

A collection task does not always return data because sometimes there isn't any data to return. However, when the collection task returns data, this historical information is used to determine the collection task's activity pattern or threshold.

Backup collection tasks

Most collection tasks run between every 20 minutes to 24 hours. However, one collection task, the Backup Job Completed Event, can post data several times a second at the height of the backup window, thereby setting the historical period for posting data to a very short interval. Subsequently, when the backup window is closed and no backups are being performed, the status monitoring might indicate an alert for this data collection component. If you have a backup window with heavy activity and then no or little activity, you may encounter some false positives for this component. If this component indicates it has not captured any data for more than 24 hours, then the component likely indicates an issue that requires investigation.

Backup event data collector

The Event Data Collector is the software component responsible for capturing backup event data. It is started for the following subsystems: Commvault Simpana,

EMC Data Domain, HP Data Protector, and Veritas Backup Exec. This event collection is logged to enable troubleshooting and isolation of collection issues, by processing thread. You can access the logs via the Support Tools utility: **Admin > Advanced > Support Tools**.

Data collection installation summary

1. Update the Local Hosts file on the Data Collector server or on any available client with web-browsing capabilities. This enables Portal access.
2. In the Portal, add a Data Collector, if one has not already been created.
3. In the Portal, add the vendor-specific data collector policy.
4. On the Data Collector Server, install the Data Collector software.
5. If collecting from Windows hosts, install the WMI Proxy Service on one of the Windows hosts.
6. Validate data collection is successful.

Update the Local Hosts file for data collection

1. Add the Portal IP Address to the Local Hosts file on the Data Collector server or on any available client with web-browsing capabilities.

Note: Only edit the local hosts file if a DNS entry hasn't already been set up in your enterprise to resolve both http://aptareportal.yourdomain.com and http://aptareagent.yourdomain.com to the Portal IP address.

Login to the Data Collector server.

On a Linux server:

```
vi /etc/hosts
```

Add entries for **itanalyticsportal.yourdomain.com** and aptare**agent.yourdomain.com**, both resolving to the Portal server IP address.

On a Windows server:

```
edit C:\Windows\System32\drivers\etc\hosts
```

Add entries for **itanalyticsportal.yourdomain.com** and **itanalyticsagent.yourdomain.com**, both resolving to the Portal server IP address.

2. On the Data Collector server, add entries to the local hosts file, both resolving to the Portal server IP address.

   Example:

   - 172.16.2.2 itanalyticsportal.<yourdomain>.com
   - 172.16.2.3 itanalyticsagent.<yourdomain>.com

Data Collector security and data encryption

Data collectors offer asymmetric encryption, also known as public-key cryptography. With this form of encryption, keys come in pairs - what a single key encrypts, only the other key can decrypt. This method of encryption provides additional security when data is collected.

In an upgrade scenario, you must enable asymmetric encryption for better security by generating a registration file. You can also choose to continue with the symmetric encryption method but it will be less secure than the asymmetric encryption. Registration file generation can occur at any time after an upgrade or if there is an issue such as data corruption or a key is lost.

To use this feature in either a new installation or an upgrade scenario, a registration file must be manually generated in the Portal. See "Data Collector encryption" on page 666.

Manage Data Collectors and collection policies

The **Collector Administration** view is a dashboard for data collectors and policies. In addition to the setup and management of Data Collectors and their collection policies, the **Collector Administration** view enables you to initiate, monitor, edit and validate the live status of your data collection all within the portal.

See "About validation and status" on page 658.

See "Add/Edit Data Collectors" on page 658.

See "Enable/Disable data collectors" on page 667.

See "Enable and disable data collection policy schedules" on page 668.

See "Review collectors and collection status" on page 669.

See "Deleting a data collector" on page 672.

See "Upgrade Data Collectors" on page 673.

See "Monitoring data collection status" on page 674.

About validation and status

In addition to the setup and management of Data Collectors and their collection policies, the **Collector Administration** view enables you to monitor if data collectors are running and the current status of collection. From this view, you have visibility into the status of both scheduled collection and on-demand runs, with drilldowns to more granular information available on the **Collection Status** view.

See "Monitoring data collection status" on page 674.

The **Collector Administration** view also provides functionality to quickly initiate and validate the collection run, once collectors and policies are setup. Validation methods are initiated differently based on the subsystem vendor associated with the Data Collector policy, but perform essentially the same functions.

You can also choose to view the collection logs on the portal while performing an on-demand run.

See "Review collectors and collection status" on page 669.

Add/Edit Data Collectors

To enable the Data Collector server to pass data to the Portal server, a corresponding Data Collector must be created in the Portal, along with Data Collector policies for each of the vendor-specific enterprise objects. Data Collector policies are specific to the type of data that is being collected; however, multiple types of policies often can be combined within one Data Collector.

The first step is to create a Data Collector. Once created, you can add policies to it. Often one Portal Data Collector is sufficient for adding Data Collector policies for a variety of enterprise objects such as backup servers, arrays, and switches.

See "Upgrade Data Collectors" on page 673.

To add a Data Collector

1. Select **Admin** > **Data Collection** > **Collector Administration**. The list of currently configured Portal Data Collectors is displayed. If a Data Collector has already been created, rather than creating a new Data Collector, you may want to add your collection policies to an existing Data Collector.
2. Click **Add Collector**.

   See "To Edit a Data Collector" on page 659.

To Edit a Data Collector

1. Search for Collector by name. Search returns results at the folder level and within the folder.

   See “Navigating with search” on page 35.

   Alternatively, select **Admin** > **Data Collection** > **Collector Administration** to browse for a collector. A list of currently configured Portal Data Collectors is displayed.
2. Select a Data Collector from the list.
3. Click **Edit**.
4. Enter or modify fields as necessary.
5. Click **Generate Registration File**.

Table 18-1 Field description

Field	Description
Collector Name*	The collector name cannot include a space and is case sensitive. The names should match exactly as entered in the Data Collector configuration screen and the Data Collector Installer screen. Edit the unique name assigned to this Data Collector. The Data Collector will use this value for authentication purposes. Changing the Collector ID or passcode requires manual changes to the corresponding Data Collector server. Collection will break if these corresponding changes are not made.
Passcode*	Edit the passcode assigned to this Data Collector. It can be any character sequence. Unlike other system passwords (which are encrypted and then saved) this Data Collector passcode is not encrypted prior to saving in the database and may appear as clear case in certain files. It simply is intended as a “handshake” identification between the Data Collector and the policy. Changing the Collector ID or passcode requires manual changes to the corresponding Data Collector server. Collection will break if these corresponding changes are not made. You can use the following OS-specific special characters in the passcode. Make sure the special characters you include are supported on the OS where the Data Collector is installed. You can specify from one to 128 characters in your passcode. ■ Linux: !@#%^* ■ Windows: !@#$%^&*()

Table 18-1 Field description *(continued)*

Field	Description
Short Notes	Descriptive notes associated with this Data Collector.
Enable SSL	Both secure (SSL) and non-secure Data Collectors can send data to the same Portal. Check this box to select the secure communication protocol (https) that the Data Collector will use. This check box will not appear in the dialog box if SSL is not enabled in your environment. The Portal data receiver must be listening for https traffic; for example: https://agent.mycollector.com
Auto-upgrade aptare.jar	Indicate if you want this configuration file upgraded automatically. This part of the Data Collector is responsible for event and metadata processing threads. The .jar file contains the processing and parsing logic for data collection. The latest versions can be downloaded automatically and applied to the collector during upgrades. It is recommended that this setting be set to Yes.
Auto-upgrade Upgrade Manager	Indicate if you want this configuration bundle upgraded automatically. This Data Collector component is responsible for managing Data Collector upgrades. The latest versions can be downloaded automatically and applied to the collector during upgrades. It is recommended that this setting be set to Yes.

Table 18-1 Field description *(continued)*

Field	Description
Generate Registration File	

Table 18-1 Field description *(continued)*

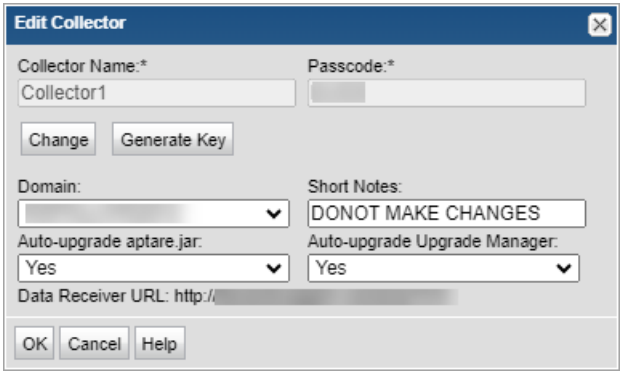
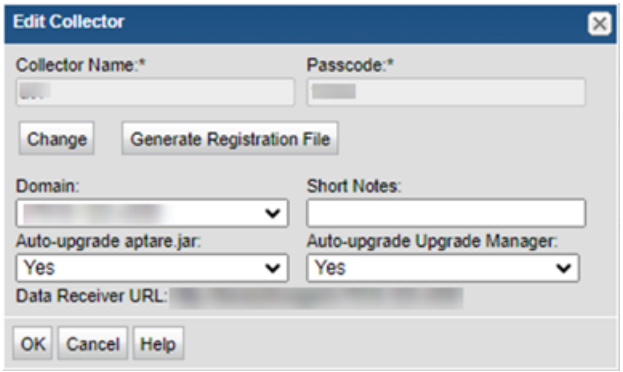
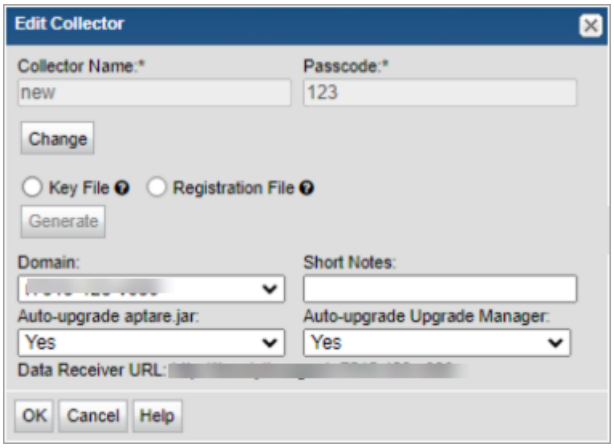
Field	Description
	This creates a registration file required for asymmetric encryption, also known as public-key cryptography. This provides additional security when data is collected. Click to generate the registration file and download it to your local system. This is required for new installations and optional for upgrade scenarios. When generating a new registration file, for example if a registration file is lost or data has been corrupted the following applies: ■ Generating a registration file disables all Data Collector policies associated with the collector and invalidates their credentials. ■ For each policy associated with the collector, you must re-enter all credentials and re-enable the policy. See "Data Collector encryption" on page 666. The **Edit Collector** pop-up displays **Generate Registration File**, **Generate Key**, or both depending on the collector configuration. The conditions when the **Edit Collector** pop-up will display **Generate Registration File**, **Generate Key**, or both are described below: ■ When the Data Collector version configured on the Portal is less than or equal to 11.2.04, the **Edit Collector** pop-up will display **Generate Key**. This will generate the `.key` file when clicked.  ■ When the Data Collector version configured on the Portal is greater than 11.2.04, the **Edit Collector** pop-up will display **Generate Registration File**. This will generate the `.json` file.

Table 18-1 Field description *(continued)*

Field	Description
	 ■ When the Data Collector is not configured on the Portal and its version is unknown, the **Edit Collector** pop-up will display **Key File** and **Registration File** options.  You will observe this typically when a new Data Collector is added on the Portal. Select the **Key File** if the Data Collector installer version is 11.2.xx or lower. Select **Registration File** when the Collector installer version is 11.3.xx or higher, or when registration needs to be performed using the NBU Web UI. **Note:** See Generate key file to configure Data Collector v11.2.xx or lower for steps to generate key file from a newly created Data Collector.

Edit IT Analytics Data Collector credentials

The following table provides the steps you need to follow when you change the Data Collector credentials, such as Collector Name and Passcode. These steps are essential to establish the Data Collector connectivity with theIT Analytics Portal and seamless data collection.

Follow the steps provided in the table after you:

1. Search for Collector by name on the Portal. Search returns results at the folder level and within the folder.

   See "Navigating with search" on page 35.

   Alternatively, select **Admin** > **Data Collection** > **Collector Administration** to browse for a collector. A list of currently configured Portal Data Collectors is displayed.
2. Select a Data Collector from the list.
3. Click **Edit**.
4. Click **Change** on the **Edit Collector** pop-up.

   A warning message is displayed that indicates the impact of making edits to the Data Collector name and passcode. Click **OK** and proceed to the next steps.

Table 18-2 Edit IT Analytics Data Collector credentials

Task	Procedure	Remarks
Change only **Passcode**	**1** Change the **Passcode** value. **2** Click **OK** to save the changes. **3** Update the new passcode value in `mbs/conf/collector.properties` file in plain text on the Data Collector server. **4** Run `updateconfig.bat` **5** Start the Aptare Agent.	This does not require generating a new registration file.

Table 18-2 Edit IT Analytics Data Collector credentials *(continued)*

Task	Procedure	Remarks
Change only **Collector Name**	**1** Change the **Collector Name** value. **2** Click **Generate Registration File** and note the `.json` file path. **3** Click **OK** to save the changes. A new registration file (.json file) gets downloaded. **4** Run `reconfigureDC.bat/sh <.json file path>` on the Data Collector server.	This change generates and downloads a new registration file.
Change both **Collector Name** and **Passcode**	**1** Change the **Collector Name** and **Passcode** values. **2** Click **Generate Registration File** and note the `.json` file path. **3** Click **OK** to save the changes. A new registration file (.json file) gets downloaded. **4** Run `reconfigureDC.bat/sh <.json file path>` on the Data Collector server.	This change generates and downloads a new registration file.

Generate key file to configure Data Collector v11.2.xx or lower

When you add a new Data Collector on the IT Analytics Portal, a Registration File is generated and downloaded by default. However, if you have a Data Collector of version 11.2.xx or lower, you require a key file to configure it with the Portal. This section describes how you can download a key file from the IT Analytics Portal to be able to install and configure a Data Collector of version 11.2.xx or lower.

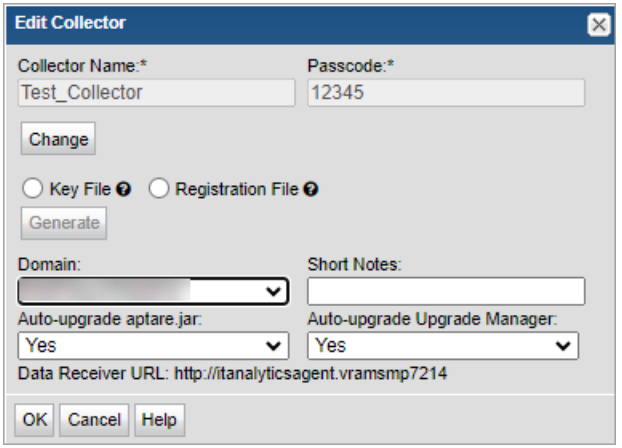
You must have a Data Collector entry already added on the Portal when you follow these steps. See “Add/Edit Data Collectors” on page 658.

To get a key file:

1. Log on to the IT Analytics Portal.
2. Go to **Admin** > **Data Collection** > **Collector Administration** to browse for a collector. A list of currently configured Data Collectors is displayed.
3. Select the Data Collector from the list. You can also use the **Search** to find the collector.

4 Click **Edit**.

The **Edit Collector** pop-up displays **Key File** and **Registration File** options.



5 Select **Key File** and click **Generate**.

A key file is downloaded locally.

6 Click **OK**.

Data Collector encryption

For new Data Collectors, asymmetric encryption requires some initial setup. When you add a Data Collector in the Portal, you download the registration file and then point to that location when you install the Data Collector software on the collector server.

For existing Data Collectors, registration file generation for asymmetric encryption can occur at any time. You can opt-in to encrypt/decrypt credentials.

Activate encryption for new Data Collectors

For a new installation, the registration file must first be generated in the Portal. Next, the registration file location must be entered as part of the Data Collector installation process.

To activate encryption for new Data Collectors

1 Navigate to **Admin** > **Data Collection** > **Collector Administration**.

2 Click **Add Collector**.

3 Click **Generate Registration File**.

4 Click **OK** to proceed and download the registration file to your local system.

5 Copy the `<collectorname>.json` file to a temporary location. You will be prompted for this location during the Data Collector installation.

6 Install the Data Collector software on the Data Collector server.

Activating Encryption for Existing Collectors

In an upgrade scenario, you can change the encryption method to asymmetric and add the extra layer of security to active Data Collector policies.

To activate encryption for existing Data Collectors

1 Stop the Data Collector.

2 In the Portal, search for Collector by name. Search returns results at the folder level and within the folder.

See "Navigating with search" on page 35.

Alternatively, select **Admin** > **Data Collection** > **Collector Administration** to browse for a collector. A list of currently configured Portal Data Collectors is displayed.

3 Select a Data Collector from the list.

4 Click **Edit**.

5 Click **Generate Registration File**.

Note: When generating a replacement registration file, for example a registration file is lost or data has been corrupted the following applies: active policies associated with the collector are disabled and their credentials are invalidated. For each policy associated with the collector, you must re-enter all credentials and re-enable the policy.

---

6 Click **OK** to proceed and download the registration file to your local system.

7 Copy the `<collectorname>.json` file to a temporary location on the collector server.

8 Execute `reconfigureDC.bat/sh <registration file path>`.

Enable/Disable data collectors

You can enable and disable a Data Collector through the Portal. When you disable a collector, it stops the service/process running on the data collector server and stops all collection. When collectors are disabled, you can still edit schedules and

parameters. To start collection after a Data Collector has been disabled, you must enable the Collector in the Portal, and start the Data Collection service/process on the Data Collector server.

Note: When you disable the Data Collector, all policies associated with it are suspended. When you enable the Data Collector, the policies resume their initial status.

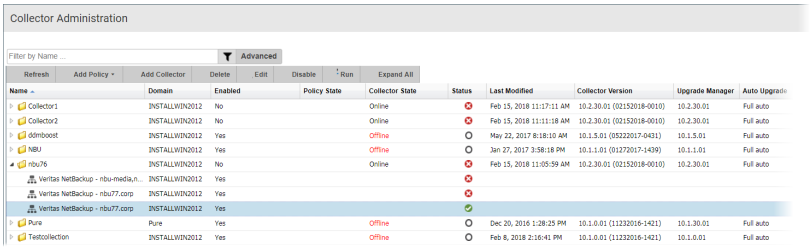
See "Enable and disable data collection policy schedules" on page 668.

To enable/disable a data collector

1. Search for a data collector by name.

   See "Navigating with search" on page 35.

   Alternatively, you can browse for a collector. Select **Admin > Data Collection > Collector Administration**. A list of Data Collectors is displayed.



2. Select the Data Collector.
3. Click **Disable** or **Enable**. When you disable a Data Collector all policies within the folder are also deactivated. You can also keep a collector enabled and disable individual policies.

Enable and disable data collection policy schedules

Enable and disable individual data collection policy schedules through the Portal. When you disable a policy, it deactivates the collection schedule. When policies are disabled, you can still edit them.

Note: When you disable a Data Collector all policies associated with it are suspended. When you enable the Data Collector, the policies resume their initial status.

See "Enable/Disable data collectors" on page 667.

To enable/disable a data collector policy

1. Search for a data collector policy by name.

   See "Navigating with search" on page 35.

   Alternatively, you can browse for a policy. Select **Admin > Data Collection > Collector Administration**. A list of Data Collectors is displayed.
2. Select a Data Collector.
3. Open the folder and select a policy.
4. Click **Enable** or **Disable**.

Review collectors and collection status

The status of your Data collection is comprised of multiple conditions, the state of the Collector, is it online or offline, are policy schedules active or stopped, is a scheduled collection run in progress or is an On-Demand run waiting? All of these conditions and others make up the current status of the Data Collector. The **Collector Administration** view provides this summary-level status. From this view, you have visibility into the status of both scheduled collection and on-demand runs, with drilldowns to more granular information available on the **Collection Status** view.

Status for collectors and policies are an aggregate value. Any collection failure on a probe, results in a Failure status for the policy, and any failure on a policy results in a Failure status for the Data Collector. When determining a status for a data collector policy, the status of individual probes are rolled up to create a value for the last time collection was attempted. For example, if collection has failed for a certain enabled probe, the status of the policy would be failed.

Known Limitations for Policy States in Collector Administration

Due to the nature of certain types of collected data, the Collector Administration view cannot always display states for all collected policies.

The following table lists the exceptions, where the collected subsystems cannot be fully represented in the view.

Collected Subsystem	Licensed Module	Policy State
- CIFS		No explicit start and finish transactions are collected, so only the Policy State displayed is Collecting or No Status.
Host Resources		The Policy State displayed is No Status.

To view collector status

1. Search for a data collector by name.

   See “Navigating with search” on page 35.

   Alternatively, you can browse for a collector. Select **Admin > Data Collection > Collector Administration**. A list of Data Collectors is displayed.

2. Select the Collector.
3. View the **Collector State** and the **Status**.
4. Click the **Status** icon to drill down to the **Collection Status** page. The information displayed on the **Collection Status** page is filtered by the Data Collector you are viewing.

Column Title	Description
Collector State	Indicates if the Collector services are running and if the handshake has occurred. Values ■ Online ■ Offline - Indicates the collector is shut down, or the collector cannot connect to the Data Receiver. To restart the Collector, it must be **Enabled** through the Portal and the Collector service must be manually started on the Data Collector server. See “Enable/Disable data collectors” on page 667.

Column Title	Description
Status	Indicates an aggregated status of the last run for each probe for a scheduled collection on all associated enabled policies. Click the **Status** icon to drill down to the **Collection Status** page for a more granular view of the status information. See “Monitoring data collection status” on page 674. Values ■ Success (green check mark) - All probes are green. ■ Warning (yellow triangle) - At least one probe is yellow. ■ Failure (red X) - Action required. If any scheduled probe run has failed, the status is red. ■ Unknown (white circle) - No status available or the collector has been turned off.

To view the policy status

1 Search for a data collector policy by name.

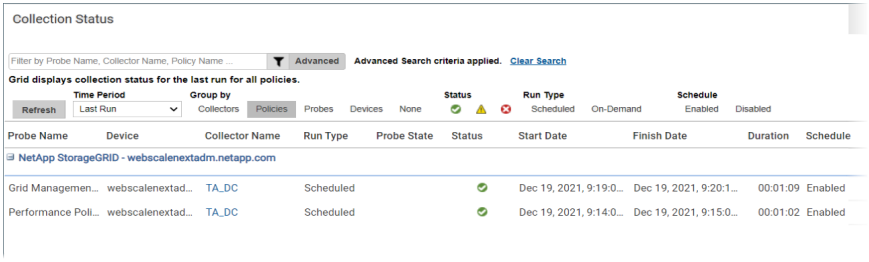
See “Navigating with search” on page 35.

Alternatively, you can browse for a policy. Select **Admin > Data Collection > Collector Administration**. A list of Data Collectors is displayed.

2 Select the Collector and click the expand icon to view the associated policies.

3 View the **Enabled, Policy State,** and the **Status**.

4 Click the **Status** icon to drill down to the **Collection Status** page. The information displayed on the **Collection Status** page is filtered by the Policy you are viewing.



Column Title	Description
Enabled	Indicates if the policy schedule is enabled or if it has been manually disabled. Values ■ Yes ■ No - No scheduled probes will run, however On-Demand runs can be started.
Policy State	Indicates the collection policy state. This can be either scheduled or on-demand. Values ■ Waiting - Collection has been initiated, but not started. If a collection is on a schedule, an on-demand run will wait for the scheduled run to complete before starting. ■ Collecting - Collection has been initiated and in-progress. **Note:** Policy State does not automatically refresh. Click **Refresh** to update the display in the grid.
Status	Indicates an aggregated status of the last set of enabled probes run for a scheduled or on-demand collection. Click the Status icon to drill down to the **Collection Status** page for a more granular view of the status information. See "Monitoring data collection status" on page 674. Values ■ Success - All probes are green. ■ Warning - At least one probe is yellow. ■ Failure - Action required. If any scheduled probe run has failed, the status is red. ■ Unknown - No status available or the collector has been turned off.

Deleting a data collector

When you delete a Data Collector from the Portal, the collector and any policies associated with it will be deleted from the Portal and the database.

To delete a data collector

1 Search for a data collector if required.

See “Navigating with search” on page 35.

Alternatively, you can browse for a collector: **Admin > Data Collection > Collector Administration**. A list of collectors is displayed.

2 Select a collector and click **Delete**. You are prompted to confirm the deletion.

See “Enable/Disable data collectors” on page 667.

Upgrade Data Collectors

The Data Collector Upgrader provides options to manually upgrade to the latest Data Collector logic. Updates can be downloaded to the Portal server. From the Portal server, you can push the updates out to individual Data Collector servers.

See “To deploy updates to collectors” on page 673..

Upgrade aptare.jar

This part of the Data Collector is responsible for event and metadata processing threads.

Update Upgrade Manager

The Upgrade Manager is responsible for all Data Collector upgrade activities.

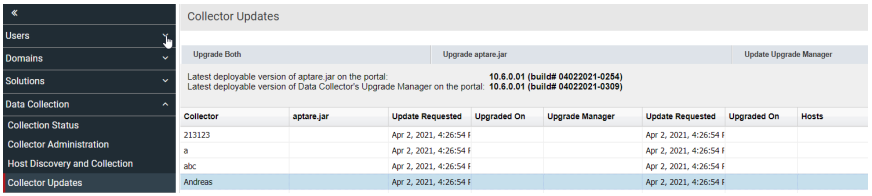
To deploy updates to collectors

- Select **Admin > Data Collection > Collector Updates**.

  Collectors that are not running the current version have the version number displayed in red.

  To download the latest version, click one of the following:

  - Upgrade Both
  - Upgrade aptare.jar
  - Update Upgrade Manager



Monitoring data collection status

Use the **Collection Status** page to monitor the health and progress of data collection. This view also contains probe runs and can be organized to suit your business requirements providing essential details that enable you to diagnose collection issues. Collection status, available at the data collector and policy level, provides results for the last time collection was attempted for enabled probes.

Organize the collection status view

Customize the view into your **Collection Status** by grouping the information in a way that is the most relevant to your business. This allows you to see what you need, when you need it, so you can efficiently analyze the data.

By default, the grid displays information for the last set of collection runs and organizes it by probe.

Use the **Advanced** Filter to build a more granular level of filtering on the viewing grid. You can create queries using all relevant fields in the database with specific operators to further refine and locate the information you need.

See “Advanced filtering” on page 40.

Note: When accessing the **Collection Status** view by drilling down from a status on the **Collectors** view, the information is already filtered based on your original selection. That is, if you drill down from a policy status, the view is filtered by the associated policy. You must **Clear the Filter** to view the entire set of information.

Quick Filters

Use Quick Filters to sort the collection status without building a query with the **Advanced** Filter. Choose from:

- Time Period - Select the data collection run time period. You can select the **Last Run** (the default) or **Last 24 hours**.

- Group by - Select how the data collection status information is grouped. You can choose from: **Probes, Policies, Collectors, Devices** or **None**. If **None** is selected, the grouping is by **Probe Name**.
- Status - Select to filter data collection information by **Status**. Choose from Success, Warning or Failure. Click the status icon in the grid to display details.
- Run Type - Select to display status information based on the type of collection run. Choose from a **Scheduled** or **On-Demand** run. All collection run types are shown by default. Note, **On-Demand** runs are not available for all vendors.
- Schedule - Select to include data collection status information for those probes, policies and collectors that have schedules **Enabled** or **Disabled**.

View data collection status

Collection status, available at the data collector, policy and probe level, provides results for the last time collection was attempted. Collection is schedule-based and can also be run on-demand for certain vendors. When policies are created, a schedule is set, but collection may not enabled. Because the status is relevant to the conditions you are troubleshooting, you can view a status if the schedule is enabled or disabled.

To View Data Collection Status

1. Navigate to **Admin>Data Collection>Collection Status**. By default, the grid is organized by **Probe Name**.
2. Organize the grid to your preference. You can use **Quick Filters** to efficiently sort the data with pre-defined options, or use the **Advanced** Filter to build a more granular query.

   See “Organize the collection status view” on page 674.
3. Evaluate the collection **Status**. Click the status icon to display details.
   - Success - Data collection has successfully completed. The last run of each probe in a policy, including an On-Demand run is successful.
   - Warning - Data collection has completed with some errors. The last run of each probe in a policy, including an On-Demand run is complete and any probe is:
     - Successful and there are error messages in the collector log
     - Canceled or skipped. If an On-Demand collection is running, the scheduled run is skipped.
   - Failure - Data collection has failed to complete. The last run of each probe in a policy, including an On-Demand run is complete and any probe is:

- Interrupted which indicates that collection was running and stopped during the run. For example, when services are stopped or the receiver monitor thread detects that a Collector is not running, any On-Demand Runs or Scheduled collections In Progress are marked as Interrupted. For scheduled runs, restarting the services, marks the previous In Progress collection as Interrupted as well. For On-Demand runs, restarting the services, results in the On-Demand run starting again.
- Canceled with an error for any probe in the policy (only applicable for On-Demand runs)
- Canceled with success for any probe in the policy and fatal error messages in the collector log (only applicable for On-Demand runs)

Troubleshoot data collection status

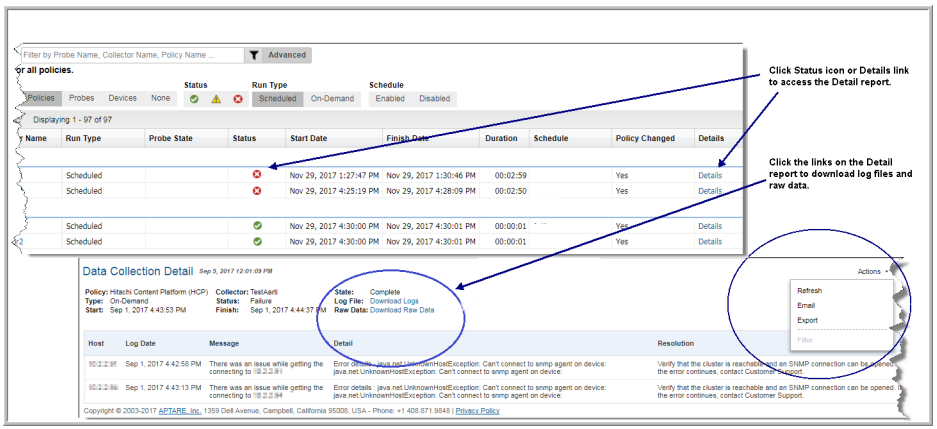
Troubleshoot issues by drilling into the **Data Collection Detail** report. This detail report can also be displayed by clicking the **Status** icon. From this report, you can view specifics about status issues as well as possible resolutions. This information can be emailed and exported.

See “Email reports and dashboards” on page 177.

See “Export reports and dashboards” on page 169.

For On-Demand runs, you can also:

- Download log files
- Download raw data



For additional information, use **Support Tools** to download and inspect data collection logs: **Admin >Advanced > Support Tools**. Refer to the following for a description of the log file naming convention.

Use reports to monitor data collection status

Use the following System Administration reports to use in conjunction with the Data Collection Status page to monitor the status of Data Collection.

- Data Collection Message Summary
- Data Collection Activity Detail
- Data Collection Hourly Activity
- Data Collection Performance Detail
- Data Collection Performance Summary
- Data Collection Schedule Summary
- Data Collector Status Summary
- File Analytics Collection Status

Chapter 19

Work with Capacity Manager host data collection

This chapter includes the following topics:

- Understand the host data collection process
- Host resources prerequisites and configurations
- Host Access Privileges, Sudo Commands, Ports, and WMI Proxy Requirements
- Host access requirements
- Command path verification
- Host discovery and collection configuration steps
- Host discovery and collection setup overview
- Host discovery and collection maintenance overview
- Before discovering hosts
- Configure/Search
- Manage credentials
- Manage WMI Proxy
- Manage paths
- Manage access control
- Host management

- Configure host discovery policies to populate the host discovery and collection view
- Execute and monitor host discovery
- Discovery processes
- Validate host connectivity
- Show errors
- Filter the host discovery and collection window - Hide/Unhide, remove
- Search and export in host discovery and collection
- Export in host discovery and collection
- Configure and edit host probes
- Propagate probe settings: Copy probes, paste probes

Understand the host data collection process

Capacity Managercan collect data and then report on storage that is allocated to and consumed by hosts in your enterprise. Host capacity and utilization reports enable you to optimize existing storage resources and more accurately forecast usage.

Host Resource Data Collection gathers information by probing hosts:

- Host Probes: Capacity (HBA, iSCSI, Volume Manager, Multi-pathing)
- Host Probes: Memory, Network, Process, Performance, System
- Application Probes: Exchange, SQL Server, Oracle, Oracle ASM
- Probes

IMPORTANT: Host Resources data collection does not require a dedicated Data Collector for each resource. If you have a Storage Array Data Collector, the Host Resources collector is inherently part of that Data Collector. However, if for some reason you do not have a Storage Array Data Collector, you can explicitly create just a Host Resources Data Collector.

Several key steps comprise the Host Data Collection Process. These steps are summarized here, with details provided in the descriptions of specific tasks.

- Add Hosts to the Host Inventory - This initial setup phase requires some pre-planning to ascertain which hosts and credentials will be needed for successful host authentication. Then, you'll take this information and create

several configuration settings--credentials, WMI proxies, paths, and access control commands--required to discover hosts in your environment. The Host Discovery process attempts to find hosts using these configuration settings and then populates the host inventory.

- Configure & Validate Hosts - Once hosts have been added to the inventory, specific probe settings can be configured to tailor the type of data to be collected from a host. The Validate step provides feedback to troubleshoot host connectivity and data collection issues. In addition, you can hide/remove hosts that do not belong in your inventory--for example, IP addresses of non-host devices such as tape drives. This is an iterative process to verify the collection settings for each host in your host inventory. Note that for the File Analytics probe, by design, the Validate option only runs a connectivity check; it does not collect File Analytics data.
- Enable & Manage On-going Collection - Once a host has been validated, enable on-going data collection. Subsequent changes to the host in your enterprise may impact data collection. As changes and collection issues arise, updates to host data collection configurations will be required.

Host resources prerequisites and configurations

Prior to configuring the system to discover your host inventory, you must identify the hosts for which you will be collecting data.

See "Understand the host data collection process" on page 679.

See "Before discovering hosts" on page 687.

Host Access Privileges, Sudo Commands, Ports, and WMI Proxy Requirements

If you are using sudo to elevate access to root privileges, update the sudoers file:

- Sudoers file: /etc/sudoers
- Use the lists of the sudo commands (per OS) that are located on the Portal server in:

  ```
  <Home>/opt/aptare/updates
  ```

- Comment out this line in the sudoers file: **Defaults requiretty**

Access Requirements by OS

Table 19-1 Host Resources Prerequisites by Operating System

Host OS	Host Access Requirements	Port Requirements	Notes
Linux RH Linux SUSE CentOS AIX Solaris	ssh or telnet must be enabled Some commands may require an account with super-user root privileges. **sudo**, **sesudo**, and **pbrun** are supported; ensure the user ID has required sudo, sesudo, or pbrun privileges.	ssh: 22 telnet: 23	Collection uses ssh/telnet to execute commands. OS and application commands require root privileges for HBA API access. The **sysstat** utility must be installed on Linux servers or storage nodes for Linux host performance data collection.
Windows	A WMI Proxy is required to collect from Windows hosts. All Windows hosts require Local or System Administrator privileges for WMI.	RPC: TCP Port 135 for WMI DCOM: TCP/UDP 1024-65535 TCP/IP 1248, if WMI Proxy server is not the same as the Data Collector server	When the Data Collector Policy is configured to include file-level data, the Data Collector and WMI need to use a Windows Domain Administrator ID.

Host access requirements

This section lists the access requirements for host resource data collection. You will use this information to populate the configurations used by the Host Discovery, Validation, and Collection processes.

- User ID & Password Credentials: Root-level, read-only access is required for host data collection.
  See “Manage credentials” on page 688.
  See “Manage access control” on page 696.
- Access Control: For security reasons, most enterprise environments mandate access control where a new non-root account is created, with temporarily elevated access to the required commands provided via an access control command, such as sudo. Otherwise, the root user is required for host access.
  See “Manage access control” on page 696.
  Files containing sudo commands per operating system can be found on the Portal server in: /opt/aptare/updates. These filenames contain both the OS and the version of the sudo commands file so that you can identify the files that

contain the latest updates; for example: hpux_9.1.01, aix_9.1.01, linux_9.1.01, solaris_9.1.01.

- Path: The system must have knowledge of the correct paths to access commands. An overview of the requirements is listed here, with the details for determining paths provided in Command Path Verification.
  See "Command path verification" on page 683.
  For Windows hosts, a path is required for fcinfo, hbacmd, and scli commands. For Linux hosts, if the Data Collector is installed on a Windows server, use plink.exe to determine the path; if the Data Collector is installed on a Linux server, determine the path by executing ssh.
  See "Command path verification" on page 683.
- **HBA Prerequisites**
- It is critical for the Data Collector to probe the HBA in order to establish a host's relationship with storage. Without the HBA information, all storage for a host will be listed as local storage. An internal probing mechanism is used to gather Host Bus Adapter (HBA) data from Windows hosts.
  - **Windows**: Either hbaverify, scli, hbacmd (required for both LUN Mapping and HBA data collection), or fcinfo.
  - **Linux**: scli or hbacmd (required only for HBA information)
  - **Solaris**: scli or hbacmd (required only for HBA information)

For Linux Hosts Only

For Linux Hosts in access control environments (such as sudo):

- If a command such as sudo is used and the path is not in the interactive ssh, identify the absolute path of the access control command.
  See "Command path verification" on page 683.

For Windows Hosts Only

- A WMI Proxy server is required for collecting data from Windows hosts. It is critical for the Data Collector to gather this data in order to establish a host's relationship with storage. Without this information, all storage for a Windows host will be listed as local storage.
- Microsoft Exchange 2003: The Data Collector uses WMI for data collection.
- Microsoft Exchange 2007 and 2010: Data collection requires PowerShell remoting to be enabled on the Exchange server. The Data Collector connects to PowerShell via the WMI Proxy to execute the PowerShell commands. For details

on remoting, see the Microsoft Administrator's Guide to Windows PowerShell Remoting.

- Verify the method of collecting Windows HBA information. Windows Hosts require one of the following to determine SAN information:
  - HBAnyware from Emulex
  - SANsurfer Command Line Interface (SCLI) for Windows from QLogic (SCLI is a separate install from the base install of SANsurfer and often is not installed with the SANsurfer utility).
  - Fibre Channel Information Tool (fcinfo) from Microsoft

Command path verification

Verify the command paths that will be used by the Data Collector.

Both Linux & Windows:

- If Volume Manager is installed on any hosts, note the path to the vxprint command.
- If any multi-pathing software is installed on hosts, note the path to the command.

Linux: Verify the non-interactive SSH path for Linux users for several sample hosts:

```
ssh <user>@<hostname> env
```

where **<user>** is the credential the collector will use to access the host.

To determine the Linux path from a Windows server, you can use a command-line interface to telnet/ssh client software. The following **example** shows Plink, which is a command-line interface to PuTTY (a telnet/ssh client):

```
plink <user>@<hostname> env
```

Example of a PATH for commands:

```
/bin:/sbin:/usr/bin:/usr/sbin:/usr/local/bin:/usr/local/sbin
```

Windows: Make a note of the paths for the executables identified for HBA data collection. Note that in Windows, multiple paths are separated by a semi-colon (;). For example:

```
C:\Program Files\Emulex\Util\HBAnyware;C:\Program
Files\QLogic\SANSurfer
```

Host discovery and collection configuration steps

This section contains:

- See "Host discovery and collection setup overview" on page 684.
  - Collector pre-requisite, Manage Credentials, Manage WMI Proxy, Manage Paths, Manage Access Control, Discover Hosts
- See "Host discovery and collection maintenance overview" on page 686.
  - Hide, Remove, Show Errors, Validate, Edit Probes, Copy Probes, Paste Probes

Read this section first for a high-level overview, then follow the specific steps in the next sections to configure your system for Host Data Collection.

See "Before discovering hosts" on page 687.

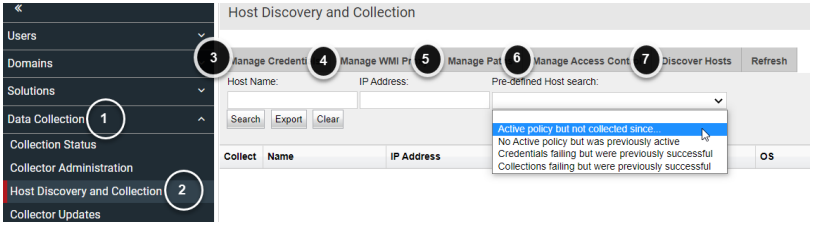
See "Host management" on page 700.

Because each enterprise has a unique inventory of hosts with specific access requirements and restrictions, the process of ensuring successful host data collection requires an assessment of the hosts in your environment (see) and several configuration steps, as summarized below.

Host discovery and collection setup overview

Note that some steps typically are required only as part of the initial configuration and rarely require additional maintenance. For the purpose of this document, these initial steps are treated as requirements.

See "Before discovering hosts" on page 687.



Note: Each step is summarized in this section. To access detailed descriptions, click the links for each step. The buttons at the bottom of the window--**Hide, Remove, Show Errors, Validate, Edit Probes, Copy Probes, Paste Probes**--are described in the section,

---

See "Host discovery and collection maintenance overview" on page 686.

1. Prior to discovering hosts, a data collector policy must be configured. You can use an existing policy--for example, a data collector policy that has been created for Storage Array data collection--or create a new data collector policy.
2. Using the Host Discovery and Collection window, you can search for hosts in the inventory; or you can set up configurations in preparation for discovering and configuring hosts. Many of the probes may not be applicable to your enterprise. It is essential that you identify the probes that are relevant to your hosts.

   See "Configure/Search" on page 687.

   See "Configure and edit host probes" on page 714.
3. Configure user IDs and passwords for authentication when the data collector is accessing hosts.
4. A WMI Proxy is required to collect data from Windows hosts. Use this option to define one or more WMI Proxies.

   See "Manage WMI Proxy" on page 691.
5. - Configure the paths that data collectors will use to execute commands on hosts.

   See "Manage paths" on page 694.
6. Data Collectors require read-only access to execute non-intrusive commands on hosts. It is strongly recommended that a separate login account, used strictly for IT Analytics, be established and using Active Directory for Windows systems and the sudo command for Linux systems, restrict the commands that IT Analytics can issue. To accommodate this security approach, you can optionally specify access control commands like sudo, sesudo, or pbrun. Files containing sudo commands per operating system can be found on the Portal server in: /opt/aptare/updates. These filenames contain both the OS and the version of the sudo commands file so that you can identify the files that contain the latest updates; for example: hpux_9.1.01, aix_9.1.01, linux_9.1.01, solaris_9.1.01.

   See "Manage access control" on page 696.
7. Host Discovery attempts to find hosts and populate your host inventory. Create Host Discovery Policies that use the credentials, WMI proxies, and paths that you configured.

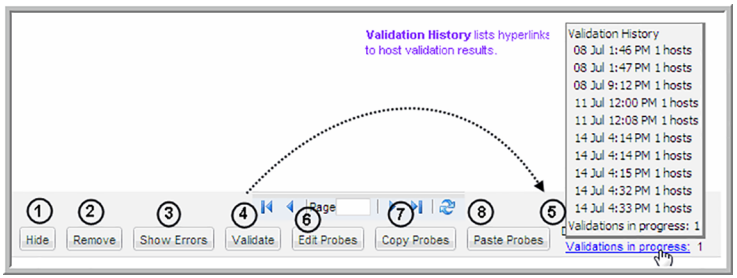
   Host validation must take into account host access for a wide variety of conditions and environments. As the discovery process accesses hosts, informative messages will provide clues to connectivity issues. In addition, devices that don't belong in a host inventory--for example, printers in the IP address range that you specified--may have been discovered and need to be hidden or removed from the inventory.

See “Configure host discovery policies to populate the host discovery and collection view” on page 701.

See “Validate host connectivity” on page 708.

Host discovery and collection maintenance overview

Once hosts have been discovered and they are listed in the Host Inventory, several options are provided to filter the list and also to manage the probes.



1. **Hide/Unhide Hosts**- Host Discovery may find devices that are not hosts that you want to manage; for example, printers.

   See “Filter the host discovery and collection window - Hide/Unhide, remove” on page 711.

2. **Remove Hosts**- Some IP addresses may be associated with devices that simply should be removed from the inventory, although if you execute a host discovery policy, the devices will return. For details,

   See “Filter the host discovery and collection window - Hide/Unhide, remove” on page 711.

3. **Show Errors** - Use this feature to troubleshoot connectivity and validation issues.

   See “Show errors” on page 710.

4. **Validate** - Use this feature in combination with the Show Errors feature to troubleshoot host data collection issues.

   See “Validate host connectivity” on page 708.

5. **Show Validations -**See “Validation history” on page 709.

6. **Edit Probes** -

   See “Configure and edit host probes” on page 714.

7. **Copy Probes** -

   See “Propagate probe settings: Copy probes, paste probes” on page 717.

8. **Paste Probes** -

   See “Propagate probe settings: Copy probes, paste probes” on page 717.

Before discovering hosts

If this is the first time you are collecting data from hosts, you will need to look at each of these steps to determine what configurations are required.

If you upgraded and you already had host collection policies in a previous version, the hosts, credentials, and access controls will have been converted and can be updated using the Host Inventory tools/features.

Before collecting host data for the first time, several configurations must be set up:

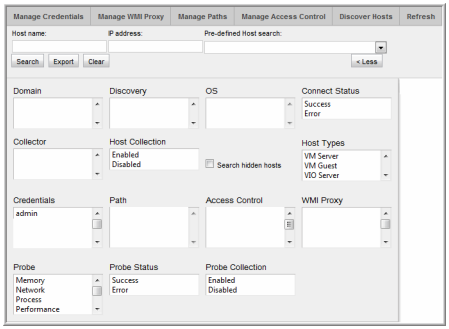
- See “Configure/Search” on page 687.
- See “Manage credentials” on page 688.
- See “Manage WMI Proxy” on page 691.
- See “Manage paths” on page 694.
- See “Manage access control” on page 696.

Configure/Search

Before Host Discovery: Use the Host Discovery and Collection window to set up configurations--credentials, WMI proxy, paths, and access control--as described in the following sections.

After Host Discovery: Use the Host Discovery and Collection window to help you find hosts in your inventory and configure probes. Also, export the list of hosts to a comma-separated-values (.csv) file.

See “Search and export in host discovery and collection” on page 712.



Note: A search with no specified criteria returns all hosts in your inventory.

Manage credentials

Multiple credential sets can be created, typically for groups of hosts with common credentials and/or hosts grouped by operating system (Linux/Windows). These credential sets are then selected and applied to specific Host Discovery policies. In fact, multiple credential sets can be listed, allowing the Data Collector to attempt authentication in a specific order until it is successful.

At the very least, you should have one credential set for Linux hosts and another for Windows hosts. Each defined set of credentials will have a name, to enable relevant selection when configuring Host Discovery policies.

See “Host resources prerequisites and configurations” on page 680.

See “Host discovery and collection configuration steps” on page 684.

To Manage Host Credentials, in the toolbar select:

1. **Admin > Data Collection> Host Discovery and Collection**.
2. In the Host Discovery and Collection action bar click **Manage Credentials**.
3. Add, Edit, or Delete credentials using the buttons at the bottom of the window.

Example of credentials for Windows hosts

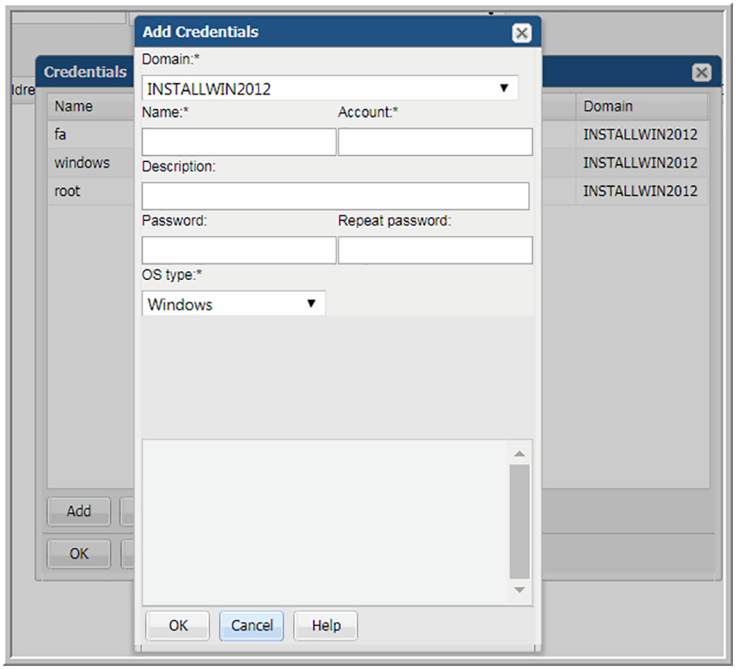


Table 19-2 Credentials for Windows Hosts

Field	Description	Sample Values
Domain*	Select the Domain from the list; for most environments, only one Domain is displayed. Multiple domains facilitate management for Managed Services Partners (MSPs).	
Name*	Assign a name to identify this set of credentials that you are defining.	

Table 19-2 Credentials for Windows Hosts *(continued)*

Field	Description	Sample Values
Account*	Enter the login account name used to log in to the hosts. If the policy includes a group of Windows hosts, use the Windows domain user id. This user id must have administrative privileges. For Linux hosts, super user root privileges are required. You also could use an access control command, such as **sudo, sesudo,**or **pbrun**. If using any of these access commands, ensure that the user ID has sudo, sesudo, or pbrun privileges. Some enterprises prefer to create a new user and provide access to commands via an access control command. Files containing sudo commands per operating system can be found on the Portal server in: /opt/aptare/updates. See "Manage access control" on page 696.	root
Description	Enter a note to help identify this type of credential	Linux logins for Corporate
Password	Enter the password for the account	Password1
OS type*	Select either Linux, Windows, or NAS.	
Windows Domain	**For Windows hosts only**: If any of the hosts specified in the Host address field are Windows hosts, you need to specify the Windows domain name. If the host is not a member of a domain, or to specify a local user account, use a period (.) to substitute the local host SSID for the domain.	win2kdomain
Private Key File	**For Linux hosts only**: If you have configured Public Key/Private Key between your Data Collector server and the Hosts you intend to monitor, use this field to specify the location of the Private Key file on the Data Collector server.	`/root/.ssh/id_rsa` or `C:\Program Files\Aptare\mbs\conf\id_rsa`

Table 19-2 Credentials for Windows Hosts *(continued)*

Field	Description	Sample Values
Known Hosts File	**For Linux hosts only**: If you have configured Public Key/Private Key between your Data Collector server and the Hosts you intend to monitor, use this field to specify the location of the Known Hosts file on the Data Collector server.	`/root/.ssh/known_hosts` or `C:\Program Files\Aptare\mbs\conf\known_hosts`

Manage WMI Proxy

Note: A WMI Proxy configuration is needed only if you are collecting data from Windows servers in your environment.

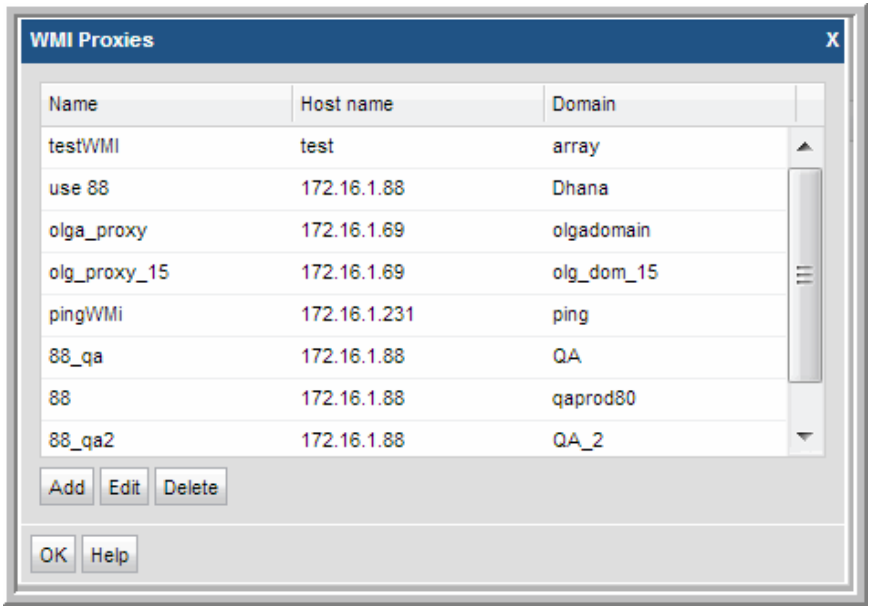
Multiple WMI Proxy settings can be created to manage access to Windows hosts.

See "Host resources prerequisites and configurations" on page 680.

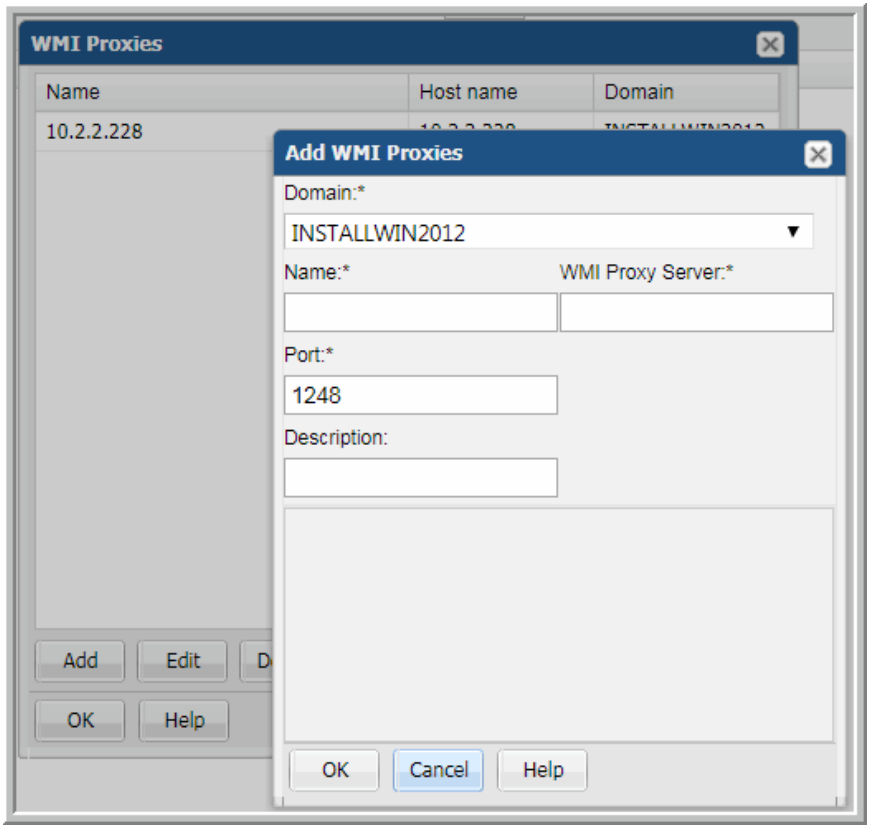
See "Host discovery and collection configuration steps" on page 684.

To Manage WMI Proxy settings select:

1. **Admin > Data Collection > Host Discovery and Collection.**
2. Click **Manage WMI Proxy**.



3. Click **Add** to configure settings and then click **OK**.



Field	Description
Domain*	Select the Domain from the list; for most environments, only one Domain is displayed. Multiple domains facilitate management for Managed Services Partners (MSPs).
Name*	Assign a name to identify this set of credentials that you are defining.
WMI Proxy Server*	This is the **server** address of the WMI proxy, which collects data on Windows hosts. Enter either the server's IP address or name.
Port*	The port that the Data Collector will use to contact the WMI Proxy; usually, there is no need to change the default setting (1248).
Description	Enter a note to help identify this WMI Proxy setting

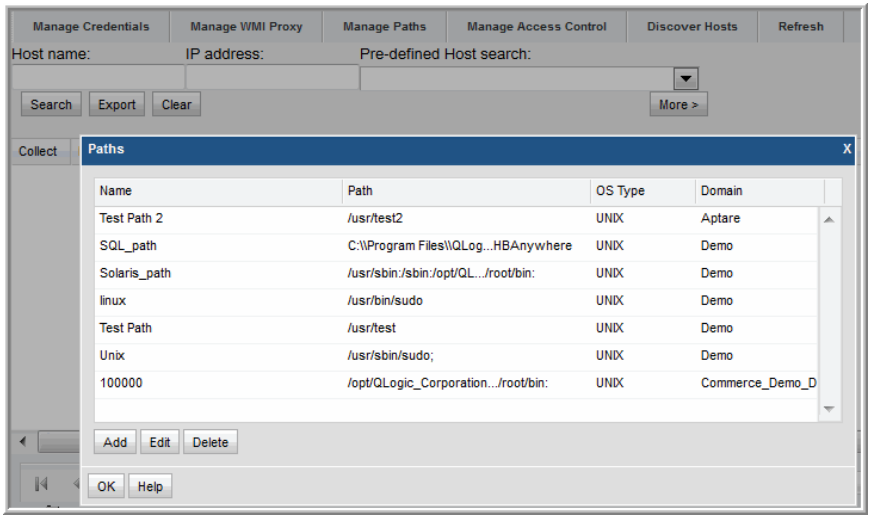
Manage paths

Multiple path settings can be created to designate specific paths to commands on hosts.The specified path is appended to the existing path and is used to search for commands (for example, /usr/bin:/usr/sbin). Certain commands, such as scli, require an absolute path.

See "Host resources prerequisites and configurations" on page 680.

See "Host discovery and collection configuration steps" on page 684.

1. Select **Admin > Data Collection > Host Discovery and Collection.**
2. Click **Manage Paths**.



3. Click **Add** to configure settings and then click **OK**.

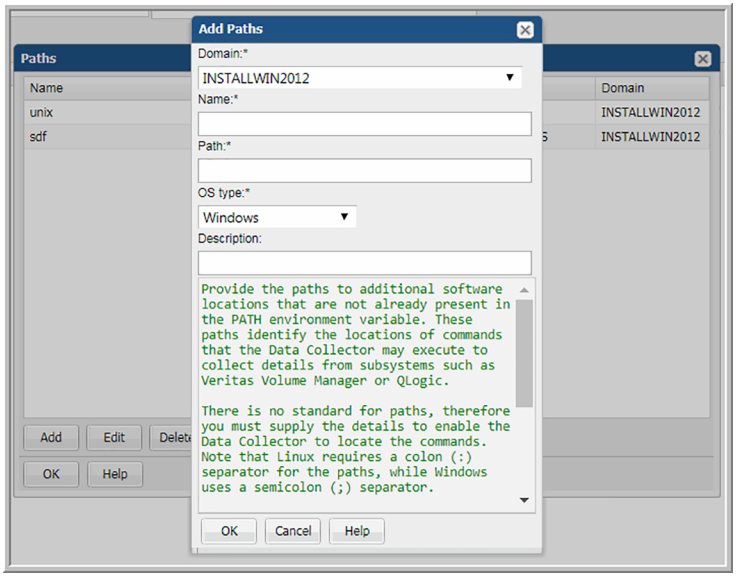


Table 19-3 Add Paths

Field	Description
Domain*	Select the Domain from the list; for most environments, only one Domain is displayed. Multiple domains facilitate management for Managed Services Partners (MSPs).
Name*	Assign a name to identify this Paths setting that you are defining.

Table 19-3 Add Paths *(continued)*

Field	Description

Path*	Provide the paths to additional software locations that are not already present in the PATH environment variable. These paths identify the locations of commands that the Data Collector may execute to collect details from subsystems such as Veritas Volume Manager or QLogic.
	There is no standard for paths, therefore you must supply the details to enable the Data Collector to locate the commands. Note that Linux requires a colon (:) separator for the paths, while Windows uses a semicolon (;) separator.
	Examples:
	Linux
	`:/opt/QLogic_Corporation/SANsurferCLI:/usr/local/sbin :/usr/local/bin:/sbin:/bin:/usr/sbin :/usr/bin:/root/bin:/opt/EMLXemlxu/bin :/usr/sbin/hbanyware:/opt/HBAnyware`
	Windows:
	`C:\Program Files\Emulex\Util\HBAnyware; C:\Program Files\QLogic\SANSurfer`

OS type*	Select either Linux or Windows
Description	Enter a note to help identify this Path setting

Manage access control

For Linux hosts, root-level privileges are required. Data Collectors require read-only access to execute non-intrusive commands on hosts. It is strongly recommended that a separate login account used strictly for IT Analytics be established and using Active Directory for Windows systems and the sudo command for Linux systems, restrict the commands that IT Analytics can issue. To accommodate this security approach, you can optionally specify access control commands like sudo, sesudo, or pbrun.

Files containing sudo commands per operating system can be found on the Portal server in: /opt/aptare/updates. These filenames contain both the OS and the version of the sudo commands file so that you can identify the files that contain the latest updates; for example: hpux_9.1.01, aix_9.1.01, linux_9.1.01, solaris_9.1.01.

Multiple Access Control settings can be created to manage access control commands for Linux hosts.

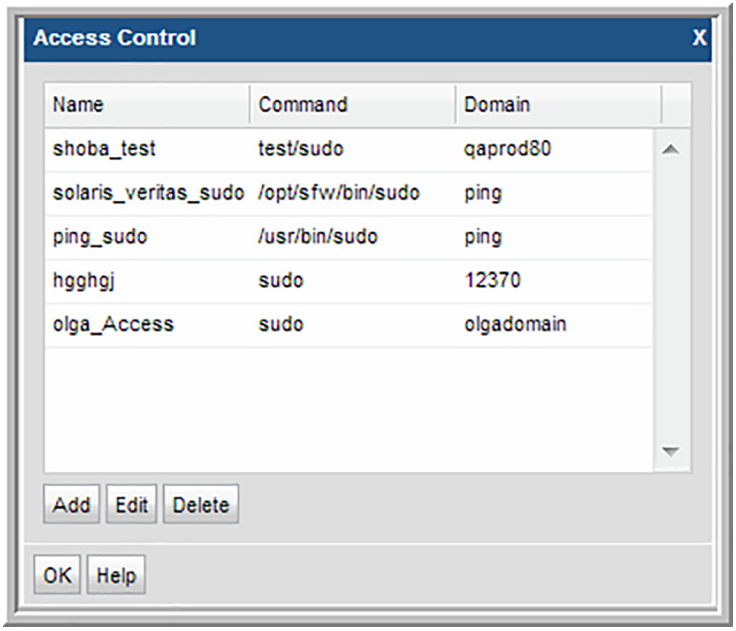
See “Host resources prerequisites and configurations” on page 680.

See “Host discovery and collection configuration steps” on page 684.

To Manage Access Control Settings

1 Click **Admin > Data Collection > Host Discovery and Collection**.

2 Click **Manage Access Control**.



3 Click **Add** to configure settings and then click **OK**.

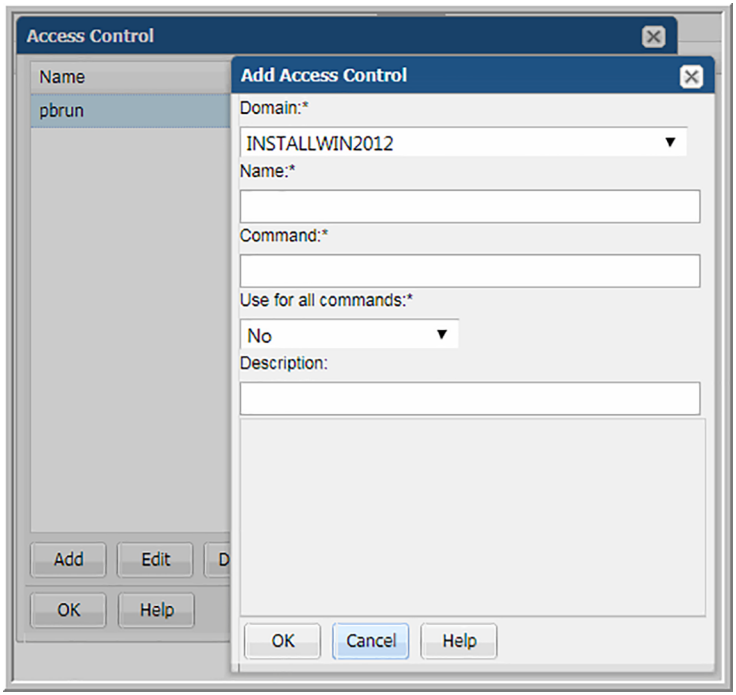


Table 19-4 Access Control Settings

Field	Description	Sample Values
Domain*	Select the Domain from the list; for most environments, only one Domain is displayed. Multiple domains facilitate management for Managed Services Partners (MSPs).	
Name*	Assign a name to identify this Access Control setting.	

Table 19-4 Access Control Settings *(continued)*

Field	Description	Sample Values

Command*	Linux hosts only: Provide the full path to the access control command, such as **sudo, sesudo,** or **pbrun**. Files containing sudo commands per operating system can be found on the Portal server in: /opt/aptare/updates. These filenames contain both the OS and the version of the sudo commands file so that you can identify the files that contain the latest updates; for example: hpux_9.1.01, aix_9.1.01, linux_9.1.01, solaris_9.1.01.
	
	You can configure sudo to prompt for a password using a custom prompt (the default is “Password”). The product expects the prompt to be “Password.” If the hosts have a custom password prompt, you’ll need to specify **-p Password** after the path to sudo. See the example to the right.
		`/usr/bin/sudo`
		
		`/usr/local/bin/sudo -p Password`

Use for all command*	Select Yes to have the Data Collector use the access command for all commands.	
Description	Enter a note to help identify this Access Control setting	

Host management

Now that you’ve set up the prerequisites, you’ll use the steps described in this section for on-going Host Management.

See “Host resources prerequisites and configurations” on page 680.

See “Host discovery and collection configuration steps” on page 684.

Once the prerequisite settings are configured, create a Host Discovery Policy to enable the process of finding hosts in your environment and populating your inventory of hosts.

The Host Management processes include:

- Configure and validate hosts in the inventory
- Enable and manage on-going collection

Several tools facilitate Host discovery and collection, as described in the following sections:

- See "Configure host discovery policies to populate the host discovery and collection view" on page 701.
- See "Execute and monitor host discovery" on page 705.
- See "Validate host connectivity" on page 708.
- See "Search and export in host discovery and collection" on page 712.
- See "Export in host discovery and collection" on page 714.
- See "Configure and edit host probes" on page 714.

Configure host discovery policies to populate the host discovery and collection view

Host Discovery begins with a Discovery Policy, which identifies the Data Collector that will gather information about hosts in your environment. In addition, a policy has an associated set of credentials, WMI proxies, and paths to access commands on the hosts.

See "Host resources prerequisites and configurations" on page 680.

See "Host discovery and collection configuration steps" on page 684.

A Discovery Policy typically is used once to initially populate your host inventory. Executing a discovery policy more than once has no effect for hosts that were previously discovered. To identify and resolve connectivity issues refer to the following.

See "Validate host connectivity" on page 708.

Although all hosts can be included in a single policy, you might want to create one or more Host Discovery Policies in the following recommended groupings:

- by OS (Windows or Linux) - This grouping is essential, as the probes and parameters are OS-specific.
- by common attributes, such as User ID, password, access control commands (sudo, pbrun, sesudo), PATH
- by application, such as Oracle or Exchange

Discovery policy considerations

If your enterprise configures hosts to lock out access after multiple failed authentication attempts, take the following tips into consideration:

- If you choose more than one credential in the Discovery Policy credentials list, you risk host authentication failure lock-out. The discovery process will try the first credentials and if they fail, discovery will try the next credentials that you've selected. Therefore, if your hosts are configured to prevent multiple authentication retries, multiple failed attempts may cause a lock-out.
- If multiple Discovery Policies are running simultaneously, with one policy using an IP address to access the host and the other policy using a name to access the host, the multiple access attempts may cause a lock-out. Note that if the authentication attempts are successful, only one host record is added to the inventory.

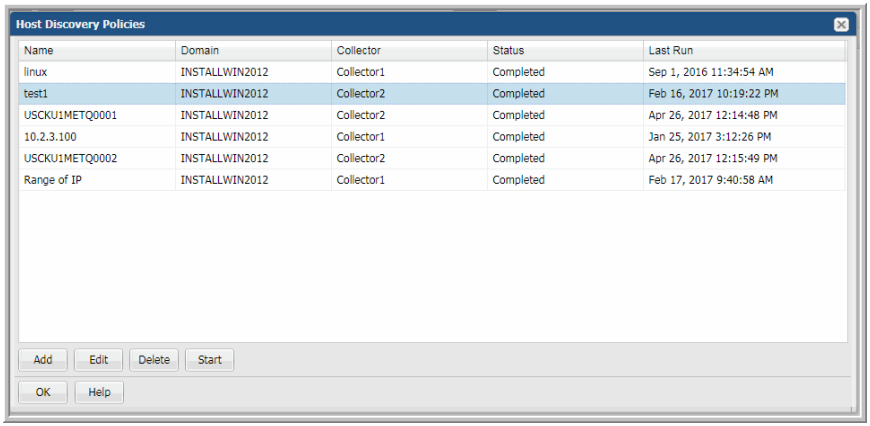
Configure a Discovery policy

Note: A Discovery Policy typically is used once to initially populate your host inventory. Executing a discovery policy more than once has no effect for the subsequent runs for hosts that have already been discovered and added to the inventory.

- See "Before discovering hosts" on page 687.
- See "Configure a Discovery policy" on page 702.
- See "Validate host connectivity" on page 708.
- See "Search and export in host discovery and collection" on page 712.

To create/edit Host Discovery Policies, in the toolbar select:

1. Select **Admin > Data Collection > Host Discovery and Collection**.
2. Click **Discover Hosts**.



3. Click **Add** to configure settings and then click **OK**.

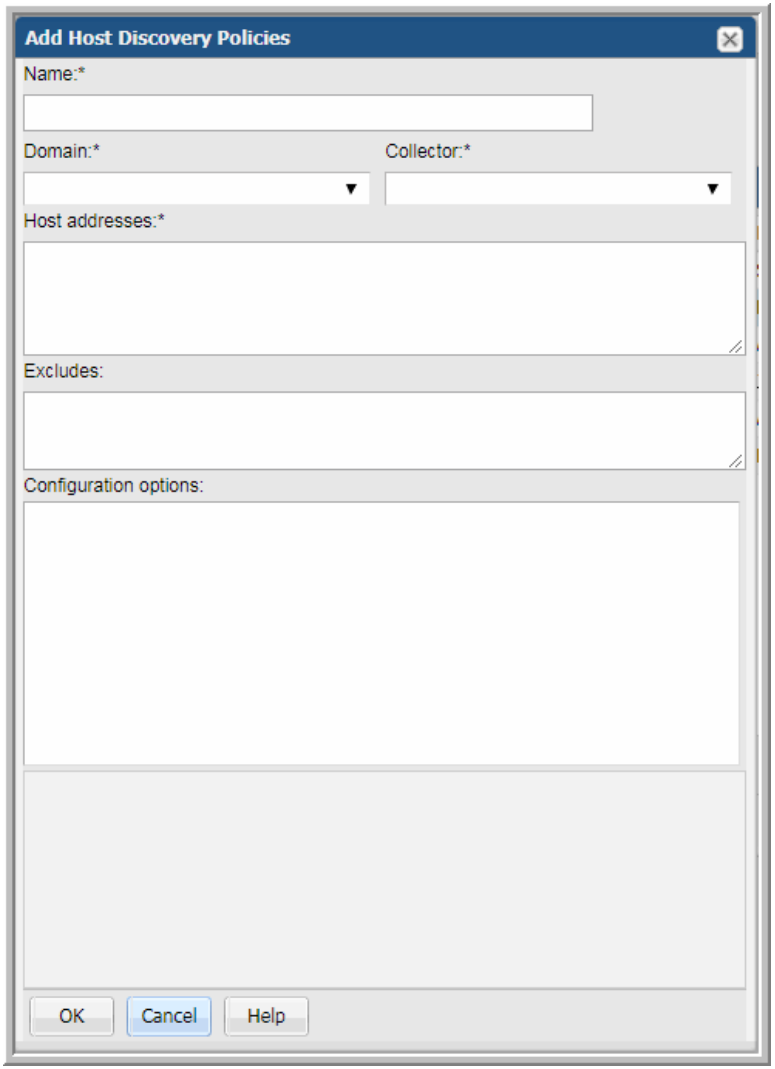


Table 19-5 Host Discovery policy settings

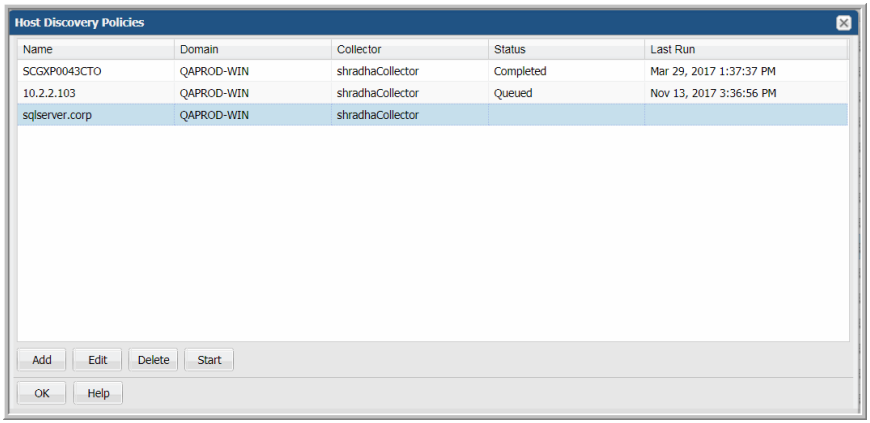
Field	Description
Name*	Assign a name to identify this Discovery Policy.
Collector*	Select the data collector from the drop-down list

Table 19-5 Host Discovery policy settings *(continued)*

Field	Description
Domain*	Select the Domain from the list; for most environments, only one Domain is displayed. Multiple domains facilitate management for Managed Services Partners (MSPs).
Host addresses*	■ A range of IP addresses can be specified ■ Hostnames and/or IP addresses can be listed, separated by commas
Excludes	List any known IP addresses that you know are not valid for host collection; for example, the IP address of a printer. IP address ranges are also supported.
Configuration options	This list gets populated when you select a Domain at the top of the Host Discovery Policies window. ■ Credentials ■ WMI Proxies ■ Paths Expand these lists to select the configurations to be used by this Discovery Policy. IMPORTANT NOTE: If you choose more than one credential in the list, you risk host authentication failure lock-out. The discovery process will try the first credentials and if they fail, discovery will try the next credentials that you've selected. Therefore, if your hosts are configured to prevent multiple authentication retries, multiple failed attempts may cause a lock-out.

Execute and monitor host discovery

1. Click **Discover Hosts** to list the Host Discovery Policies.



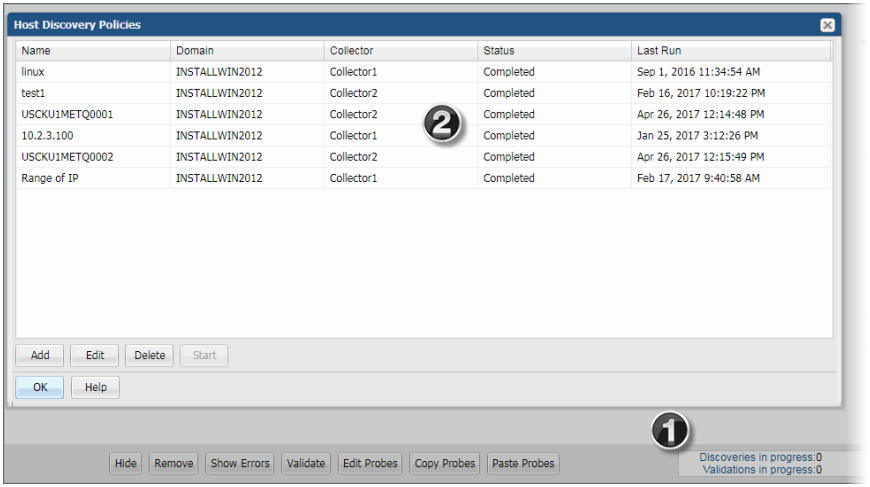
2. Select the Discovery Policy.
3. Click **Start**. At this point, the discovery process begins. It may take a few minutes for the background processes to initiate the discovery.
4. Click **OK**.
5. Refresh the view you launched in step 1 or use another method to verify that the discovery process is running.

See " Discovery processes" on page 706.

Discovery processes

Several methods can be used to progress:

Method 1



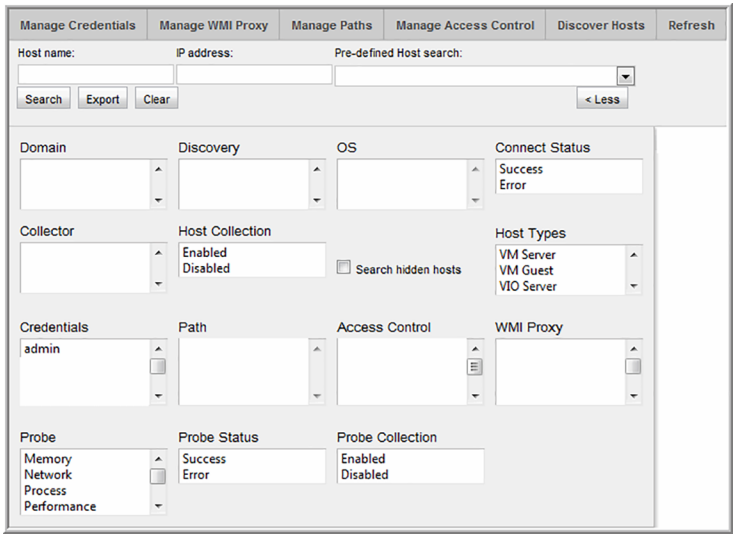
1. At the bottom of the **Host Discovery and Collection** window, double-click the **Discoveries in progress** link to launch the Host Discovery Policies window. Note the Status at the right of the window.
2. Double-click the Discovery Policy to view the settings.

Method 2

1. Using the Host Discovery and Collection

   See “Advanced Parameters” on page 713.

   function, search for hosts associated with a Discovery policy to see what hosts have been found.



Validate host connectivity

The Validate step executes the necessary validation steps and provides a summary of the overall success/failure. The Validation process steps through a handshake process, executing the preliminary steps that will occur during data collection. The informative messages enable you to pro-actively identify issues prior to initiating the data collection process.

Validate hosts

The Validation process identifies issues such as:

- Credential Validation Failures - Verify account IDs and passwords.
- Probe Errors - For example, an HBA probe may fail on a host that does not have an HBA. Other similar errors include iSCSI port not found or LUN not found.
- Connection Failures - The host may not be reachable.
- DNS Lookup Failures - IP addresses may not have been configured correctly.
- OS Verification Failures - Check the Access Control or WMI Proxy settings. In addition, verify that the paths are valid for the host's operating system.
- Command Errors - Verify that the Path and Access Control settings are correct.

Validation automatically occurs when the Data Collection processes are initiated; however, you can manually start the processes to get immediate feedback so that you can troubleshoot issues.

1. In the **Host Discovery and Collection** window, search for hosts. You can search by a Discovery Policy to see the results of a discovery.

   See " Discovery processes" on page 706.

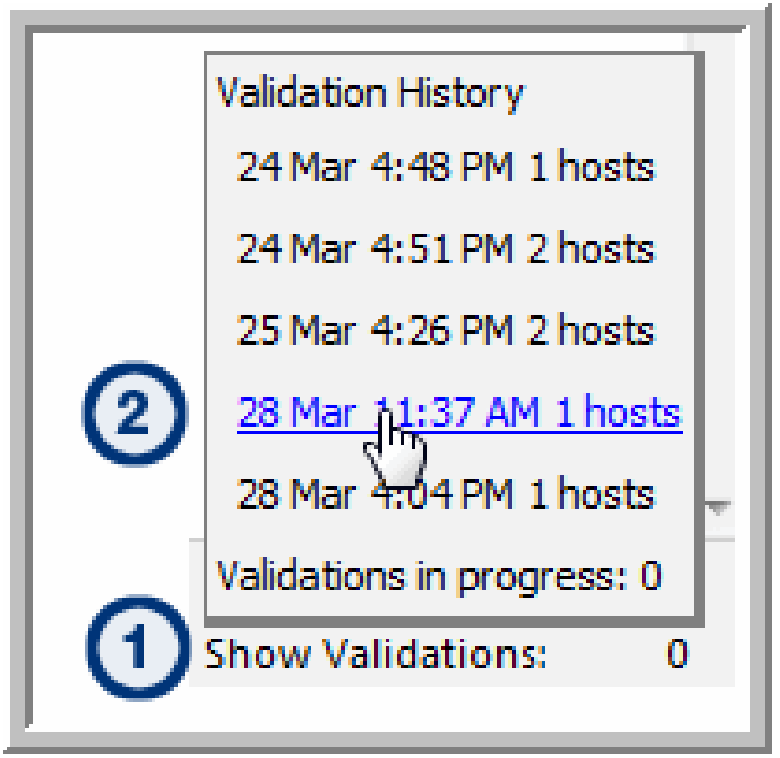
2. Click one or more hosts and then click **Validate** at the bottom of the window.
3. Click **Show Validations**at the bottom right of the window to verify that the validation has begun. A list of the hosts that are currently being validated are displayed.

   See "Validation history" on page 709.

Validation history

Once a set of hosts have been selected for validation the current status, as well as the history, can be viewed in a pop-up box.

1. At the bottom right of the **Host Discovery and Collection** window, click **Show Validations**.



- **Validation History**: A list of the past 10 validations is displayed. These are hyperlinks that can be used to access the list of hosts associated with that validation process.
- **Validations in progress**: Click this link to view the status of the current validation process.

2. Click the link to either display validations in progress or the hosts that were included in previous validations.

Show errors

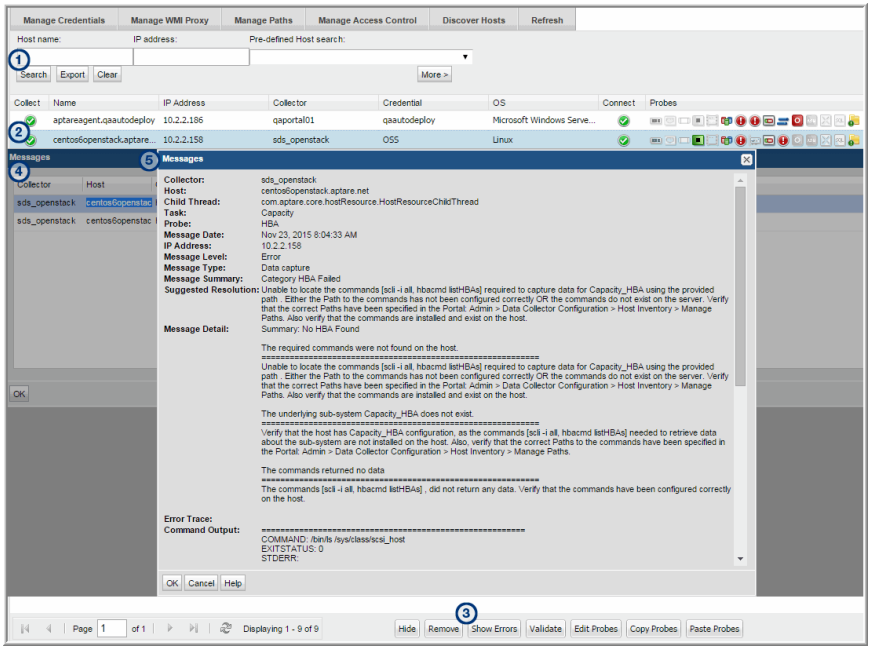
Before host data can be successfully collected, a number of configuration steps need to be taken. The **Show Errors** button enables you to identify details to help you troubleshoot host inventory collection issues.

Show Errors lists issues specific to:

- Connectivity
- Probes

- Validation

Use the following example and steps to view troubleshooting messages.



1. Search the **Host Discovery and Collection** window to view a list of hosts.
2. Select a host in the list that displays failure icons (in the above example, three probes have exclamation points in red circles).
3. Click **Show Errors** to display the Messages window for the selected host.
4. Double-click a message in the Messages window to view the details.
5. Take the recommended steps provided in the message details to rectify the issue. Then, re-validate the host.

Filter the host discovery and collection window - Hide/Unhide, remove

The Host Discovery process populates your inventory with hosts it finds. Often, discovery policies are designed to discover an IP address range. Host Discovery creates a record for every IP address in the range, even if it's not in use. Therefore, invalid IP addresses will appear in your Host Inventory. In addition, Discovery may

find printers, routers, or switches, or other devices that aren't relevant for host data collection.

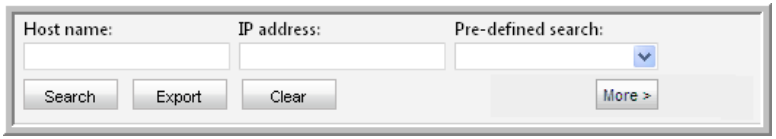
To filter your Host Discovery and Collection window to include only hosts for which you want data collected, use the following options:

- Hide - Select a host in the inventory and click **Hide** at the bottom of the window. When you Hide a host, it will not appear in your search results.
- Unhide - If you list hosts that have been hidden, the **Hide** button will be toggled to Unhide. Use the option, Search hidden hosts, to view a list of hidden hosts. See "Advanced Parameters" on page 713.
- Remove - Select a host in the inventory and click **Remove** at the bottom of the window. When you choose to remove a host, if you execute the Discovery Policy again, it will re-add it to the inventory. You may have an IP address that is now associated with a device that is different from the one that was discovered.

Search and export in host discovery and collection

The Host Discovery and Collection window offers a Search feature to help you find hosts that have been discovered.

Basic Search



Note: A search with no specified criteria returns all hosts in your inventory.

Pre-Defined Search

Several pre-defined searches enable easy access to host lists that are useful for troubleshooting.

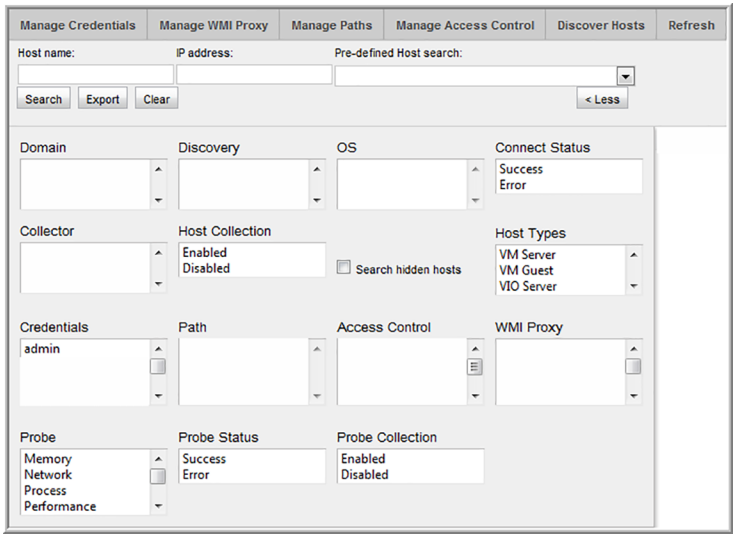
- Active policy but not collected since... (**Note**: When you select this option, a calendar pop-up enables date selection.)
- No active policy but was previously active (This means host data was successfully collected at an earlier time.)

- Credentials failing but were previously successful
- Collections failing but were previously successful

For more specific search parameters, click **More** (once clicking More, the button displays **Less**) to enter

See "Advanced Parameters" on page 713.

Advanced Parameters



- Select specific search criteria to narrow the list of hosts displayed in the inventory. Use Ctrl+Click to select multiple values.
- Search criteria within a category is mutually exclusive, so if you select multiple values, it functions as an OR statement. Search criteria across categories is an AND statement. For example, if you select the values **Memory** and **Network** from the **Probe** category and the value **Error** from **Probe Status**, the search query is: (Probe = Memory OR Probe = Network) AND Status =Error
- When you check the Search hidden hosts box, only hidden hosts will be displayed in the Host Inventory window. Also, the Hide button in the Host Inventory window will toggle to Unhide.
- When searching on Probes, if a probe was at some point activated, but then de-activated, it will appear in the search results because there is an entry in the database table.

Export in host discovery and collection

To export the details of the **Host Discovery and Collection** to a comma-separated-values file (.csv):

1. Search without supplying any values in the search criteria fields.

   See "Search and export in host discovery and collection" on page 712.

2. In the Search area at the top of the Host Discovery and Collection list, click **Export**.

The resulting file will include values for the status of each of the available probes. For example, the values will be similar to N/U or Y/S, as described in the following table.

E	Error
F	Failure
N	No - Not Active
S	Success
U	Unknown
W	Warning
Y	Yes - Active Probe

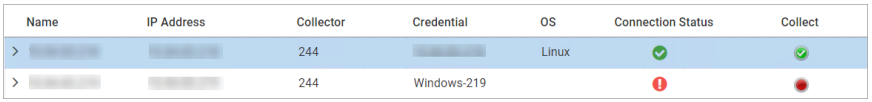
Configure and edit host probes

Host Data Collection can gather the following information by probing hosts:

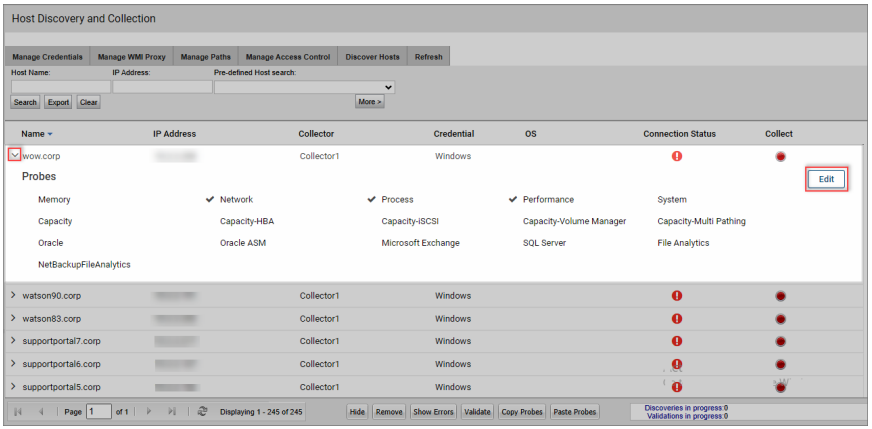
- Host Probes: Capacity (HBA, iSCSI, Volume Manager, Multi-pathing)
- Host Probes: Memory, Network, Process, Performance, System
- Application Probes: Exchange, SQL Server, Oracle, Oracle ASM
- Probes

To configure probes for a host:

1 Search for hosts in the **Discovery and Collection** list. A search with no specified criteria returns all hosts in the **Discovery and Collection** list.

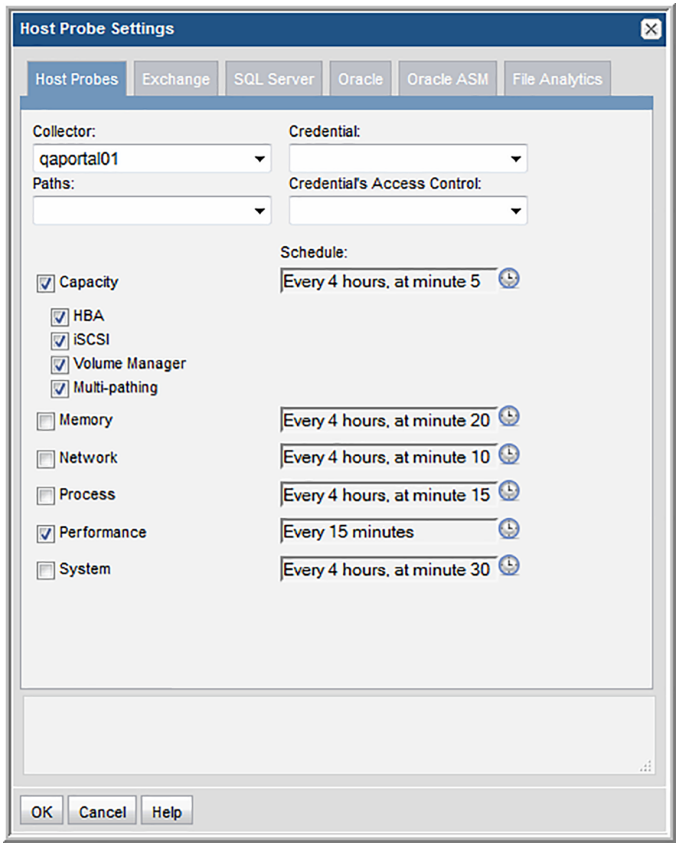


2 Click **>** next to the host name to expand the row. The probes will be visible in the expanded view.



3 Click **Edit** to configure/view the Host Probe Settings window. For details:

See "Probe Settings" on page 718.



4 Click each tab to updated the configuration settings for the specific probes.

5 For the SQL Server and Oracle probes, you can create multiple instances, using the following steps:

Click **Add**.

Enter the mandatory configuration.

Click **OK**.

Propagate probe settings: Copy probes, paste probes

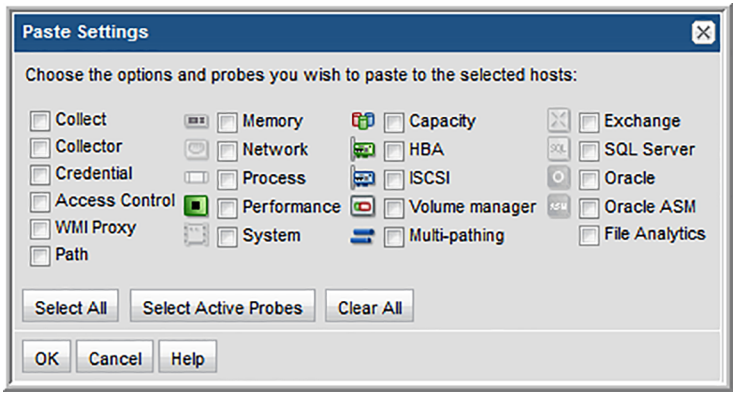
Whenever you have hosts with common attributes, you can save time by configuring probe settings for one host and then copying and pasting those settings to other hosts.

A key advantage to using the probe copy/paste feature is the ability to propagate the probe schedules to multiple hosts. In addition, you can explicitly select the probes you want to activate.

Note: You only can copy/paste probes that are within the same Domain. This mainly impacts Managed Services Partners with multi-domain environments. Use the Advanced Search function to list probes within a specific Domain.

Example of Probe Copy/Paste

1. Search for all Linux hosts.
2. Configure the probes for one of the hosts in your Linux list and click **Copy Probes**.
3. Finally, select the remaining Linux hosts and click **Paste Probes**.



- The icons of configured Probes will be highlighted in the Paste Probes window; however, you must explicitly check those probes to copy the probe schedules and to activate the probes. Use the **Select Active Probes** button to select active probes.
- By default, the probe check boxes are unchecked, enabling you to explicitly select the probes that you want to paste. Or, click Select All to turn on all the probes for the selected host.

Probe Settings

Table 19-6 Probe Settings

Probe Type	Parameters	Description
Capacity, HBA, iSCSI, Volume Manager, Multi-pathing	Probe schedule*	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
Memory	Probe schedule*	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
Network	Probe schedule*	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
Process	Probe schedule*	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
Performance	Probe schedule*	For Windows, only CPU performance will be collected. For Linux, both CPU and device performance will be collected. CPU performance includes CPU allocation and usage. Device performance includes throughput and latency for disk partitions. **Note:** The sysstat utility must be installed on the Linux servers or storage nodes for Linux host performance data collection. A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”

Table 19-6 Probe Settings *(continued)*

Probe Type	Parameters	Description
System	Probe schedule*	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”.
Exchange	Collect	Check this box to activate collection on an on-going basis. When it is unchecked, only initial validation will attempt this probe.
	Probe schedule	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
	Active Directory Host	Host name or address
	Active Directory Port	For example: 389
	Active Directory Base DN*	The starting point for the Active Directory. For example: CN=Services,CN=Configuration, DC=contoso2003,DC=com Several tools are available to help you identify the Base DN: **Ldp.exe** - http://support.microsoft.com/kb/224543 **adsiedit.msc** - http://technet.microsoft.com/en-us/library/cc773354(WS.10).aspx
	Active Directory User Name	Active Directory User Name This username must have privileges to search under the base DN within the Active Directory. Typically, this is an Administrator.
	Password	Active Directory Password
SQL Server	Collect	Check this box to activate collection on an on-going basis. When it is unchecked, only initial validation will attempt this probe.
	Probe schedule	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”

Table 19-6 Probe Settings *(continued)*

Probe Type	Parameters	Description
	Database*	The name of the database within the SQL server.
	Instance	The system identifier to identify the SQL server database instance--for example: BKUPEXEC. Specify either an instance name or a port. If an instance name is not specified, MSSQLSERVER is substituted.
	Port	To identify the SQL server instance, provide either an instance name or a database port number; for example: 1433. If a port number is not specified, the port is determined automatically from the instance name.
	Account*	Database access user name The data collector requires a user account with permissions to execute the stored procedures
	Password*	Database access password
	Windows Authentication	Check this box if you want Windows authentication rather than SQL server authentication.
Oracle	Collect	Check this box to activate collection on an on-going basis. When it is unchecked, only initial validation will attempt this probe.
	Probe schedule	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
	SID*	The system identifier to identify the database instance.
	Port*	Database port number; default: 1521
	Username*	The Oracle user must have the following role granted: SELECT_CATALOG_ROLE To grant this access, use: GRANT SELECT_CATALOG_ROLE TO ‘user’ where user is the database Username that you’ll provide here.
	Password*	Database access password

Table 19-6 Probe Settings *(continued)*

Probe Type	Parameters	Description
Oracle ASM	Collect	Check this box to activate collection on an on-going basis. When it is unchecked, only initial validation will attempt this probe.
	Probe schedule	A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”
	Account*	The Oracle user privileges required: SYSDBA privilege if 10g sysasm in 11g
	Password*	Database access password
	Port*	Database port number; default: 1521
	ASM Instance*	The name that identifies the database instance.
	Collect	Check this box to activate collection on an on-going basis. When it is unchecked, only initial validation will attempt this probe.
	Probe schedule	Default is once a month. A schedule in cron format; for example: */20 9-18 * * * which translates to “every 20 minutes between the hours of 9 a.m. and 6 p.m.”

Chapter 20

Discovery policies for Veritas NetBackup

This chapter includes the following topics:

- Task overview: Configure and monitor discovery policies
- Discovery policies overview
- About Discovery types
- Activate a discovery license
- Exclude devices from discovery policies
- Activate discovery probes in the NetBackup Data Collector policy
- Monitor discovery processes
- View client protection status

Task overview: Configure and monitor discovery policies

To configure Discovery, perform the following sequence of steps:

Table 20-1 Configure and monitor discovery policies

	Task	For Instructions
1.	Learn about how Discovery policies can help you protect your data.	See "Discovery policies overview" on page 724.

Table 20-1 Configure and monitor discovery policies *(continued)*

	Task	For Instructions
2.	Purchase and activate your Discovery license. Two of the three Discovery processes require a license: See "Client drive discovery" on page 724. See "Backup policy coverage" on page 725.	See "Activate a discovery license" on page 725.
3.	Enable SNMP, if you are enabling these Discovery types: See "Client drive discovery" on page 724. See "Backup policy coverage" on page 725.	See the *System Administrator Guide*.
4.	Determine the primary server that requires the policy that you are about to create, and identify the Discovery type(s) that you want to enable on this primary server.	See "About Discovery types" on page 724.
5.	If necessary, exclude specific network devices from your policies.	See "Exclude devices from discovery policies" on page 726.
6.	Turn on Discovery probes in the NetBackup Data Collector policy.	See "Activate discovery probes in the NetBackup Data Collector policy" on page 726.
7.	Regularly monitor the status of Discovery processes.	See "Monitor discovery processes" on page 726.
8.	View the Client Protection Summary report to see how well your data is being protected.	See "View client protection status" on page 727.
9.	If significant changes in your environment warrant a fresh view, rebuild the Discovery database.	
10.	Tune Discovery by modifying time out settings for probes.	See the IT Analytics System Administrator Guide for details.

Discovery policies overview

The Discovery module, specific to Veritas NetBackup, uses Discovery policies to illuminate risk and exposure within the corporate IT backup and recovery environment. The Discovery module is a separately licensed feature.

See "About Discovery types" on page 724.

See "Activate a discovery license" on page 725.

Discovery policies provide answers to the following questions:

- Where is my data protected? (for example, disk-to-disk, disk-to-tape, or disk-to-disk-to-tape)
- What is the extent and coverage of my data protection?
- Are all my clients and applications protected?
- Is every data set on every client and every application protected?

Discovery finds hosts on a corporate network and compares those hosts with the policies of the underlying backup and recovery software. Discovery performs the following steps:

1. Identifies orphan clients that are not being protected.
2. Probes and determines the file systems or drives of the hosts.
3. Compares and contrasts the file systems to the equivalent policies within the underlying backup and recovery software.

Use Discovery policies if:

- Your IT infrastructure, applications, and servers are rapidly changing.
- Your backup solution cannot detect your backup servers and cannot provide information about successful or unsuccessful backups.

About Discovery types

Three different Discovery types can be configured to collect additional NetBackup data. To configure and manage Discovery refer to the following.

Client drive discovery

This feature requires a Discovery license and SNMP. This Discovery process seeks out hosts and devices in your environment. The process identifies all hosts in your environment, in particular those that are not currently stored in the reporting database

and are therefore potentially not being backed up. This probe uses SNMP to probe the IP address range for drive utilization; therefore, SNMP must be enabled.

Media server disk discovery

This Discovery process probes all the media servers associated with the management server to gather disk-based information such as capacity and free space on the media server file systems. This information is then displayed in the Disk Usage and Performance report. If the Media Server Disk Discovery process is not enabled, disk-based information will show as Unknown in reports. If you have several primary servers in your environment, and they have media servers and disk storage units attached to them, you must enable the Media Server Disk Discovery module on each of the primary servers.

Backup policy coverage

This feature requires a Discovery license and SNMP. This Discovery process, probes all the NetBackup clients known to the NetBackup database that are associated with the management server. It queries NetBackup to discover if there are backup policies that cover the client. A client is determined to be associated with the NetBackup management server if it belongs to a policy associated with the management server. This probe uses SNMP to probe for drive utilization; therefore, SNMP must be enabled.

Activate a discovery license

You need to activate your Discovery license so that you can access the additional discovery features beyond the Media Server Disk Discovery component.

A Discovery license is required for the following Discovery types:

- Client Drive Discovery
- Backup Policy Coverage

To activate the Discovery license

1. Go to the utilities directory.

   Linux: **/opt/aptare/utils**

   Windows: **C:\opt\aptare\utils**

2. Run the following license utilities to view the status of your current license or to install your updated license.

   Linux:

./printLicense.sh

./installLicense.sh

Windows:

printlicense.bat

installlicense.bat

Exclude devices from discovery policies

An exclude list is a list of names or IP addresses that will not be probed by any of the Discovery policies configured for a given management server. Each management server maintains its own exclude list.

To exclude devices from Discovery policies

1. In the **Discovery Administration** window, enter a comma-separated list of the IP addresses that you want to exclude.
2. Click **OK**.

Activate discovery probes in the NetBackup Data Collector policy

The NetBackup Data Collector Policy lists probes that can be turned on to collect different types of data. Three of these probes are specific to Discovery.

- See “Client drive discovery” on page 724.
- See “Media server disk discovery” on page 725.
- See “Backup policy coverage” on page 725.

Monitor discovery processes

To monitor a Discovery process

1. From the Portal toolbar, view the Discovery Administration window by selecting **Admin > Reports > Discovery Policies**.
   - **Inactive**. Indicates that there are currently no active policies for the particular Discovery process.

- **Active**. Indicates that there is at least one active policy for the particular Discovery process. To access the individual Discovery processes, click on the management server row.

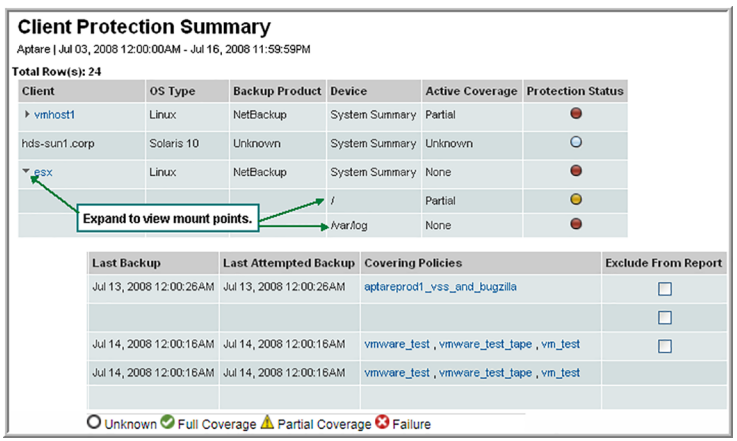
2 For each active policy, double-click on the management server that is responsible for running a particular policy.

3 Using the **last run status** field, determine the status of the Discovery process that last ran:

- Failed. Indicates a problem during the execution of the policy or a problem with saving the data to the Reporting Database. Check the **mbs/logs/crontab.log** file for detailed information about the failure.
- Partial. Indicates one or more probes time out and a response was not received.

View client protection status

The Client Protection Summary report provides a view of the protection status of clients that you think are being backed up by NetBackup.



See the *Report Reference Guide* for details about this report.

Chapter 21

View and manage system notifications

This chapter includes the following topics:

- Overview
- Viewing system alerts
- Defer or suppressing notifications

Overview

At login, an alert indicator notifies the Administrator to any IT Analytics issues that require attention. This is shown as a badge on the main toolbar. You can click the badge to launch the System Notifications window or navigate to the window using the menu.

Viewing system alerts

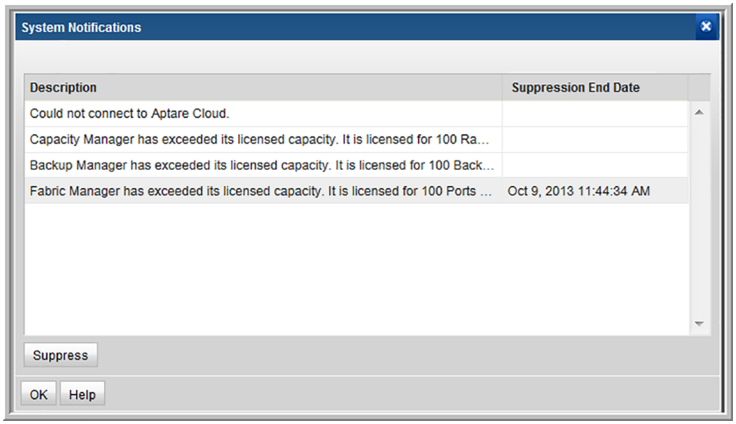
The System Notifications window lists system events that require your attention:

- Pending license expiration
- Collection reached license limitations by product
- Database events--rejected clients, ESX servers, arrays--due to license overage
- Version notification--newer version of data collector software/files is available
- Unable to connect for the latest software updates (proxy not configured)

To view system notifications

1 Click the badge or navigate to **Admin > Advanced > System Notifications** to read and acknowledge notifications.

The **System Notifications** window is displayed.



2 Click **OK** to acknowledge the list of notifications.

Clicking **OK** does not remove the notification. The next time you log in, the notification is displayed again, unless you suppress it.

See "To suppress or defer a notification" on page 730.

Often the message is a notification that a newer software version is available for download.

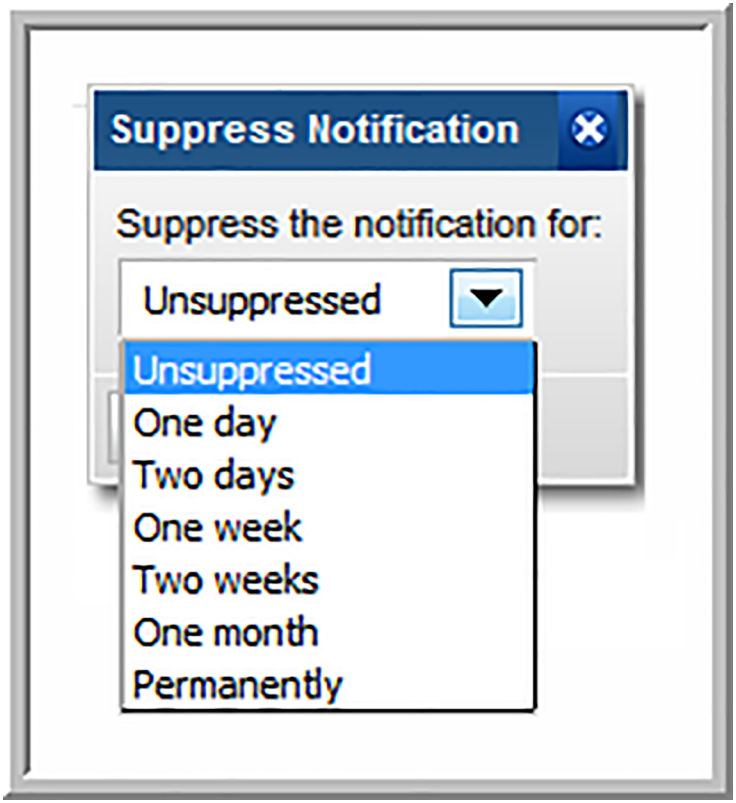
Defer or suppressing notifications

Not all notifications require your immediate attention. In this case, you can defer taking action.

To suppress or defer a notification

1 Select a notification in the list and click **Suppress**.

A drop-down list provides options.



2 Select the suppression period the message and click **OK**.

The notification will display in your browser until either the time period has elapsed or the problem is corrected.

Chapter 22

Customize with advanced parameters

This chapter includes the following topics:

- Overview of advanced parameters
- Use cases for advanced parameters
- Adding an advanced parameter
- Access control advanced parameters
- General Data Collection advanced parameters
- Cloud data collection advanced parameters
- Host discovery and collection advanced parameters
- Backup Manager advanced parameters
- Capacity Manager advanced parameters
- File Analytics advanced parameters
- Virtualization Manager advanced parameters
- Fabric Manager Advance Parameters

Overview of advanced parameters

Advanced Parameters should be configured only when directed by Support, to enable/disable functionality for certain circumstances, such as improving performance or gathering details to troubleshoot data collection.

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

A mechanism is provided for customizing internal parameters to meet the needs of your organization. These parameters can be configured and restricted to the following levels:

- Domain - This scope is useful for Managed Services Partners where the parameter can be applied to all data collectors within a specific environment
- Collector - Restricts the parameter change to a specific Data Collector
- Operating System - Linux, Solaris, Windows, or AIX
- Servers - Defines the pre-existing hosts with the Domain or any new hosts
- Command - Defines the command or list of commands to which the parameter value is restricted

Example

SSH port - A change could be made at the domain level and then could also be set for a collector and even for a specific host.

Additional Details

- See “Use cases for advanced parameters” on page 732.
- See “Adding an advanced parameter” on page 734.
- See “Access control advanced parameters” on page 736.
- See “General Data Collection advanced parameters” on page 738.
- See “Cloud data collection advanced parameters” on page 745.
- See “Host discovery and collection advanced parameters” on page 748.
- See “ Backup Manager advanced parameters” on page 752.
- See “Capacity Manager advanced parameters” on page 758.
- See “File Analytics advanced parameters” on page 761.
- See “ Virtualization Manager advanced parameters” on page 763.

Use cases for advanced parameters

The following examples provide simple use cases for using advanced parameters.

Domain/Collector

- Set the access control prompt to new value
- Control the protocols used - WMI / SSH / Telnet
- Control the 3rd-party SSH used for all hosts
- Control the ciphers used for JSch for all hosts
- Increase SSH debugging
- Set the access control prompt for all hosts
- Set the standard out error strings for all hosts
- Set the access control command requirement for TTY
- Set the access control required error string for an OS
- Set the number of min, max threads for host resources
- Set connection time-out for connections to hosts
- Set socket time-out for connections to hosts
- Set the SSH port for hosts

Operating system

- Set the collection to load from raw data for an OS
- Set the access control prompt for an OS
- Set the standard out error strings for an OS
- Set the access control command requirement for TTY
- Set the access control required error string for an OS
- Set connection time-out for an OS
- Set socket time-out for an OS
- Set SSH port for an OS

Host

- Set the collection to load from raw data for a specific host
- Set the access control prompt for a specific host
- Set the standard out error strings for a specific host
- Set the access control command requirement for TTY

- Set the access control required error string for a specific host
- Set connection time-out for a specific host
- Set socket time-out for a specific host
- Set SSH port for a specific host

Commands

- Set the time-out for a specific command
- Set the time-out for all commands for a host
- Set the time-out for all commands for an OS
- Set the time-out for a specific command for a host
- Set the time-out for a specific command for an OS
- Set the time-out for all commands

Adding an advanced parameter

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

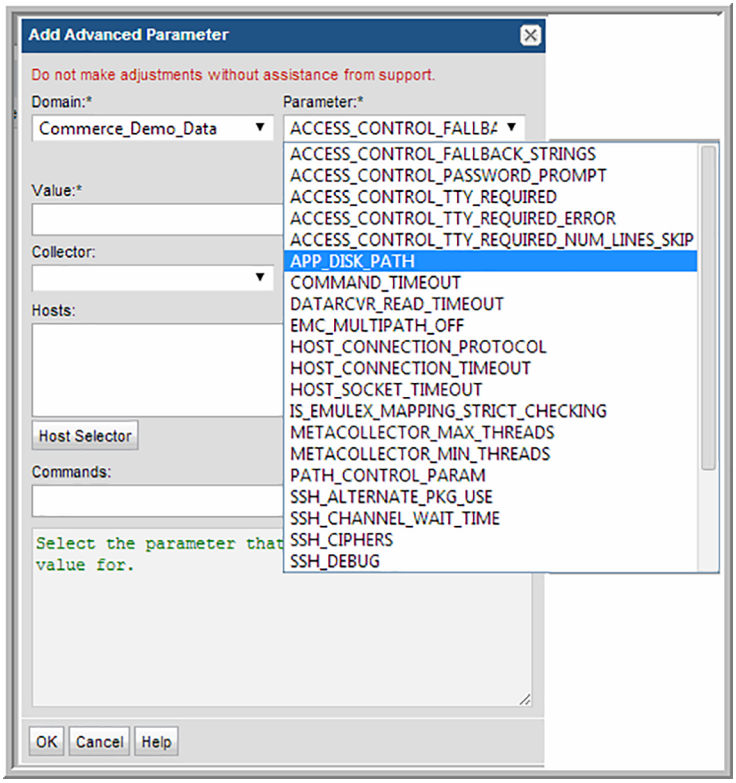
Advanced Parameters by Function and Product

- See “Access control advanced parameters” on page 736.
- See “General Data Collection advanced parameters” on page 738.
- See “Cloud data collection advanced parameters” on page 745.
- See “Host discovery and collection advanced parameters” on page 748.
- See “ Backup Manager advanced parameters” on page 752.
- See “Capacity Manager advanced parameters” on page 758.
- See “File Analytics advanced parameters” on page 761.
- See “ Virtualization Manager advanced parameters” on page 763.

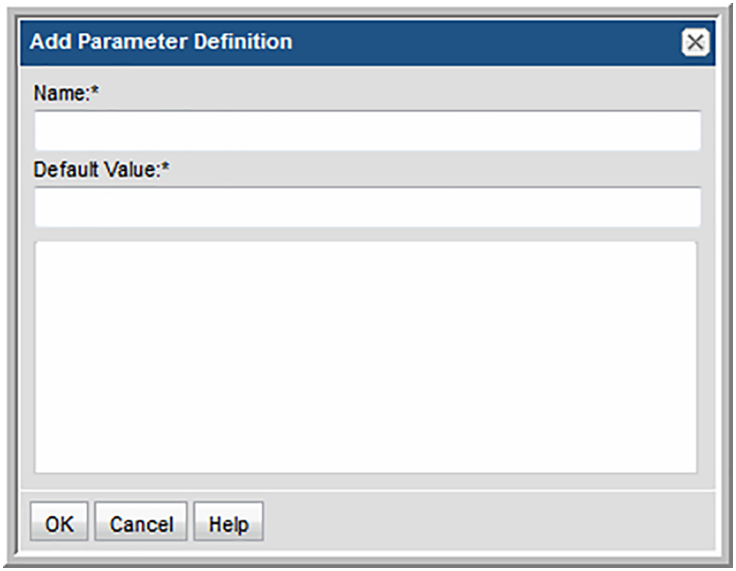
A standard set of Advanced Parameters is shipped with the Portal software. Other parameters can be added, using the following steps.

1. Select **Admin > Advanced > Parameters**.

2. Click **Add** to incorporate an advanced parameter listed in one the sections listed in the following.



3. Click the **Parameter** drop-down list to view the possible out-of-the-box parameters that can be configured. Note that a custom parameter can be added by clicking **Add** (underneath the **Parameter** list).
4. In the Add Parameter Definition window, enter the specific advanced parameter name. Note that the name is case-sensitive and must precisely match the supported parameter name, as listed in this documentation.



Access control advanced parameters

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

The following advanced parameters apply to host data collection.

Admin > Advanced > Parameters

See “Adding an advanced parameter” on page 734.

ACCESS_CONTROL_FALLBACK_STRINGS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text Strings (Default = See below)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS, or for a specific host.

For Host Resource data collection, when this parameter is configured with error text strings, if stdout and stderr contain any of these strings, an error is reported. These errors indicate that a command has failed with sudo configuration problems.

Default text strings for this advanced parameter:

```
is not allowed to execute|sudo: A file or directory in the path name
 does not exist| not found|ermission denied|is not allowed to run
sudo
```

ACCESS_CONTROL_PASSWORD_PROMPT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text (Default = See below)

Use this parameter to control the password prompt for the SSH access control command.

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS, or for a specific host.

Default text strings for this advanced parameter:

```
Password:|password:|Password for|password for
```

Example:

```
assword:|assword for"  //  [Pp]assword: or ?[Pp]assword for
```

ACCESS_CONTROL_SKIP_COMMAND_OUTPUT

Valid Values: Text

Scope: This parameter can be set at the Domain or Data Collector level, for an OS, or for a specific host.

Enter any command that you want to skip when running sudo. This can be used to prevent security violations.

ACCESS_CONTROL_TTY_REQUIRED

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (default = N)

Scope: This parameter can be set at the Domain or Data Collector level, for an OS, or for a specific host.

Use this parameter if your access control commands requires TTY. For example, recent versions of sudo require this functionality.

ACCESS_CONTROL_TTY_REQUIRED_ERROR

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text (Default = See below)

Scope: This parameter can set the access control requirements for TTY, at the Domain or Data Collector level, for an OS, or for a specific host.

Use this parameter to configure the TTY-required error strings that would allow the Data Collector to run the program to use TTY dynamically.

Default text strings for this advanced parameter:

```
you must have a tty to run sudo|no tty present and no askpass program
 specified
```

ACCESS_CONTROL_TTY_REQUIRED_NUM_LINES_SKIP

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric (default = 1)

Scope: This parameter can set the access control requirements for TTY, at the Domain or Data Collector level, for an OS, or for a specific host.

Use this in conjunction with the following.

See “ACCESS_CONTROL_TTY_REQUIRED” on page 737.

See “ACCESS_CONTROL_TTY_REQUIRED_ERROR” on page 738.

When TTY is enabled, sudo adds extra data that is not related to the relevant output. This parameter will help to eliminate the extra lines in the output.

General Data Collection advanced parameters

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

The following advanced parameters apply to a variety of general collection functions, not necessarily associated with a particular subsystem.

To configure the parameters used for Data Collection, navigate to **Admin > Advanced > Parameters.**

See "Adding an advanced parameter" on page 734.

DC_START_CONNECT_RETRIES

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value (Default = 120).

Scope: This parameter can be configured at the Domain or Data Collector level, for an OS, or for a specific host.

Use this parameter to modify the connection attempts for a Data Collector to connect to the portal when starting the APTARE agent. The Data Collector attempts to connect to the portal for specified value times with 5 seconds of wait time between each attempt.

This might be necessary when Data Collector and Portal system both are re-started at almost same time.

DC_UPGRADE_HANDLE_UTILITY_PATH

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Folder Path in Data Collector Filesystem (Default = C:\Program Files\Handle)

Scope: This parameter can be set for all Data Collectors or for a specific Data Collector.

Use this parameter to specify the location of Handle utility, which is used to kill the file-locking processes during Data Collector upgrade.

DC_UPGRADE_KILL_FILE_LOCKING_PROCESSES

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (Default = Y)

Scope: This parameter can be set for all data collectors or for a specific data collector.

Use this parameter to specify whether to kill file-locking processes during data collector upgrade automatically. If its value is set to N and file-locking processes are detected during upgrade, then upgrade will not proceed further. You will have to kill such processes manually before starting the upgrade again.

COMMAND_TIMEOUT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in milliseconds (Default = 30000)

Scope: This parameter can be configured at the Domain or Data Collector level, for an OS, or for a specific host.

Use this parameter to modify the wait time for a command to send a reply. This might be necessary when data collection times out and cannot complete. Typically, this advanced parameter will be used in large collection environments where factors such as network performance and the sheer amount of data can impact responses.

DCLOGS_CLEANUP_LIMIT_IN_MONTHS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Scope: This parameter can be set for all data collectors or for a specific data collector level.

Default value is 6 (indicating 6 months). Value can be any positive integer. The aptare_agent service runs a logs cleanup task on service startup and after every 24 hours from the current run. Cleanup task deletes DC Logs under <AnalyticsHome>/mbs/logs folder, which are older than the configured value at the time when Cleanup task runs.

DATARCVR_READ_TIMEOUT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in seconds (Default = 600)

Scope: This parameter can be configured at the Domain or Data Collector level, for an OS, or for a specific host.

Use this parameter to modify the wait time for the data receiver to send a reply before the process times out. This might be necessary when data collection times out and cannot complete. Use this parameter if you are seeing data collection errors such as, Read timed out. You can increase the default time-out value for the

connection time (in seconds) between the data collector and the data receiver. Assign this parameter to the problem collector. Typically, this advanced parameter will be used in large collection environments where factors such as network performance and the sheer amount of data can impact responses.

LOGGING_LEVEL

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: ERROR, WARN, INFO (Default), DEBUG, TRACE

Scope: This parameter can be set at the Domain, Data Collector, and Host levels.

Use this parameter to set the minimum logging level for the specified collector subsystem. When configured without selecting any hosts, it sets the minimum logging level for all subsystems in the domain/collector. When hosts are selected, it sets the minimum logging level for when the selected hosts are being processed. Defaults to INFO. If set to a more restrictive level than the logging configuration it will have no effect. The use case is to temporarily increase the logging level for analysis and debugging.

METACOLLECTOR_MAX_THREADS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Integer that represents the maximum number of threads (Default = 20)

Scope: This parameter can be set at the Domain and Data Collector levels.

Use this parameter to modify the maximum number of threads within a metadata collector. This thread pool manipulation enables changes to improve performance.

METACOLLECTOR_MIN_THREADS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Integer that represents the minimum number of threads (Default = 5)

Scope: This parameter can be set at the Domain and Data Collector levels.

Use this parameter to modify the minimum number of threads within a metadata collector. This thread pool manipulation enables changes to improve performance.

PATH_CONTROL_PARAM

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text value of **find** or **all** (Default = **find**)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS, or for a specific host.

Use this parameter to control the path-finding mechanism in the Data Collector for host resources data collection.

- find = The Data Collector finds the command paths.
- all = The Data Collector tries all the command paths until successful and the find command will not be executed.

RAWDATA_CLEANUP_LIMIT_IN_MONTHS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Scope: This parameter can be set for all data collectors or for a specific data collector level.

Default value is 6 (indicating 6 months). Value can be any positive integer. The aptare_agent service runs a rawdata cleanup task on service startup and after every 24 hours from the current run. Cleanup task deletes rawdata under <AnalyticsHome>/mbs/rawdata folder, which are older than the configured value at the time when Cleanup task runs.

SSH_ALTERNATE_PKG_USE

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level or for a specific host.

Use this parameter designate which SSH package should be used to connect to a host.

SSH_CHANNEL_WAIT_TIME

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in milliseconds (Default = 1000)

Scope: This parameter can be set at the Data Collector level or for a specific host.

Use this parameter to specify the time to wait for the SSH channel to return data. Typically, this is used in large environments with heavy data collection.

SSH_CIPHERS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text (Default is a null value)

Scope: This parameter can be set at the Data Collector level or for a specific host.

Use this parameter to supply explicit ciphers to be used for the session encryption for SSH connections.

SSH_DEBUG

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level or for a specific host.

Use this parameter to enable additional debugging for SSH.

SSH_PORT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value (Default = 22)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS, or for a specific host.

This parameter enables you to change the port used for SSH connections to a non-standard port.

Example:

```
<server>:1234
```

STDOUT_ERROR

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text Value (Default = See below)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS or host.

This parameter enables you to supply a list of messages that could be found in the standard output, which would be considered error messages.

Default messages for this advanced parameter:

```
      command not found|Must have SuperUser privileges to execute
this command|invalid options|No such file or directory|not allowed
to execute|ermission denied|No authority to execute|No Adapters
Found|need to be root|program must be run by root|Failed to open
adapter port|Functionality may be unavailable|function not
supported|cannot execute|Insufficient user privilege|SCSI failure
      <server>:1234
```

STDOUT_FILTER_LIST

Valid Values: Text

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS or host.

This parameter enables you to supply a list of strings that are to be filtered from the STDOUT of a command. The value takes ~~ as a word delimiter and | as a line delimiter.

Example:

```
df:~~failed|error
```

In this example, a line containing (df: **and** failed) **or** error will get filtered from STDOUT.

THREAD_IDLE_TIMEOUT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in seconds (Default = 360)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific host.

Use this parameter to modify the time-out period for data collection thread processing.

WMI_PROXY_VERSION_TIMEOUT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in seconds (Default = 7200000)

Scope: The scope for this parameter is at the Domain or Data Collector level. Domain is required, Data Collector is optional.

Use this parameter to modify the minimum time-out period for WMI Proxy data refresh for the collector.

HISTORIC_TIME_INTERVAL_HOURS

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid values: Specify the range in valid hours, in Integer. (Default value = 12).

Scope: Configure this parameter for NetBackup Historic Data Collection.

For example: The date range is start date12/08/2024 00:00:00 and end date as 14/08/2024 00:00:00.

If the parameter value is 6, then for every 6 hours, IT Analytics will collect the data for the specified date range.

If no value is configured, then the default value 12 will be considered and for every 12 hours the data will be collected.

Cloud data collection advanced parameters

Note: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

This section contains Advanced Parameters for the following:

- See “Amazon Web Services (AWS)” on page 746.
- See “Microsoft Azure” on page 746.

The following advanced parameters apply to a variety of general collection functions, not necessarily associated with a particular subsystem.

Admin > Advanced > Parameters

See "Adding an advanced parameter" on page 734.

Amazon Web Services (AWS)

AWS_BILLING_LOOKBACK_MONTHS

Valid Values: Numeric value in months (Default = 1)

Scope: This parameter can be configured at the Data Collector level.

On the first collection of Amazon Web Services (AWS) billing data, the Data Collector collects one month's worth of data. To override this period, set the AWS_BILLING_LOOKBACK_MONTHS advanced parameter to the number of months of billing history that should be retrieved during the first collection.

Microsoft Azure

AZURE_BACKUP_LOOKBACK_DAYS

Valid Values: A positive integer

Scope: This advanced parameter applies to Microsoft Azure backup data collection.

Set this parameter to define the maximum number of days to go back when collecting for backup jobs. This is applicable the first-time collection begins or if the last collection poll date cannot be determined.

When the data collection last poll date can be determined, and if this date is before the number of days represented by this advanced parameter, then this advanced parameter takes precedence.

For example, if the last poll date was 20 days ago, but the value of AZURE_BACKUP_LOOKBACK_DAYS is 3, only the last 3 days of data will be collected.

AZURE_BACKUP_LOOKBACK_OVERRIDE_DAYS

Valid Values: A positive integer

Scope: This advanced parameter applies to Microsoft Azure backup collection.

AZURE_BACKUP_LOOKBACK_OVERRIDE_DAYS defines an override value to set the maximum number of days to go back when collecting backup jobs. This value takes precedence over all other lookback day settings. If this parameter is

set, both a last poll date and the AZURE_BACKUP_LOOKBACK_DAYS are ignored as values.

AZURE_BILLING_LOOKBACK_DAYS

Valid Values: A positive integer. Default value 30 days.

Scope: For the first time, Azure Billing Collector will collect billing records for the number of days as set by this parameter. If this parameter is not set, Azure assumes 30 days as the default value.

Azure billing collection chooses the start time as the last usage time collected. However, if this start time is more than the number of days (as set by this parameter), AZURE_BILLING_LOOKBACK_DAYS takes the precedence.

Example:

The last usage time collected for billing is Dec 15, 2016, and the current date is Jan 20, 2017.

If the parameter AZURE_BILLING_LOOKBACK_DAYS is set as 7 days, then Azure billing collection will only collect billing information for the past 7 days.

If this parameter is set as 50 days, then Azure billing collection will collect the billing information from Dec 15, 2016, until Jan 20, 2017.

AZURE_BILLING_LOOKBACK_OVERRIDE_DAYS

Valid values: A positive integer.

Scope: This advance parameter applies to Microsoft billing collection.

AZURE_BILLING_LOOKBACK_OVERRIDE_DAYS is a parameter to define the value to override the maximum number of previous days when collecting billing information. This value takes the precedence over all the other look back days configurations.

Note: If this parameter is configured, the values of both last poll date and AZURE_BILLING_LOOKBACK_DAYS is ignored.

AZURE_STORAGE_MAX_RESULTS

Valid values: A positive integer.

Scope: This advanced parameter applies to Microsoft Azure Storage Collection.

Using the AZURE_STORAGE_MAX_RESULTS advance parameter, user can configure the value of `maxResults` parameter in Microsoft Azure REST API call. This API parameter specifies the number of items to return on each page.

Note: Default value is 1500. Also, if the value is configured to greater than 5000, pagination is up to 5000.

Host discovery and collection advanced parameters

Note: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

The following advanced parameters apply to a variety of general collection functions, not necessarily associated with a particular subsystem.

Admin > Advanced > Parameters

See "Adding an advanced parameter" on page 734.

HOST_COLLECTION_ORACLE_CONN_STR

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: `jdbc:oracle:thin:@//localhost:${port}/${sid}`.

Scope: The Oracle probe was designed only for 12c Oracle database. Hence, for 19c Oracle with the user system, it can only connect to the container and not the pluggable database since the connection string used was of 12 oracle db.

APP_DISK_PATH

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text value (Default is a null value)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS, or for a specific host.

Applications, such as Oracle ASM, may have devices created from disk aliases located in any directory. This parameter enables you to provide a list of directories that will contain the alias devices that are of interest.

A colon-separated list of paths for disk locations provides a mapping to non-standard locations.

Example: APP_DISK_PATH=**/dev/oracleasm/disks:/dev/raw**

HBAVERIFY_NUM_DELAY_SECS

Valid Values: Integer (in seconds)

Scope: This parameter can be set at the Data Collector level.

HBAverify command output varies with every run and may or may not provide all the Physical Disk drives. The command will be executed until all the LUN-derived disks are found. The logic excludes Direct Attached Storage disks while running hbaverify. This advanced parameter can be set to configure the number of extra runs for the command and the delay between the subsequent execution.

HOST_CONNECTION_PROTOCOL

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Text value (Default = **WMI|SSH**)

Scope: This parameter can be set at the Data Collector level only.

Use this parameter to control which protocol is used for the host connection.

HOST_CONNECTION_TIMEOUT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in seconds (Default = 60 seconds)

Scope: This parameter can be set at the Data Collector level only.

Use this parameter to set the period of time allowed before the host connection times out.

HOST_SOCKET_TIMEOUT

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Numeric value in seconds (Default = 180 seconds)

Scope: This parameter can be set the socket time-out for connections to hosts, for an OS, or for a specific host.

Use this parameter to update the command time-out for Windows servers for Windows host data collection. A corresponding time-out must be configured for the WMI Proxy server.

To set the time-out for the WMI Proxy server:

1. Edit the file:

   ```
   <HOME>\APTAREWMIServer\conf\aptarewmiserver.properties
   ```

2. Modify or add this line to set the value in seconds: **command_timeout=<value>**
3. Restart the WMI Proxy server.

Example:

To set the time-out to 10 minutes, set the HOST_SOCKECT_TIMEOUT advanced parameter to 600 and modify the WMI Proxy file and set **command_timeout=600.** Then, restart the WMI Proxy server.

IGNORE_FILE_SYSTEMS

Use this parameter to filter a list of file systems for host resources data collection.

Valid Values: Example: **/proc|/ahafs|/testfs**

Scope: This parameter can be set at the Domain or Data Collector level.

This advanced parameter can be used to avoid reporting a data collection error for the File System category of a Host Resource collection, when a file system does not have a capacity value in the output of the **df -k**command.

More than one file system can be specified. In this case, file systems need to be separated by a pipe character (|).

IS_EMULEX_MAPPING_STRICT_CHECKING

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Domain or Data Collector level or for a specific OS, or for a specific host.

Emulex HBA mapping in Windows environments requires the SCSI bus-LUN-target mapping to strictly match the physical drive. Sometimes there may be a data collection issue where the LUNs are mapping to the wrong device names. This can happen if there are multi-paths using the same target and SCSI OS LUN in the output.

Use this parameter to turn on strict checking of the mapping.

NUM_HBAVERIFY_RUN

Valid Values: Integer

Scope: This parameter can be set at the Data Collector level.

HBAverify command output varies with every run and may or may not provide all the Physical Disk drives. The command will be executed until all the LUN-derived disks are found. The logic excludes Direct Attached Storage disks while running hbaverify. This advanced parameter can be set to configure the number of extra runs for the command and the delay between the subsequent execution.

STRICT_HOST_KEY_CHECKING

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Domain or Data Collector level only.

Use this parameter in the context of connecting to hosts through SSH via the use of private keys (that is, no password authentication).

- When the value is N, the “known hosts” file will automatically be updated when the Data Collector connects to host for the first time, or when the host’s key has changed.
- When the value is Y, the “known hosts” file will have to be manually maintained.

USE_VIO_RESTRICTED_SHELL

Valid Values: TRUE or FALSE (Default = FALSE) - Note that this is a case-insensitive value.

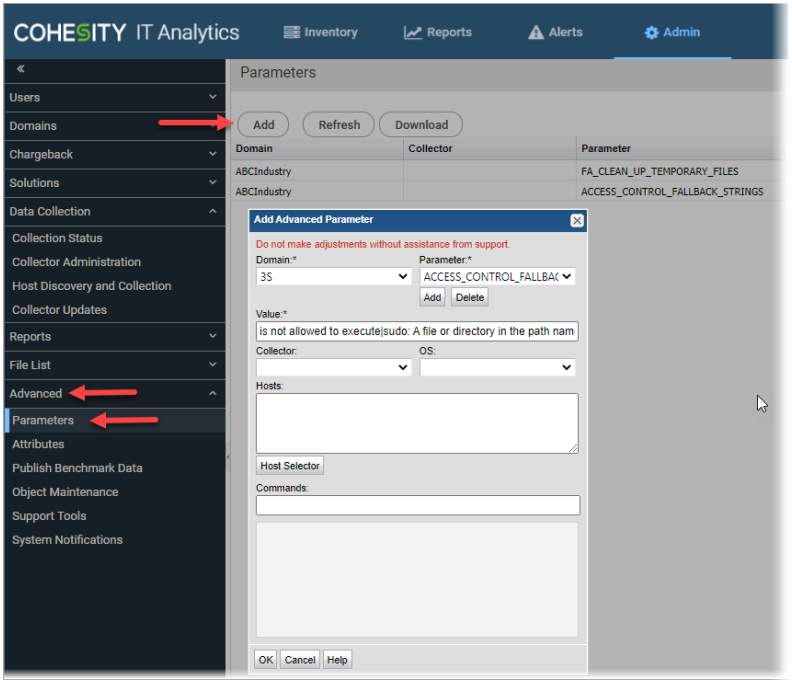
Scope: This parameter can be set at the Data Collector level.

Collection of host resources capacity data from IBM VIO/AIX servers can now use the VIO restricted shell to execute VIO commands, as supported by IBM. This collection method is enabled via this advanced parameter.

Backup Manager advanced parameters

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

To configure advance parameter, navigate to **Admin > Advanced > Parameters.**. Advance parameter page is displayed. Click "+Add" icon to add a new advance parameter.



Click Add to configure new parameter. Add Parameter definition dialog box is displayed.

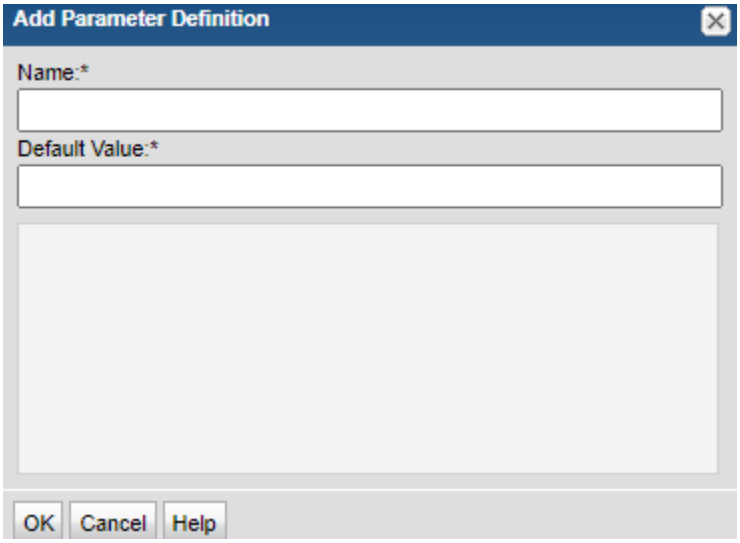


Table 22-1 Advance parameters for data collections.

Parameter	Permitted values	Default value	Description
ENABLE_MINUS_T_OPTION	Y / N	Y	This parameter can be set at the Data Collector level. This advanced parameter applies to NetBackup data collection. Based on advanced parameter value, the schedule for the job details probe in Netbackup policy is determined as follows : ■ If the value is set to 'Y', then the schedule is 5 mins. ■ If the value is set to 'N', then the schedule is 35 mins. The schedule configuration is displayed in Veritas NetBackup Data Collector >> Active Probes >> Job Details option. Set this parameter to Y to indicate "Job Details" probe to collect only those NetBackup jobs which were modified/created since the last jobs collection. **Note:** Setting this parameter to Y can potentially reduce CPU load as well as improve job collection performance.

Table 22-1 Advance parameters for data collections. *(continued)*

Parameter	Permitted values	Default value	Description
NBSTL_UTIL_WITH_U_OPTION	Y / N	N	Advanced parameter to control execution of command with and without -U. This parameter can be set at the Data Collector level for NetBackup SLP Job details probe. Setting the parameter value to 'N', the nbstlutil command, which is used for SLP Job Details Probe, is executed without -U option. **Note:** Parameter value with option 'N' improves performance. If the parameter value is set to 'Y', then the same command will be executed with -U option.
NBU_ENABLE_HISTORICAL_INCOMPLETE_SLP_JOB_COLLECTION	Y / N	Y	Set this parameter at the Data Collector level. This parameter applies to the SLP Job Details probe of the Veritas NetBackup Data Collector policy. Set this parameter to **Y** to let the SLP Job Details probe collect historical incomplete SLP jobs. Historical collection refreshes images and copies that remain incomplete in IT Analytics after they fall outside the standard SLP lookback window. Set this parameter to **N** to restrict the probe to the standard SLP lookback window.
NBU_HISTORICAL_INCOMPLETE_SLP_JOB_COLLECTION_INTERVAL_DAYS	A positive integer, in days.	7	Set this parameter at the Data Collector level. This parameter defines how often the SLP Job Details probe runs historical incomplete SLP job collection. IT Analytics stores the date of the last historical collection. The probe starts the next historical collection only after the configured number of days has passed. This parameter does not affect regular SLP job collection, which continues to run on the probe schedule. This parameter takes effect only when NBU_ENABLE_HISTORICAL_INCOMPLETE_SLP_JOB_COLLECTION is set to **Y**.

See “Adding an advanced parameter” on page 734.

COMMVAULT_OLDEST_JOB_HOURS

Valid Values: 0 or higher (Default = 0 to force collection of all jobs)

Scope: This parameter can be set at the Data Collector level.

This advanced parameter applies to CommVault Simpana data collection. Specifies how many hours to retrieve historical jobs on the first collection only. If you only want to collect jobs from the last week, you don't need this advanced parameter. You can also specify positive values for this number of hours; for example, 168 = 1 week.

MMINFO_MOVE_BACKWARD_MIN

Valid Values: 0 or higher (Default = 30)

Scope: This parameter can be set at the Data Collector level.

This advanced parameter applies to NetWorker data collection. The data collection event thread runs every 10 minutes to get all the finished jobs that have been created after the last jobs that were captured during the last collection cycle, plus 30 minutes of additional time to account for long-running jobs.

For long-running backups, it may be necessary to configure the MMINFO_MOVE_BACKWARD_MIN Advanced Parameter to ensure that all NetWorker savesets are collected successfully. The default value is 30 minutes, but for long-running backups, that time may need to be adjusted because 30 minutes may not be long enough to capture all the jobs. For example, some jobs can last several hours.

Note that if a value of less than 30 minutes is configured, the Data Collector uses the default 30 minutes.

NBU_AUDIT_LOOKBACK_DAYS

Valid values: A positive integer to indicate days (Default=3)

This parameter is used only for the first time collection of NetBackup Audit events and by default, events from last 3 days will be collected for the first time. Change the value of this advanced parameter to collect events that are anything other than 3 days.

NBU_DUP_JOB_PARTIAL_STATUS_OVERRIDE

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

This advanced parameter applies to NetBackup data collection. Set this parameter to Y to set NetBackup backup job statuses to the associated duplication job status. Note that this can cause issues when the Duplication job is duplicating multiple backup jobs and some of these jobs fail (causing all backup job statuses to be set to 1 - Partial success).

NO_AVAMAR_CLI

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

This advanced parameter applies to Avamar data collection. Due to security restrictions, some customers do not want to run command-line interface (CLI) commands on their Avamar servers. To turn off CLI commands, set this parameter to Y. This should be set on a per-collector basis.

RMAN_BACKUP_LOOKBACK_DAYS

Valid Values: A positive integer. (Default = 1)

Scope: During the first collection, the RMAN collector will collect session and backup details for the number of days as set by this parameter. If this parameter is not set, the default value is 1 day.

The RMAN collector chooses the start time as the last "session time" collected for. However, if this start time is more than the number of days (as set by this parameter), RMAN_BACKUP_LOOKBACK_DAYS takes the precedence.

Example:

The last "session time" collected for RMAN is Feb 5, 2018 and the current date is Mar 12, 2018.

If this parameter is set as 5 days, then the RMAN collector will collect the session and backup information only for the past 5 days.

If this parameter is set as 50 days, then RMAN collector will collect the session and backup information from Feb 5, 2018 until Mar 12, 2018.

RMAN_BACKUP_LOOKBACK_OVERRIDE

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

When this advanced parameter is set to "Y" the value set for RMAN_BACKUP_LOOKBACK_DAYS is used ignoring any previous collection date. This can be used to collect historical data.

After running a collection with this advanced parameter set to "Y" ensure that it is reset back to "N" to avoid excessively-long collection cycles.

SUPPRESS_KERBEROS_PROMPT

Valid Values: Y or N (Default = Y)

Scope: This parameter can be set at the Data Collector level.

Data collection may trigger a Kerberos authentication prompt that can cause the collector to hang until the authentication passes.

USE_ALT_NBU_INCL_EXCL

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

This advanced parameter can be configured for collection of NetBackup include/exclude lists from Unix clients. By default, the collector uses the NetBackup-recommended command syntax to retrieve the lists. If the lists are not collected successfully, set this advanced parameter to Y, which instructs the collector to use an alternative command syntax for list data retrieval from Unix clients.

USE_NTML_V2

Note: This advanced parameter was deprecated in release version 10.1. If it was configured in a previous software version, it is now ignored as it is no longer required to enable NTMLv2 authentication.

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

This advanced parameter applies to Backup Exec data collection. When set to Y, the Backup Exec Data Collector uses NTMLv2 authentication.

VEEAM_BACKUP_LOOKBACK_DAYS

Valid Values: A positive integer. (Default = 1)

Scope: During the first collection, the Veeam collector will collect session and backup details for the number of days as set by this parameter. If this parameter is not set, the default value is 1 day.

The Veeam collector chooses the start time as the last "session time" collected for. However, if this start time is more than the number of days (as set by this parameter), VEEAM_BACKUP_LOOKBACK_DAYS takes the precedence.

Example:

The last "session time" collected for Veeam is Feb 5, 2018 and the current date is Mar 12, 2018.

If this parameter is set as 5 days, then the Veeam collector will collect the session and backup information only for the past 5 days.

If this parameter is set as 50 days, then Veeam collector will collect the session and backup information from Feb 5, 2018 until Mar 12, 2018.

VEEAM_BACKUP_LOOKBACK_DAYS_OVERRIDE

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

When this advanced parameter is set to "Y" the value set for VEEAM_BACKUP_LOOKBACK_DAYS is used ignoring any previous collection date. This can be used to collect historical data.

After running a collection with this advanced parameter set to "Y" ensure that it is reset back to "N" to avoid historic data collection on future scheduled or On Demand runs.

Capacity Manager advanced parameters

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

Warning: The following advanced parameters apply to Capacity Manager data collection and can be added using:

Admin > Advanced > Parameters

See “Adding an advanced parameter” on page 734.

EMC_MULTIPATH_OFF

This parameter is shipped in the Portal software and is listed in the Add Advanced Parameter window.

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

Use this parameter to turn off EMC PowerPath Multi-Pathing Collection. Servers can have EMC PowerPath and Device Mapper multi-paths at the same time on Linux. If your environment has file systems on Device Mapper multi-pathed disks, it is required to turn off EMC PowerPath.

EMC_VMAX_INCLUDE_ALL_MASKINGVIEWS

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

Use this parameter for EMC Symmetrix VMAX array collection when the path between the initiator and target are missing. Data Collector parsers which construct the storage path use the command:

```
symaccess -sid <sid>list view -v -detail
```

When the Data Collector parser constructs the storage path, it assumes that masking view names are unique, so it does not process information from 2nd instance of the same view name. This causes the path to be ignored. Set this advanced parameter to have the parser processes paths from all instance of masking views.

Set this parameter only when

```
symaccess -sid <sid>list view -v -detail
```

is showing multiple instances of the masking view.

HDS_HTM_JPCCTRL_PATH

Valid Values: Text string

Scope: This parameter can be set at the Domain and Data Collector levels.

Used for Hitachi Array performance collection (Hitachi Tuning Manager).

It defines the relative path for executing the JPCCTRL command. This is required when the executables for **jpcctrl** and **jpcrdef** are at locations that are different from what is defined in the Data Collector policy.

For example, if the value for HDS_HTM_JPCCTRL_PATH is set to **/opt**, the command path for jpcctrl will be: **/opt/jp1pc/tools/jpcctrl**

To use this parameter, take the steps similar to this example:

1. Identify the location of the **jpcrdef** and **jpcrpt** files/commands (HTnM location). For example, if these two files (commands) exist in **C:\ProgramFiles\HiCommand\TuningManager\PerformanceReporter\tools\,** then the HTnM install location would be **C:\ProgramFiles(x86)\HiCommand\TuningManager**.
2. Check if the **jpcctrl** file (command) exists in the location: **HTNM_LOCATION\jp1pc\tools\.** If it does, then the Data Collector should run without any issues. If it does not, create the HDS_HTM_JPCCTRL_PATH Advanced Parameter with a value of: **C:\ProgramFiles(x86)\HiCommand\TuningManager**

REMOTE_RAID_AGENT

Valid Values: Y or N (Default = N)

Scope: This parameter can be set at the Data Collector level.

Used for Hitachi array performance collection (Hitachi Tuning Manager). Configure this advanced parameter when there are distributed RAID agents (remote RAID agents).

If the value is set to Y, the Data Collector displays the configuration and status of Collection Manager and Agent services on all the remote hosts. Set the REMOTE_RAID_AGENT parameter to Y when we want to get the status of the agent services on all the hosts.

If no value is set, or the value is set to N, the Data Collector checks the operating status of all services on the local host of the Tuning Manager system.

SYMSTAT_COUNT

Valid Values: Numeric (Default = See below)

Scope: This parameter can be set at the Data Collector level.

Used for EMC Symmetrix array performance collection.

Use this advanced parameter to gather more granular performance data. This parameter sets the count of data sets for the custom interval. Use this in conjunction with the following.

See "SYMSTAT_INTERVAL_SECONDS" on page 761.

Default Values:

If these advanced parameters are not set, then the defaults are set to:

- count = 1

- interval is derived from the polling frequency of the performance probe.

SYMSTAT_INTERVAL_SECONDS

Valid Values: Numeric in seconds (Default = See below)

Scope: This parameter can be set at the Data Collector level.

Used for EMC Symmetrix array performance collection.

Use this advanced parameter to gather more granular performance data. This parameter customizes the data collection interval. Use this in conjunction with the following.

See "SYMSTAT_COUNT" on page 760.

Default Values:

If these advanced parameters are not set, then the defaults are set to:

- count = 1
- interval is derived from the polling frequency of the performance probe.

The reason for using CLIENT is that in some cases the CIM server appears to not return all requested relationships. Using CLIENT will require more disk space on the collector (but not excessive - 10s or 100s of MB) and will tend to take longer before persisting at the start of collection, but less time at the end.

File Analytics advanced parameters

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for help.

The following advanced parameters apply to File Analytics data collection and can be added using: **Admin > Advanced > Parameters**

See "Adding an advanced parameter" on page 734.

Note: All Windows hosts require Local or System Administrator privileges for WMI.

FA_CLEAN_UP_TEMPORARY_FILES

Valid Values: Y or N (case-insensitive)

This advanced parameter controls the deletion/retention of work files created by each File Analytics data collector. If the parameter is not configured, the default behavior is to clean up the temporary files.

Note: This parameter should be set only when directed by Cohesity Support.

AFS_PATH_EXCLUDE_DIRECTORY_NAME

Valid Values: Enter text to specify the directory to be excluded from data collection.

This advanced parameter can be used to exclude certain folders from File Analytics data collection. For example, you may not want to collect snapshot directories, as they contain copies of replicated folders. To exclude folders, take the following steps.

- Add the advanced parameter, AFS_PATH_EXCLUDE_DIRECTORY_NAME, and for the value of the parameter, specify the directory to be excluded.

FA_HOST_CAPTURE_REMOTE_SHARES

Valid Values: Y or N (Default = N)

- If this advanced parameter is set to Y along with a target server host name, the File Analytics collector logs a statement in Metadata_HostResource_10.log stating "FA Collection Capture Remote Shares: true on Host: <server host name>".
- If this advanced parameter is not set, this log statement will not appear in the log and the File Analytics collector excludes from collection any remote shares attached to the target host.

FA_HOST_DIRECTIVE

This advanced parameter is for File Analytics collection initiated through Host Inventory for Linux hosts. FA_HOST_DIRECTIVE allows the user to associate it with the target host and specify the following:

[scan_root_dirs={dir1 to scan}[~[dirs to exclude]][:dir n to scan]] [capture_remote_shares={Y/N} [root_scan_depth={1..10}]

Value: scan_root_dirs=/opt/aptare~/inode2*:/filetest/rnddirfiles capture_remote_shares=Y root_scan_depth=3

- scan_root_dirs - By default the root directory is “/”, with this directive, one can override the root directory. You can specify one or more root directories (separated by :), helpful if all you are interested in are specific file share

directories.You can specify one or more excludes for each of the root directory. ~ separates the excludes from the root, and each exclude is separated by ","

- capture_remote_shares - By default it is N, when set to Y the code will capture remote shares.
- root_scan_depth - The default is 1, but the recommendation is to set it to 3 or 5 if scanning a large filesystem.

Virtualization Manager advanced parameters

Warning: Making changes to Advanced Parameters should not be undertaken unless the user understands the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Supportor visit the Community web site for further guidance.

The following advanced parameters apply to Virtualization Manager data collection and can be added using:

Admin > Advanced > Parameters

See "Adding an advanced parameter" on page 734.

VMWARE_COLLECT_UNUSED_PROPERTIES

Valid Values: Y or N (Default = N)

Set to Y to collect all data for objects retrieved from VMware. This is used to collect raw data suitable for analysis for new features and enhancements.

VMWARE_DATASTORE_TIMEOUT_SECONDS

Valid Values: 0 or higher (Default = 0)

Number of seconds after which to time out datastore listings. The default is no time-out, but VMware will time out in 30 minutes irrespective of the value that is set for this parameter. This parameter only applies to the datastores data collection probe.

Note: This parameter applies per directory in the datastore. The collection will stop scanning a datastore as soon as a single directory times out. Any files collected before the time-out will be persisted.

Fabric Manager Advance Parameters

Observe caution when you make changes to Advanced Parameters. Make sure you have assessed the impact of the parameter changes. If undesired results occur, revert the settings and contact Cohesity Support or visit the Community web site for further guidance.

To configure theFabric Manager Advance Parameters for Data Collection, navigate to **Admin > Advanced > Parameters.**

CISCO_REST_CONNECTION_TIMEOUT

Valid Values: Positive integer, seconds (Default = 60 if the parameter is unset, not an integer, or ≤ 0)

Scope: Maximum time to establish the HTTP connection to the Cisco device for REST (NX-API) collection. Values are converted to milliseconds internally. Increase if connections is slow or distant networks often time out while connecting.

CISCO_REST_READ_TIMEOUT

Valid Values: Positive integer, seconds (Default = 300 if the parameter is unset, not an integer, or ≤ 0)

Scope: Maximum time to wait for a response after the connection is established for REST (NX-API) collection. Values are converted to milliseconds internally. Increase for slow CLI-show responses or large outputs that exceed the default wait.

CISCO_SWITCH_DISCOVERY_VIA_CREDENTIALS_OVERRIDE

Valid Values: Y or N (Default = N)

Scope: Set to **Y** to treat each switch listed in the credentials override additional parameter as a discovery target for the Cisco Switch REST collector, in addition to the primary server address. When N or unset, override entries are used only for credentials when connecting to switches already found through topology; they do not add new switches to the discovery queue. Only the uppercase Y is recognized.

Chapter 23

Manage your Portal environment

This chapter includes the following topics:

- Root folder and domains
- Add/Configure a domain
- Configure alerting for a domain
- Community performance statistics for a domain
- Configure the domain password policy
- Manage your profile and set a language preference
- Change your login password
- View Downloads

Root folder and domains

During installation, a Root Folder and a Domain is set up to support the host group hierarchy structure.

Root folder

By default, a root folder is set up. This folder has a unique group identifier (300000). You can rename this host group, but you cannot change its ID. As you create host groups, the Portal adds them into the hierarchy, in the position that you specify.

Domain

A Domain provides a way to “partition” the reporting database into separate, private realms. It is used primarily to implement security controls for multi-tenancy systems.

See “Multiple domains (for managed services partners only)” on page 766.

A Domain is a unique entity associated with the top level of your host group hierarchy. The domain name is supplied during the installation process and the Portal assigns it to the root folder.

The domain is used by the Data Collector for:

- Authentication: The primary server record must exist in the host group hierarchy for the domain. (Veritas NetBackup only)
- Host Searches: The Data Collector searches the domain’s host group hierarchy checking for hosts associated with the backup data it is gathering. If no host is found, a new host is added to the root-level host group folder for the domain.

Enterprise environments typically have only one domain. When you add (or delete) host groups or attributes, you do so globally for your domain and all the host groups in that domain. Unless you are a Managed Services Provider (MSP), you should not have to specify a domain when you add or delete host groups or attributes through the Portal.

Multiple domains (for managed services partners only)

If you are a Managed Services Partner, you need the capability of managing multiple, independent hierarchies--one for each of your client companies. As an MSP, you will define a unique domain for each of your customers. When you add or delete attributes, you can do so for all domains or you can select specific domains to apply changes.

A domain is associated with a host group hierarchy and all newly discovered hosts are added to the root host group associated with the domain. Each MSP customer will have a separate domain with its own hierarchy.

Note: A host group can only serve as the root for one domain. For example, you could have a host group defined for Acme Corp and then create an Acme Domain that uses the host group as the root of its host group hierarchy. Once a domain is associated with a host group, this host group cannot become the root of any other domain.

Add/Configure a domain

A domain identifies the top level of your host group hierarchy.

If you are a Managed Services Partner, you will create a domain for each of your customers. Each domain must be associated with a host group that serves as the root of the customer's host group hierarchy and a license that governs the domain's entitlement. In this way, you can manage each customer environment separately.

See "Root folder and domains" on page 765.

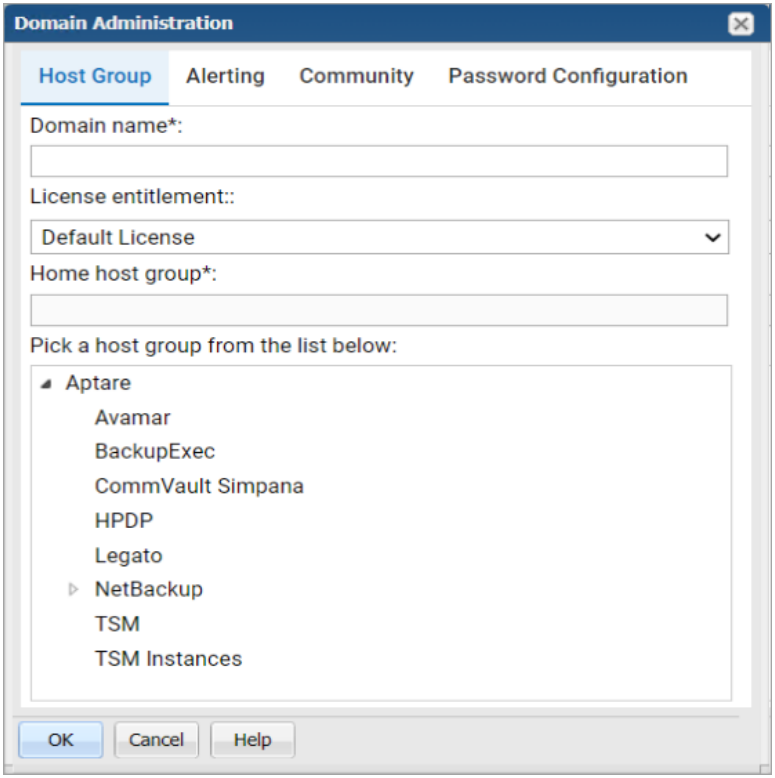
To create a new domain

1 Choose **Admin** > **Domains** > **Domains**.

The **Domains** page will display the configured domain with its respective host group.

2 Click **Add** to display the **Domain Administration** window where you can configure the Domain's home host group, Alerting (SNMP trap details), and community performance statistics for a domain.

See "Community performance statistics for a domain" on page 769.



3 In the **Domain Administration** window:

- Enter a domain name.
- Select a license from **License entitlement**. The licenses displayed depend on the license set for the parent domain and its entitlements become applicable to the domain and also to the users or user groups assigned to the domain.

  A confirmation prompt is displayed if you change the license of a sub-domain. Also, you cannot change the license of the top-level domain.
- Expand the host group list, then click on a host group that will be the home host group for this domain--that is, the host group that must be the root of the domain's hierarchy.
- Click **OK** to create the domain.

Note: A host group can serve as the root for only one domain. For example, you could have an Acme domain and then assign the Acme Corp host group to this domain. Once a domain is associated with a host group, this host group cannot become the root of any other domain.

Additional references:

- See "Configure alerting for a domain" on page 768.
- See "Community performance statistics for a domain" on page 769.
-

Configure alerting for a domain

The over arching rules for alerting are defined at the domain level.

To configure Alerting at the domain level

1. Select **Admin** > **Domains** > **Domains**.
2. On the **Domain Administration** window, click the **Alerting** tab.
3. Enter the **Port**, **Community**, and **Management servers** to be used as the default values when an SNMP Alert is configured for a saved report.
4. Configure the check boxes for the delivery methods that you want to be associated with this domain. By default, only email is enabled for all domains. Email is not an optional setting shown in the Domain Administration list.

   The following methods must be granted permission at the domain level:

   - Script

- SNMP - Allows you to route the alert-based and report-based notification through a SNMP server.
  When this box is selected, you can select any one or both from **Legacy** and **Enhanced** as SNMP payload format. For any new domain, select **Enhanced**. **Legacy** is for backward compatibility.
- Native Log - When this box is checked, a log entry is written to the OS-specific log: either the Windows event log or the Linux syslog.
- Syslog - Allows you to route alert-based and report-based notifications through a Syslog server. This permission is essential to enable Syslog notifications for report-based alerts.
  See "Syslog configuration" on page 487.

5 Click **OK** to save the configuration.

Community performance statistics for a domain

To configure community participation at the domain level

Before you can share array performance statistics with similarly configured arrays in the broader community to gauge your environment's performance, you must enable Community participation.

Performance profiles are securely transmitted (over https) as anonymous and aggregated with other customers' profile data in Profile Central, which is then imported into a customer's profile for reporting purposes. This import/export task occurs in a single, daily scheduled Portal process. Using the aggregated community profiles, companies can better gauge if the metrics collected in their environments are within a normal performance range. Profile data cannot be associated with any contributor. No company or environment-specific details, such as storage array or hosts, are transmitted. No personally identifiable information is collected, used, or disclosed.

The Administrator can enable participation in two different ways--either will enable participation for the arrays in the specified Domain:

- Community Performance Statistics: This enables participation in the Cloud community for array statistics sharing for the Domain associated with your User ID and Home Host Group.
- Configure Community Performance Profiling: This enables a multi-domain organization (such as a Managed Services Provider) to configure participation on a Domain-by-Domain basis, selecting only the Domains that are authorized to participate.

IMPORTANT: To enable participation in Community Performance Profiling Cloud Policies, an authorized representative of your company must opt-in.

1. Select **Admin > Domains > Domains**.
2. Select a domain name.
3. Click **Edit.**
4. Click the **Community** tab in the **Domain Administration** window.
5. Click **Participate in community program**, then click **OK**.

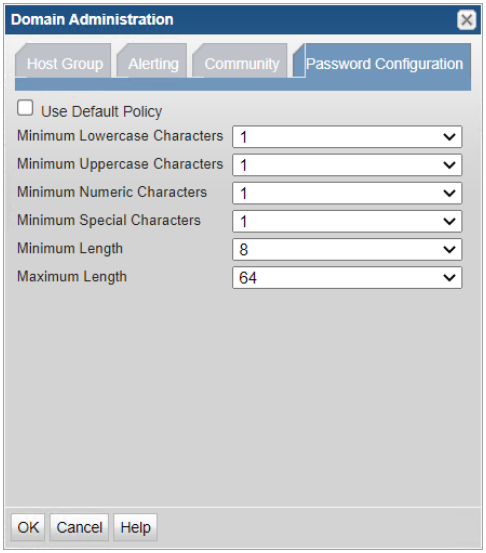
Configure the domain password policy

This procedure allows you to configure a domain password policy. In IT Analytics Portal version 11.1 and later, user password is governed by the default settings of the password policy mentioned in the table below. You can configure these settings further to set a custom password policy based on your organization guidelines.

To configure the domain password policy:

1 Select **Admin** > **Domains** > **Domains**.

2 On the **Domain Administration** window, click the **Password Configuration** tab.



3 Specify the field values based on the descriptions in the table below and click**OK**. Clear the **Use Default Policy** checkbox to enable the configuration.

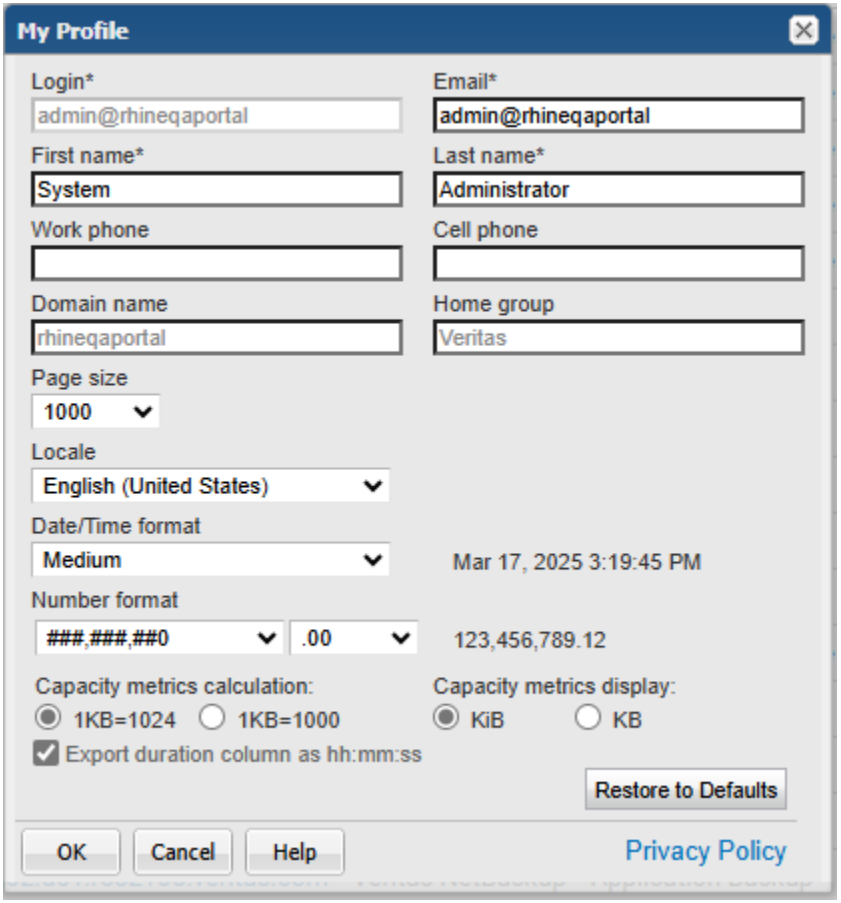
Table 23-1 Configure the domain password policy

Field	Range	Default value	Description
Use Default Policy	None	Selected	When selected, applies default values to the password policy. To override the values of the parent domain, uncheck this box and set the custom values. If this field already has custom values, selecting this field will reset all the fields to their default values.
Minimum Lowercase Characters	1 to 5	1	Defines the minimum number of lowercase characters required in the password.
Minimum Uppercase Characters	1 to 5	1	Defines the minimum number of uppercase characters required in the password.
Minimum Numeric Characters	1 to 5	1	Defines the minimum number of numeric characters required in the password.
Minimum Special Characters	1 to 5	1	Defines the minimum number of special characters required in the password. Acceptable characters: @#$%^&+=!
Minimum Length	8 to 15	8	Defines the minimum number of characters required in the password.
Maximum Length	16 to 64	64	Defines the maximum allowable characters in the password.

All the existing passwords before the current update will be honored. Any future attempt to change the password will have to honor the password policies defined for the user domain. Also, users whose credentials are reset by the administrator will be required to again set their passwords. Lastly, while resetting passwords through the portal, users will have to set passwords that comply with the policy configured using the above procedure.

Manage your profile and set a language preference

While the Application Administrator has the responsibility of setting up login access, you can update other aspects of your profile.



To update your profile

1. Click User Account menu and select **My Profile**.
2. Modify any of the following fields:
   - Login - This field is required.
   - Email - This field is required.

- First name - This field is required.
- Last name - This field is required.
- Work phone
- Cell phone

3 Review the read-only fields:

- Domain name - Displays the domain name associated with the user account.
- Home group - Displays the home group for the user account.

4 Select your preferences for the UI and reports:

- Locale - In general, this field defines the language for the display date.
  If language support is available for a Locale, select your language preference. Selecting the Locale switches the Portal to display in that language after logging out and logging in again. Help is also displayed in the selected language.
  If the selected locale is not supported, English is displayed in the UI, but other locale settings are honored, for example the date format. Language of locale controls user language irrespective of region. **Restore Defaults** returns the UI to the initial locale. If browser settings do not specify a region, the default is en_US.
- Date/Time - Select from:
  - Short - for example: 6/12/13 9:57 AM
  - Medium - for example: Jun 12, 2013 10:01:01 AM
- Number format - Defines the decimal/comma placement for numbers.
- Capacity metrics calculation - When 1024 is used for calculations, it refers to binary multiples: kibibyte (KiB), mebibyte (MiB), gibibyte (GiB), tebibyte (TiB), and pebibyte (PiB). When 1000 is used for calculations, it refers to decimal multiples: kilobyte (KB), megabyte (MB), gigabyte (GB), terabyte (TB), and petabyte (PB). By default, the units for values in reports are displayed as binary multiples. This default setting can be changed by modifying a user's profile.
- Capacity metrics display - Select the label to represent the unit of measure. This label is displayed in reports.
- **Export duration column as hh:mm:ss** - This check box is not visible by default. An administrator can configure **portal.enableDurationColumnExportOption=true** parameter in portal.properties file (/opt/aptare/portalconf) and restart the tomcat-portal

services. Once available, select this option to export the reports in Excel format. The **Duration** column in the report is exported in the selected format.

Note: You can always click **Restore to Defaults** if required.

Change your login password

If you forget your password, you can click the **Forgot your password?** link on the login page. Your password is immediately reset and the new password is sent to you in an email. You can change your password once you login. You are allowed five attempts at a password before the account is locked. An administrator must reactivate the account before you can login again.

To assign or change a user password:

1. Click the user account menu and select **Change Password**.
2. In the **Change Password** dialog, enter the old password, the new password and confirm it. Click **OK** to update the password.

View Downloads

The **Downloads** item form the user profile menu opens a Downloads list on the portal. The Downloads list displays all the downloads and exports performed during the browser session. You can expand, minimize, or close the **Downloads** list while your are working on the Portal. The list of downloads is cleared if you refresh or close the portal browser.

You can open the **Downloads** list from User Profile menu > **Downloads**.

Chapter 24

Manage ransomware scorecard

This chapter includes the following topics:

- Ransomware Scorecard overview
- Ransomware administration
- Answer a ransomware query
- Override a data query
- Add a ransomware question
- Utility to add Ransomware Scorecard questions and answers in bulk
- Edit a ransomware question
- Add ransomware data query
- Edit ransomware data query

Ransomware Scorecard overview

The Ransomware Scorecard provides a view of the preparedness of your environment with respect to ransomware resilience and recoverability. In addition, it also recommends improvements and best practices that ensure strong ransomware defense and accurate data recovery. The scorecard displays its results based on the data collected from your environment and the user responses to the evaluation queries present on the scorecard. You can also add custom queries if the default set of queries do not adequately evaluate your environment. Out-of-the-box, the scorecard derives a set of data points only from NetBackup, plus the ability to add your own custom reports for an all-round evaluation of the environment.

Quick start for Ransomware Scorecard

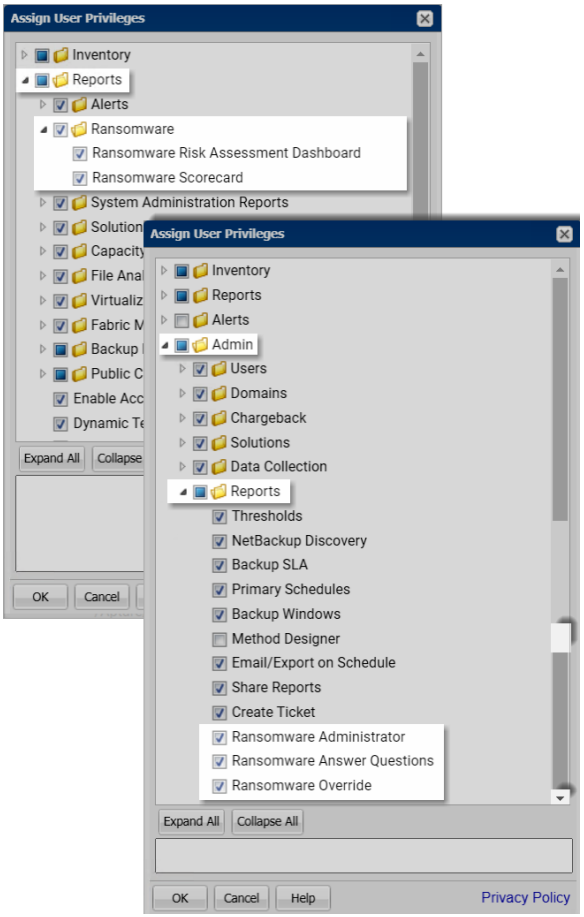
Follow these recommendations to ensure you get a realistic evaluation of your environment with respect to ransomware resilience and recoverability:

- The higher the number of responses and data points, the more realistic is the evaluation.
- Since the scorecard derives a set of data points exclusively from NetBackup out-of-the-box, ensure Cohesity NetBackup Data Collector policy is configured and you have a data collection at least for a month.
- Run the Ransomware Scorecard from the Ransomware folder of the **Reports** tab.
- While setting the scope for the scorecard, choose the correct domain and select **All** in the other scope fields.
- Initially, the Ransomware Score and Complete Queries (%) are zero. The scores start showing up as data stats are received from reports and as users post responses to the queries.
- Sort the **Risk** column to identify the high-risk items. (The default sort order is highest risk at the top.)

Ransomware Scorecard access permissions and privileges

The role-based access permissions and privileges available to different portal users of the Ransomware Scorecard are described below.

An administrator can set privileges and permissions for other users for the Ransomware Scorecard from **Admin** tab > **Users** > **Users and Privileges** > **Privileges**. While the Ransomware dashboard and scorecard access is controlled through **Reports** > **Ransomware**, the privileges to manage, answer, or override scorecard queries are assigned from **Admin** > **Reports**.



See “Assigning user privileges” on page 578.

Table 24-1 Scorecard Roles

Role	Description
Ransomware Administrator or Super User	A Ransomware Administrator is expected to configure the scorecard for the organization. This role can add custom queries and disable or enable some of the built-in queries.
User	A User can answer the queries on the scorecard and override Data type queries. Since these are two separate privilege levels, a portal administrator must enable these privileges for the user separately to allow the user to perform these actions. Both have separate privilege levels but collectively impact the total score.

Table 24-1 Scorecard Roles *(continued)*

Role	Description
Viewer	A viewer can simply view the Ransomware Scorecard and share it outside IT Analytics Portal through email, as HTML, PDF or other export format. Typically, a viewer is someone who is primarily interested in the results and recommendations of the scorecard.

The above-mentioned roles translate to privileges and/or restrictions to perform the following actions. These actions are available through the **Actions** menu on each row of the scorecard.

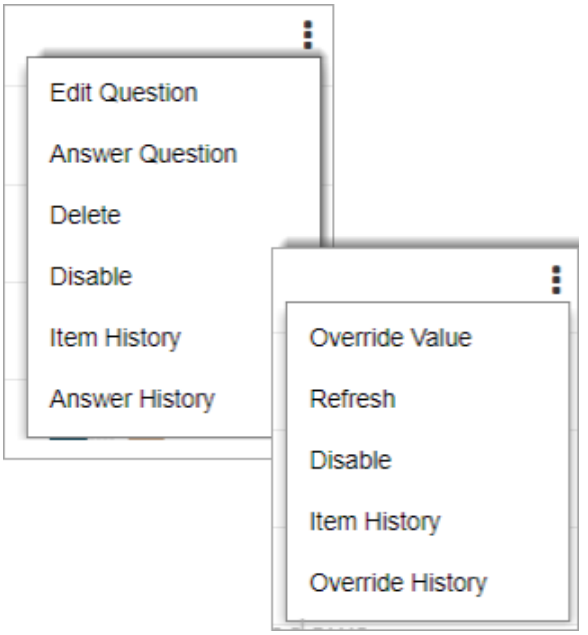


Table 24-2 Role-based user privileges

User actions	Description	User privileges		
		Ransomware Administrator (Y/N)	**Ransomware Answer Questions (Y/N)**	**Ransomware Override (Y/N)**
Add question Add data query	Add a new query.	Y	N	N
Edit question Edit data query	Edit only the respective user-created query or question.	Y	N	N
Delete	Delete only the respective user-created query or question.	Y	N	N
Disable/Enable	Toggles between exclusion and inclusion of the data query or question in the scorecard calculation. Some default data queries or questions are mandatory and do not have this option.	Y	N	N

Table 24-2 Role-based user privileges *(continued)*

User actions	Description	User privileges		
		Ransomware Administrator (Y/N)	**Ransomware Answer Questions (Y/N)**	**Ransomware Override (Y/N)**
Answer Question	Opens the answer form to submit a new answer or edit the previous one. This menu option appears only for Question type queries.	N	Y	N
Answer History	Opens the answer history of a query. The report contains the answer trail of the query. This menu option appears only for Question type queries.	N	Y	N
Item History	Shows the audit trail of the query. The report provides change history of the query as well as its responses.	Y	N	N

Table 24-2 Role-based user privileges *(continued)*

User actions	Description	User privileges		
		Ransomware Administrator (Y/N)	**Ransomware Answer Questions (Y/N)**	**Ransomware Override (Y/N)**
Override Value	Shows the form where you can specify the override value. Override value is the percentage value by which you can offset a report-based result. You can use this option to minimize inaccuracies; however, this impacts the overall Ransomware Score. This menu option appears only for Data type queries.	N	N	Y
Override History	Shows the trail of historical overrides of the query, including its notes. This menu option appears only for Data type queries.	N	N	Y

Table 24-2 Role-based user privileges *(continued)*

User actions	Description	User privileges		
		Ransomware Administrator (Y/N)	**Ransomware Answer Questions (Y/N)**	**Ransomware Override (Y/N)**
Refresh	Refreshes the query result by refreshing the associated underlying report to fetch the latest figures. This menu option appears only for Data type queries.	Y	Y	Y
More Info	If the **More Info** link is configured for the query, it launches the link in a new browser tab.	Y	Y	Y

Access Ransomware Scorecard

Users having access permissions to the Ransomware Scorecard, can access the it as follows:

1. Select **Reports** > **Ransomware** > **Ransomware Scorecard**.
2. Select the **Domain**, **Query Status**, **Query Visibility**, and **Query Type** from the scope selector based on the description below and click **Generate**.

Field	Description
Domain	Domain for which you want to create the Ransomware Scorecard

Field	Description
Query Status	Filters the scorecard view based on query status as follows: ■ **All**: Displays both answered and unanswered queries on the scorecard. ■ **Completed**: Displays only the answered queries on the scorecard. ■ **Uncompleted**: Displays only the unanswered queries on the scorecard.
Query Visibility	Filters the scorecard view based on its visibility status as follows: ■ **All**: Displays both enabled and disabled queries on the scorecard. ■ **Enabled**: Displays only the enabled queries on the scorecard. ■ **Disabled**: Displays only the disabled queries on the scorecard. The authority to enable or disable a query lies with the super user.
Query Type	Filters the scorecard view based on the query type as follows: ■ **All**: Displays all queries of all types on the scorecard. ■ **Question**: Displays queries of type Question on the scorecard and hides all other queries. ■ **Data**: Displays queries of type Data on the scorecard and hides all other queries.

The Ransomware Scorecard is generated based on your scope selection.

Ransomware Scorecard components

Ransomware Scorecard card components are described in the table below.

Table 24-3 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.

Table 24-3 Ransomware Scorecard components *(continued)*

Component	Description
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. These are hidden until the user answers the query.
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. Low or Lowest 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.

Scorecard component histories

The Ransomware Scorecard provides a history view of the following components:

- Answer History: Displays the historical responses received for a query of type Question. The history contains details such as Answer (user response), Deleted (Yes/No), Event, Notes as added by the user, and the name of the users who modified the query.
- Override History: Displays the history of override action performed by the user on the query. The history contains details such as Event, Notes, Modified by, and Modification date of the override.
- Item History: Displays the change history of the query and also its responses. The report details include the trail of edits made to the query and another table

that displays a trail of changed responses over time. Each detail is captured with its time stamp.

See “Add a ransomware question” on page 787.

See “Add ransomware data query” on page 796.

See “Edit a ransomware question” on page 795.

See “Edit ransomware data query” on page 798.

See “Answer a ransomware query” on page 785.

See “Override a data query” on page 786.

Ransomware administration

Navigate to the **Actions** menu of each setting displayed on the view which is accessible from **Admin** > **Advanced** > **Ransomware** to access the Ransomware administration settings.

Note: Access to ransomware administration is enabled through the user privileges or user group privileges. See “Assigning user privileges” on page 578.

Table 24-4 Ransomware administration settings

Field name	
Ransomware Pool Space usage threshold (%)	The percentage threshold up to which the storage pool space can be utilized for reporting on ransomware.
Ransomware RPO RTO lookback for job recovery duration (days)	The lookback duration in days for ransomware RPO-RTO for job recovery.
Ransomware RTO backup restore (min 48 hours)	The time duration in hours within which the service disruption must be restored.
Ransomware minimum version for NetBackup Primary	The minimum supported version of Cohesity NetBackup primary server required to generate data for the ransomware scorecard.

Answer a ransomware query

The Ransomware Scorecard displays your preparedness based on the answers submitted for the ransomware queries.

To answer a query:

1. Select **Reports** > **Ransomware** > **Ransomware Scorecard**.
2. Select the **Domain**, **Query Status**, **Query Visibility**, and **Query Type** from the scope selector to access the scorecard containing your target query and click **Generate**. The Ransomware Scorecard is generated based on your scope selection and the queries are denoted as **Question** in the **Type** column of the scorecard.
3. Click the **Actions** menu against your query and select **Answer Question** from the menu. The **Answer Question** window is displayed.
4. Select your answer and add your comments in the **Notes** field. Adding **Notes** is optional. You can post your additional thoughts about your chosen answer.
5. Click **OK** to save your answer and exit.

   Your answer is saved and the Ransomware Scorecard is updated based on your answer. The **Best Practice Recommendation** column also displays the recommendation text specific to the query.

View ransomware query recommendations

Each query that you answer on the Ransomware scorecard is likely to have a description or a best practice recommendation. The recommendation describes the most desirable configuration specific to the ransomware risk discussed in the query. This recommendation can be seen once you or an authorized user has answered the ransomware query.

Ensure you have answered the query before you follow these steps to view the query recommendations:

1. Click the **Actions** menu against the query that you have already answered and select **View recommendation**. This menu option is visible after you have answered the ransomware query.

   The **Answer Question** window is displayed but the fields to answer the query are disabled. However, you can view the recommendations below your answers on the**Answer Question** window below the answered questions.
2. Click **OK** or **Cancel** to exit the **Answer Question** window.

Override a data query

Override option can be used to prevent the scorecard from showing skewed or inaccurate scores resulting from a report-based query. For example, if there are several hosts that are not collected by your backup solutions and you know from the drill-down report that most of them are in your UAT environment, you can set

a proportionate override percentage that you feel accurately represents your total environment. This option is available only for Data type of queries, which derive results from their associated reports.

To override a data query on the Ransomware Scorecard:

1. Click the **Actions** menu against your Data query and select **Override Value** from the menu. The **Override Value** window is displayed.
2. Enter the percentage value by which you want to override the result of the data query and add your comments in the **Notes** field. The **Notes** field is optional and helps you to add comments about the override.
3. Click **OK** to save your changes.

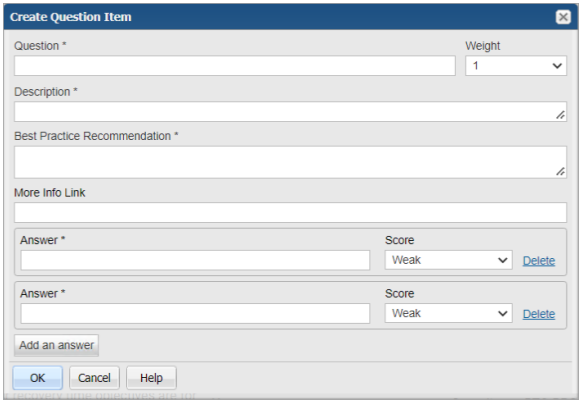
   Your override is set and the Ransomware Scorecard is updated based on the override percentage.

Add a ransomware question

As a user with Ransomware Administrator privileges, you have the privilege to add queries to the default set of queries. This procedure provides the steps you can follow to add custom queries to the Ransomware Scorecard.

To add a question on the Ransomware Scorecard:

1. Select **Reports** > **Ransomware** > **Ransomware Scorecard**.
2. Select the **Domain**, **Query Status**, **Query Visibility**, and **Query Type** from the scope selector to access the scorecard containing your target query and click **Generate**. The Ransomware Scorecard is generated based on your scope selection and the queries are denoted as **Question** in the **Type** column of the scorecard.
3. Click **Add question** on the scorecard and fill the **Create Question Item** form based on the description below:



Field	Description
Question	Enter the query text.
Weight	Assign a weight to your query in a range of 1 to 5, where 1 indicates the least and 5 indicates the most important query. Weight is used as a multiplier of the Score to calculate Risk, which eventually impacts the Ransomware Score. Hence, assign the highest weight to the most important query.
Description	Enter any additional information about the query. This may include the details about the expected responses or guidance about how to answer the query.
Best Practice Recommendation	Describe the recommended best practice for ransomware preparedness with respect to the query. Provide guidance to implement the best practice. This recommendation hidden until the user has answered the query.
More Info Link	Provide a hyperlink to any additional information relevant to your query. This field is optional.
Answer	Enter the answer to you query. If your query has multiple answers, click **Add an answer** to add a row for your answers.
Score	Select a ransomware resiliency rating from the options below: ■ **Lowest** ■ **Low** ■ **Middle** ■ **High** ■ **Highest**

Field	Description
Add an answer	Click to add an answer to your query.

4 Click **OK** to save the query. The new query gets added to the Ransomware Scorecard and becomes available to other users to register their responses.

Utility to add Ransomware Scorecard questions and answers in bulk

The `loadRswQuestionAnsFile` utility to add ransomware questions and answers in bulk provides an efficient method of loading ransomware questions and answers with user adding the question and answer one at a time. This utility takes as input a comma-separated values (CSV) file.

To add questions and answers in bulk:

1. Create two CSV files - one for questions and one for answers.
2. Run the load utility.
3. Verify the questions and answers added in bulk.
4. Verify the questions are available in Ransomware Scorecard.

See "Add Ransomware Scorecard questions in bulk" on page 789.

See "Add answers in bulk to Ransomware Scorecard questions" on page 792.

Add Ransomware Scorecard questions in bulk

You must first create a CSV file containing the questions and add them to the Ransomware Scorecard using the utility.

Create the CSV file of questions

This CSV file becomes the master document of record for custom Ransomware questions and answers. It is the key reference for loading the Ransomware questions and mapping their respective answers on the scorecard. Hence, you must preserve it in a working directory for future updates.

To create the CSV file of questions, create a table with headers and format described below and save it as a CSV file (for example `questions.csv`) in a working directory.

Table 24-5 CSV headers to add questions in bulk

QuestionID	Question	Weightage	Description	Recommendation
Each row under this column header must contain a unique question ID assigned to the question.	Enter the actual Ransomware question.	Enter the numeric value of the weightage assigned to the question.	Enter a brief description or note about the question.	Enter the ideal value or most preferred configuration that is expected as an answer to the question.

A sample of the CSV file of questions can appear as below:

Table 24-6 Example of CSV (`questions.csv`) to add questions in bulk

QuestionID	Question	Weightage	Description	Recommendation
QUES001	Are there processes in place for backups?	3	Backups help in restoring lost or corrupt data.	There should be backups in place to recover data loss or corruption.
QUES002	Are the passwords encrypted?	5	Passwords in the system must be encrypted.	All passwords in the database must be encrypted.
QUES003	Is the source code in proper repository?	2	The software code in proper repository.	The source code must be protected always.
QUES004	Is there a policy to have strong passwords?	4	User passwords.	The password policy must compel all users to keep a strong password.
QUES005	Is NetBackup Appliance used?	4	NetBackup Appliance is safe.	Organizations must use NetBackup Appliance.

Run the utility to add questions

The bulk add utility must be run in SQL Plus as APTARE user.

The load_package utility is located in:

- `/opt/aptare/database/stored_procedures` on Linux
- `\opt\oracle\database\stored_procedures` on Windows

To load questions and answer to ransomware score card

1. Login to the Portal server.
2. At the command line, run:

   ```
   su -aptare
   ```

3. At the command line, launch SQL Plus

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb
   ```

   For example:

   ```
   sqlplus portal/portal@//localhost:1521/scdb
   ```

4. Run this command at the SQL prompt.

   ```
   SQL> Execute load_package. loadRswQuestionAnsFile ('domain_name',
    'pathname_and_filename','log_path_name', 'log_file_name',
   'Type');
   ```

   Where:

domain_name	Name (enclosed in single straight quotes) of the domain in which the host groups and hosts reside. Example: 'DomainEMEA'
pathname_and_filename	Full path + filename (enclosed in single straight quotes) of the CSV file. ■ Windows Example: 'c:\temp\questions.csv' ■ Linux Example: '/tmp/questions.csv'
log_path_name	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to `scon.log` and `scon.err`. To omit this parameter, enter: **'c:\temp'** or **'/tmp'**

log_file_name	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to `scon.log` and `scon.err`. To omit this parameter, enter: **'questions.log'**
type	'Questions' for loading the questions in bulk.

Example:

```
SQL> load_package.loadRswQuestionAnsFile ('INSTALLWIN2012','/tmp/question
```

5. Check the log file for status and errors.
6. Restart the Portal services so that the newly added questions are available in Ransomware Scorecard.
7. Verify that the bulk adding of questions was successful: In the Portal, go to **Reports** > **Ransomware** -> **Ransomware Scorecard**.

Add answers in bulk to Ransomware Scorecard questions

This procedure helps you to add answers to the Ransomware Scorecard questions in bulk and also associate the correct answer to each question. You must first create a CSV file of answers and then add it to the Ransomware Scorecard using the utility

The `loadRswQuestionAnsFile` utility provides an ability to bulk load Ransomware answers corresponding to the questions you loaded. This utility takes as input a comma-separated values (CSV) file. You have to load the questions before you can add the answers for the question..

Create the CSV file of answers

This CSV file becomes the key document of record for custom Ransomware questions and answers. Hence, you must preserve it in a working directory for future updates.

To create the CSV file of answers, create a spreadsheet table in the format shown below. Save it as a CSV file (for example `answers.csv`) in a working directory. This file is essential to map the answers to their respective Ransomware question accurately.

Table 24-7 Headers and description of CSV file for answers

QuestionID	AnswerID	Answer	Score
Enter the question identifier to which the answer identifier has to be mapped. This question identifier must match with the identifier value in the CSV created for questions.	Enter the answer identifier that you want to link to the QuestionID. The QuestionID and AnswerID must form a unique pair. If there are multiple answers to the same question, create a unique AnswerID for each answer and associate it with the same QuestionID.	Enter the answer to the question. If there are multiple answers, enter each value on a separate row so that it associates with a unique AnswerID.	Enter the score associated with each answer.

A typical CSV file to add multiple answers in bulk can appear as below.

Table 24-8 Headers and description of CSV file for answers

QuestionID	AnswerID	Answer	Score
QUES001	ANS001	Yes	4
QUES001	ANS002	No	0
QUES002	ANS001	Yes	4
QUES002	ANS002	No	0
QUES003	ANS001	Yes	4
QUES003	ANS002	No	0
QUES004	ANS001	Yes	4
QUES004	ANS002	Kind of	3
QUES004	ANS003	Not that strict	2
QUES004	ANS004	No	0
QUES005	ANS001	Yes	4
QUES005	ANS002	No	0

Run the utility to add answers

The bulk add utility must be run in SQL Plus as APTARE user.

The load_package utility is located in:

- `/opt/aptare/database/stored_procedures` on Linux
- `\opt\oracle\database\stored_procedures` on Windows

To load questions and answer to ransomware score card

1. Login to the Portal server.
2. At the command line, run:

   ```
   su -aptare
   ```

3. At the command line, launch SQL Plus

   ```
   sqlplus <pwd>/<pwd>@//localhost:1521/scdb
   ```

   For example:

   ```
   sqlplus portal/portal@//localhost:1521/scdb
   ```

4. Run this command at the SQL prompt.

   ```
   SQL> Execute load_package. loadRswQuestionAnsFile ('domain_name',
    'pathname_and_filename','log_path_name', 'log_file_name',
   'Type');
   ```

   In the above command:

domain_name	Name (enclosed in single straight quotes) of the domain in which the host groups and hosts reside. (Example: 'DomainEMEA')
pathname_and_filename	Full path + filename (enclosed in single straight quotes) of the CSV file ■ Windows Example: 'c:\temp\answers.csv' ■ Linux Example: '/tmp/answers.csv'

log_path_name	Full path (enclosed in single straight quotes) where the log file will be created/updated; verify that you have write access to this directory. Optional: If a log path and filename are not specified, log records are written to `scon.log` and `scon.err`. To omit this parameter, enter: Example: 'c:\temp' or '/tmp'
log_file_name	Log file name enclosed in single straight quotes. Optional: If a log path and filename are not specified, entries are written to `scon.log` and `scon.err`. To omit this parameter, enter: '' Example: answers.log'
type	'Answers' of the questions loaded in bulk.

Example:

```
SQL> load_package.loadRswQuestionAnsFile
('INSTALLWIN2012','/tmp/questions.csv','/tmp',
'questions.log','Answers');
```

5 Check the log file for status and errors.

6 Restart the Portal services so that the newly added questions are available in Ransomware Scorecard.

7 Verify that the bulk adding of questions was successful: In the Portal, go to **Reports** > **Ransomware** -> **Ransomware Scorecard**. Select a questions you want to view and click 'Answer Question'.

Edit a ransomware question

You can edit only those ransomware queries that you have added as a ransomware administrator.

To edit a query on the Ransomware Scorecard:

1 Click the **Actions** menu against your query to view the additional options.

2 Click **Edit Question** from the menu and update the **Edit Question** form with the required changes.

3 Click **OK** to save the edits. The changes appear on Ransomware Scorecard and become available to other users to register their responses.

See "Add a ransomware question" on page 787.

See “Ransomware Scorecard overview” on page 775.

Add ransomware data query

Data query enables you to perform a data-driven evaluation of your ransomware preparedness. You can add a data query around a critical report value and evaluate your ransomware score and trend. To display accurate Ransomware score and its trend, the report instance used in the data query must contain only a single row, column, and numeric value. For this reason, your report instance must span your entire domain. The report instance may contain an average, a sum total, or a percentage value derived from multiple sources. Hence, this value can be typed either as a **Number**or **Percentage** and a default score of lowest, low, medium, high, or highest is assigned to it. User responses are compared to this default score to determine the ransomware score and trend. For example, a data query to show the deployment status of a specific NetBackup EEB requires a report instance that calculates the percentage of NetBackup primaries with the EEB installed.

As a user with Ransomware Administrator privileges, you have the privilege to add data queries to the default set of queries. This procedure provides the steps you can follow to add custom data queries to the Ransomware Scorecard.

To add a ransomware data query on the Ransomware Scorecard:

1. Select **Reports** > **Ransomware** > **Ransomware Scorecard**.
2. Select the **Domain**, **Query Status**, **Query Visibility**, and **Query Type** from the scope selector to access the scorecard containing your target query and click **Generate**. The Ransomware Scorecard is generated based on your scope selection and the data queries are denoted as **Data** in the **Type** column of the scorecard.
3. Click **Add data query** on the scorecard and fill the **Add Data Query** form based on the description below:

Field	Description
Question	Enter the query text.
Weight	Assign a weight to your query in a range of 1 to 5, where 1 indicates the least and 5 indicates the most important query. Weight is used as a multiplier of the Score to calculate Risk, which eventually impacts the Ransomware Score. Hence, assign the highest weight to the most important query.
Description	Enter any additional information about the query. This may include the details about the expected responses or guidance about how to answer the query.

Field	Description
Best Practice Recommendation	Describe the recommended best practice for ransomware preparedness with respect to the query. Provide guidance to implement the best practice. This recommendation hidden until the user has answered the query.
More Info Link	Provide a hyperlink to any additional information relevant to your query. This field is optional.
Saved Report Instance	Select the report instance from which the query can source values to determine the data item score. This list displays the names of all report instances saved under **My Reports**. The report you select must be based on a report template that returns only one value of type number or percentage. Preferably, the value must have a drill down view that displays the sources from which it was derived. If your report does not include a drill down, you will be asked if you want to continue, as having a drill down is a best practice in most cases.
Schedule	Schedule the time when the data query must run.
Data Type	Select the type of value that the report instance must return. You can choose **Number** or **Percentage** as type of value from the list.
Default Score	Set the default value to use if the report returns a value outside the mapped range. Default Score options: ■ **Lowest** ■ **Low** ■ **Middle** ■ **High** ■ **Highest**

Field	Description
Range	Define the value range and assign a resiliency score for the defined range from the options below. User response will be assigned a score from these options based on which numeric range in which their responses fit into. You can define multiple ranges depending on the anticipated user responses. If the **Data Type** is set to **Percentage**, all ranges must be specified between 0 and 100. ■ **Start**: Enter the starting value of the score range. This value is inclusive ■ **Less than**: Enter the end value of the score range. This value is exclusive. ■ **Score**: Assign a score from the list to the defined range. This score is compared with the default score to determine the ransomware score and trend.
Add Result	Click to add another score-mapping row for your data query.
Delete	Click to remove a score-mapping row.

4 Click **OK** to save the data query. The new data query appears on the Ransomware Scorecard and becomes available to other users to register their responses.

Edit ransomware data query

You can edit only those ransomware data queries that you have added as a ransomware administrator.

To edit a query on the Ransomware Scorecard:

1 Click the **Actions** menu against your query to view the additional options.

2 Click **Edit Data Query** from the menu and update the **Edit Data Query** form with the required changes.

3 Click **OK** to save the edits. The changes appear on Ransomware Scorecard and become available to other users to register their responses.

See “Add ransomware data query” on page 796.

See “Ransomware Scorecard overview” on page 775.

Chapter 25

Analyze files

This chapter includes the following topics:

- Overview
- File categories
- Adding/Editing file categories
- File list export
- File list export output in CSV format

Overview

File Analytics collects volume and share metadata to enable you to identify and manage unstructured data that may be consuming storage in your enterprise. This data is profiled and categorized so that you can easily identify areas that need attention.

File categories

A set of default file categories is provided. Although not recommended these categories can be customized (by a Super User only) to add or remove file extensions.

File Analytics profiles and categorizes collected data by file type (filename extension). This categorized data is listed in the **File Types** report, which lists the top 50 file types for the selected volumes/shares.

Note: The **File Types** report is an aggregation at the volume level of only the top 50 file extensions. For a complete list of all profiled file types, use the **Export List** checkbox, available in the **File Types** report and several other File Analytics reports; or from the Tools menu list, select: **Admin > File List > Export.** File List Export represents all filenames profiled for the Portal. Typically, if multiple volumes/shares are selected in the report scope, the **File Types** report will list less types than what is listed in the exported output.

Table 25-1 File categories

Category	Filename Extensions
Adobe Acrobat	pdf
Archives	zip, tar, cab, gz, rar, z
Audio	aiff, au, mp3, mp4, wav, wma
Backups	bak, bkp, bkf
CD/DVD Images	iso, nrg, img
Desktop Publishing	qxd
Email Archives	pst
Hard Drive Images	tib, gho, ghs
Images	bmp, gif, jpg, jpeg, tif, tiff, png, psd
Installers	msi, rpm
Log Files	log
Lotus Notes	nsf, ns2, ns3, ns4, box, ncf, ntf
MS Office Documents	doc, docx, docm, dot, dotx, dotm, xls, xlsx, xlsm, xlt, xltx, xltm, xlsb, xlam, ppt, pptx, pptm, potx, potm, ppam, ppsx, ppsm, mdb, accdb, accde, accdt, accdr
System Files	exe, dll, bin
Text Files	txt, csv
Video	mpg, mpeg, avi, mov, m4v
VMware	vmdk, nvram, vmx, vmxf, vmtm, vmem, vmsn, vmsd, hlog

Table 25-1 File categories *(continued)*

Category	Filename Extensions
Ransomware	

Table 25-1 File categories *(continued)*

Category	Filename Extensions
	emc, #, ##encrypted_by_pablukl0cker##, #locky, 0000, 0wn3dyou, 0x0, 0x004867, 0x009d8a, 1999, 1btc, 1cbu1, 1txt, 2cxpcihgsvxb3, 3301, 3ncry, 3ncrypt3d, 490, 491, 492, 4rwcry4w, 4x82n, 63vc4, 666, 6fkr8d, 707, 725, 726, 73i87A, 7h9r, 7zipper, 8637, 8lock8, 96e2, @decrypt2017, AngleWare, BarRax, CCCRRPPP, CCCRRRPPP, Dexter, EnCiPhErEd, LeChiffre, MERRY, MRCR1, PEGS1, PoAr2w, R16m01d05, RARE1, RMCM1, SecureCrypted, VforVendetta, Whereisyourfiles, _AiraCropEncrypted, __dilmav1, _airacropencrypted) not found., _ryp,
	a19, a5zfn, a604af9070, a990, a9v9ahu4, aaa, aajf, abc, actum, adk, aes!, aes-ni, aes256, aes_ni, aes_ni_0day, aesir, aga, airacropencrypted!, akaibvn, akira, alcatraz, aleta, allcry, alosia, amba, amnesia, andonio, android, angelamerkel, anon, antihacker2017, anubi, ap19, area, areyoulovemyrans, areyoulovemyransfile, armadilo1, asasin, asdasdasd, atlas, au1crypt, axx, azer,
	b0ff, b10cked, b5c6, b89b, bagi, bam!, bart, basslock, bbqb, beef, beep, beethoven, belgian_cocoa, better_call_saul, big1, bitkangoroo, bitstak, bleep, bleepyourfiles, blind, blind2, bloc, blocatto, bloccato, block_file12, blocked, blocked2, bloked, bmcode, bonum, braincrypt, breaking bad, breaking_bad, breeding123, brickr, bript, brt92, btc, btc-help-you, btcbtcbtc, btcware, bugware, bunny, bush,
	c0rp0r@c@0xr@, canihelpyou, cawwcca, cbu1, ccc, ceber3, cerber, cerber2, cerber3, cerber6, cerbersyslocked0009881, cesar, cezar, cfk, chak, chartogy, chifrator@qq_com, chip, christmas, cifgksaffsfyghd, ck, clinton, cloud, cobra, code, coded, coin, comrade, conficker, conflicker, corrupted, country82000, coverton, cr020801, crab, cradle, crh8, crime, crinf, cripted, criptiko, criptokod, cripttt, crjocker, crjoker, crptd, crptrgr, crptxxx, crrrt, cry, cry128, cry36, cry9, cryeye, crying, cryp1, crypt, crypt1, crypt12, crypt38, crypte, crypted, crypted000007, crypted_file, cryptedopps, crypto, cryptoboss, cryptobyte, cryptojoker, cryptolocker, crypton, cryptoshiel, cryptoshield, cryptotorlocker2015!, cryptowall, cryptowin, crypttt, cryptwalker, cryptz, crypz, crysis, cspider, ctb2, ctbl, ctbl2, cyberdrill, cybersoldiersst, cyclone, cyron, czvxce, czvxe,
	d2550a49bf52dfc23f2c013c5, d4nk, dCrypt, da_vinci_code, dale, damage, damoclis, darkcry, darkness, dcry, decrypt2017, decryptional, deria, derp, deuscrypt, dg, dharma, diablo6, dian, die, disposed2017, dle, dlenggrl, dolphin, domino, doomed, doxes, ds335, duhust, dviide, dwbiwty, dxjay, dxxd, dyatel@qq_com,
	e4m, ebay, ecc, eclr, edgel, eky, empty, enc, encedrsa, encencenc, enciphered, encoderpass, encr, encrptd, encrypt, encrypted,

Table 25-1 File categories *(continued)*

Category	Filename Extensions
	encryptedaes, encryptedped, encryptedrsa, encryptedyourfiles, encryptile, enigma, enjey, epic, error, eternity, evil, evillock, executioner, exotic, exploit, explorer, exte, exx, ezz,
	facebook, failedaccess, fairytail, fake, fantom, fartplz, fat32, fbuvkngy, fil0locked, file0locked, filegofprencrp, fileiscryptedhard, filesfucked, filock, firecrypt, fix, flat, flux, flyper, fmoon, freefoam, frivolity, frtrss, fuck, fuck_you, fuck_you_av_we_are_not_globe_fake, fucked, fucku, fuckyourdata, fun,
	gangbang, ganklocked, gdcb, gefickt, gembok, getrekt, ghost, gigahertz, globe, gocr, godra, gommemode, good, gorilla, gotham, gotya, gr3g, granny, grt, grux, gruzin@qq_com, gryphon, gsupport3,
	h3ll, h_f_d_locked, ha3, hacked, hakunamatata, hannah, happ, happenencedfiles, happy, happydayzz, harzhuangzi, hasp, haters, hb15, hcked, heisenberg, hello, helpdecrypt@ukr_net, helpmeencedfiles, herbst, heroesofthestorm, hncrypt, hnumkhotep, hnyear, howcanihelpusir, hrm, htrs, hush,
	i"want money, iaufkakfhsaraf, id-3044989498_x3m, ifuckedyou, igotyou, ihsdj, illnest, imsorry, incanto, infected, infinite, info, insane, ipcrestore, ipygh, isis, iwant, iwanthelpuuu,
	jaff, jeepers, jey, jezroz, justbtcwillhelpyou,
	k0stya, katipuneros, kee, keepcalm, kencf, kernel_complete, kernel_pid, kernel_time, keybtc@inbox, keybtc@inbox_com, keyh0les, keyholes, keyz, kgpvwnr, kilit, kill, killedxxx, kimchenyn, kimcilware, kirked, kk, kkk, kok, korea, korrektor, kostya, kr3, kra, krab, kraken, kratos, krypted, kryptonite, kuntzware, kyra,
	l0cked, lalabitch,, lambda_l0cked, lamo, lckd, lcked, legion, lego, leon, lesli, letmetrydecfiles, lfk, lightning, lime, lock75, lock93, lockd, locked, locked-by-mafia, locked3, locked_by_pablukl0cker, locked_file, lockify, locklock, lockme, lockout, locky, lokitus, lol!, loli, loptr, lordofshadow, losers, lost, loveransisgood, lovewindows, loveyou, loveyouisreal, ltml, lukitus, lukitus-tiedostopäätettä,
	madebyadam, magic, magic_software_syndicate, maktub, malki, maniac, matrix, maxicrypt, maya, maysomware, medal, mention9823, micro, mikoyan, mind, mole, mole00, mole01, mole02, mole03, mole04, moments2900, mordor, mp3, mtc, mtk118, mychemicalromance4ever, myransext2017,

Table 25-1 File categories *(continued)*

Category	Filename Extensions
	nalog@qq_com, napoleon, neitrino, nemesis, netn6, news, nm4, no_more_ransom, noblis, nochance, node0, noob, nopasaran, noproblemwedecfiles, noproblemwedecfiles, notfoundrans, nsmf, nuclear, nuclear55, nuke55, numberdot, ocean, odcodc, odin, ogonia, ogre, ohno!, okokokokok, oled, omg!, one-we_can-help_you, oni, onion, only-we_can-help_you, onyon, oops, oor, oplata@qq_com, ordinal, oshit, osiris, otherinformation, otr, owned, oxr, p5tkjw, pa-siem, pablukcrypt, pabluklocker, padcrypt, panda, pay, paybtcs, paycyka, payday, payfordecrypt, payforunlock, paym, paymds, paymrss, paymrts, payms, paymst, paymts, payransom, payrms, pays, paytounlock, pdcr, pec, pedo, perl, petya, phantom, phobos, pirate, pizda@qq_com, pizdec, pizdosik, pky, planetary, plauge17, plin, pnr, pohu, porno, pornoransom, poshkoder, potato, powerfulldecrypt, powned, pr0tect, privat66, prosperous666, pscrypt, pubg, purge, pwned, pzdc, qwerty, qwqd, r3k7m9, r3store, r4a, r4bb0l0ck, r5a, raas, radamant, raid10, ramen, ranranranran, ranrans, ransed, ransom, ransommine, rapid, rastakhiz, razarac, razy, razy1337, rdm, rdmk, rdwf, readme_txt, reagan, realfs0ciety@sigaint.org.fs0ciety, reco, rekt, relock@qq_com, remind, resurrection, revenge , revolution, reyptson, rip, rmd, rnsmwr, rnsmwre, rokku, rose, rrk, rsnslocked, rsplited, rtyrtyrty, ruby, rumblegoodboy, ryp, s1crypt, sage, saherblueeagleransomware, salsa222, same, samsung, sanction, satan, scarab, scl, scorpio, securecrypte, senrus17, serp, serpent, server, sevendays, sexy, sgood, shadow, shark, shifr, shinigami, shino, shit, shutdown57, sifreli, silent, sinta, skjdthghh, skunk, skvtb, slvpawned, snake, son, spectre, spider, spora, sport, srpx, sshxkej, stn, stop, stroman, styx, supercrypt, supported2017, suppose666, suprise, surprise, sux, svn, symbiom_ransomware_locked, sysdown, szesnl, szf, tastylock, technicy, tesla, test, tgif, thetrumplockerf, thetrumplockerp, theva, theworldisyours, thor, toxcrypt, trans, triple_m, trmt, troyancoder@qq_com, true, trump, trun, ttt, tzu, udz2j8mv, uiwix, unavailable, unlis, usr0, vForVendetta, vbransom, vcrypt1, vdul, velikasrbija, velso, vendetta, venusf, venusp, viiper, viki, vindows, visioncrypt, vpgvlkb, vrmrkz, vscrypt, vvv, vxlock,

Table 25-1 File categories *(continued)*

Category	Filename Extensions
	wallet, wamarlocked, wana decrypt0r trojan-syria editi0n, warn_wallet, wcry, wcryt, wdie, weapologize, weareyourfriends, weencedufiles, wflx, whatthefuck, whycry, wincry, windows, windows10, wlu, wmfxdqz, wncry, wncrypt, wncryt, wndie, wnry, wooly, wowreadfordecryp, wowwhereismyfiles, write, write_us_on_email, wrny, wsmile, wtdi, wuciwug, www, wxdrjbgsda, wyvern, x0lzs3c, x1881, x3m, x3mpro, xbtl, xcrypt, xdata, xfile, xhspythxn, xiaoba1, xiaoba10, xiaoba11, xiaoba12, xiaoba13, xiaoba14, xiaoba15, xiaoba16, xiaoba17, xiaoba18, xiaoba19, xiaoba2, xiaoba20, xiaoba21, xiaoba22, xiaoba23, xiaoba24, xiaoba25, xiaoba26, xiaoba27, xiaoba28, xiaoba29, xiaoba3, xiaoba30, xiaoba31, xiaoba32, xiaoba33, xiaoba34, xiaoba4, xiaoba5, xiaoba6, xiaoba7, xiaoba8, xiaoba9, xmdxtazx, xncrypt, xolzsec, xorist, xort, xrnt, xrtn, xtbl, xxx, xyz, xzzx, yakes, ykcol, yl, youransom, yourransom, ytbl, yyto, z3r0, z81928819, zablokowane, zayka, zbt, zc3791, zcrypt, zendr4, zepto, zilla, zimbra, zino, zorro, zuzya, zxz, zycrypt, zyklon, zzz, zzzz, zzzzz, xdata

Adding/Editing file categories

Note: WARNING: When you edit any file category, you will lose the historical category data because the revised category no longer represents the data that already has been collected and categorized.

Add a file category

Access this Super User feature from:

Admin>Reports>File Categories

1. On the **File Categories Administration** window, click **Add**.
2. Complete the following fields in the **Add Category** window.

Name	Name of the file category
Extensions	Enter comma-delimited file extensions, without the period; for example, mif

Edit a file category

When editing a file category, simply add or delete file extensions in the comma-separated list.

Note: WARNING: When you edit any file category, you will lose the historical category data because the revised category no longer represents the data that already has been collected and categorized.

File list export

For File Analytics, the File List Exporter provides a mechanism for extracting the File Analytics collected metadata into a comma-separated values (.csv) file, enabling further analysis of files that may be wasting storage. Up to 20K rows can be exported to this CSV file.

See "File list export output in CSV format" on page 808.

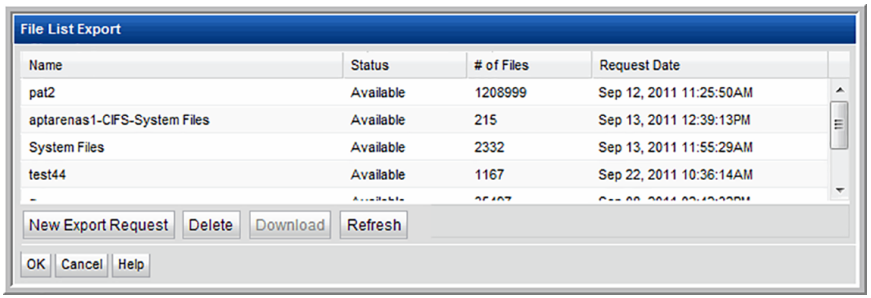
From the Portal toolbar, choose **Admin > File List > Export.**

In addition to accessing this tool using the **Admin** menu list, several File Analytics reports include a check box that launches the File List Exporter:

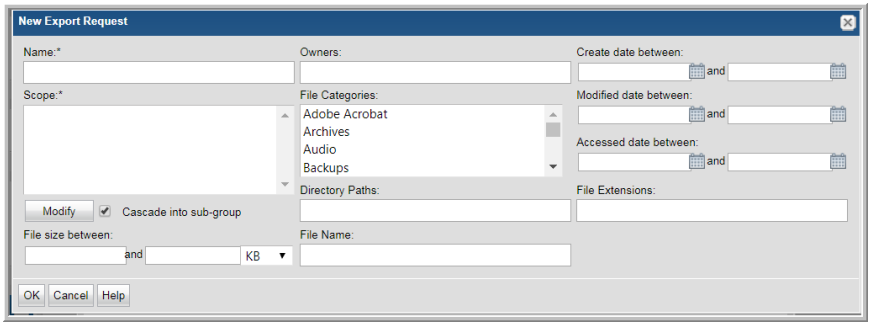
- Usage by Owner
- File Types
- File Categories

The **File List Export** window lists the export processes that are currently in progress, as well as those that are available for download.

Note: If you select an export process that is in a Processing state and then click **Delete**, the request will be removed from the list and the current processing will be terminated.



To initiate a File List Export process, click **New Export Request**and enter the filtering criteria for the files to be investigated.



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
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.
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 > File Categories. See “File categories” on page 799.
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)
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.

File list export output in CSV format

The data exported with this tool is available in a comma-separated-values (CSV) file, in the following format:

Domain ID	The domain in which the host is located.
Host	IP address or host name on which the file is located.
Path	Full path and file name, such as: `"services/Docs/Collateral/DS-ProLaunch_2011.pdf"` This value is enclosed in quotes to accommodate spaces within the path and file name.
File Size	The size of the file, in bytes.
Owner	The file's owner.
Create Date	Date and time the file was created, in the format: `YYYY-MM-DD HH:MM`
Modified Date	Date and time the file was last modified, in the format: `YYYY-MM-DD HH:MM`
Access Date	Date and time the file was last accessed, in the format: `YYYY-MM-DD HH:MM`
Attributes	File attributes, if attributes are enabled for the file, are listed. A single character represents each attribute. This value can contain multiple, concatenated attributes, such as **SRE**. D = Directory S = System R = Read-only H = Hidden E = Encrypted C = Compressed A = Archive

Exported CSV file example

Domain Id	Host	Path	File Size	Owner	Create Date	Modified Date	Access Date	Attributes	Backup Policies	Backup Time
100000	172.16.1.27	PS/Doc/ DS_Store	6148	Unknown	1/3/2010 1:49	2/13/2010 2:04	12/13/2010 1:49	AH	*Standard, Hyper-V*	10/6/2020 3:06:00
100000	172.16.1.27	PS/Doc/ Services Tasks .docx	23428	Unknown	4/1/2011 5:40	4/11/2011 15:40	4/11/2011 15:40	A	*Standard*	10/6/2020 3:06:00
100000	172.16.1.27	PS/Doc/ Migration .docx	18387	Unknown	8/7/2013 8:38	10/7/2013 18:38	10/7/2013 18:38	A	*Standard*	10/6/2020 3:06:00
100000	172.16.1.27	PS/Proc/ Disaster Recov .docx	305792	Unknown	3/5/2013 2:30	9/15/2013 20:30	10/15/2013 20:30	A	*Hyper-V*	10/6/2020 3:06:00

Chapter 26

Troubleshoot the Portal

This chapter includes the following topics:

- Support tools
- Requesting a support package
- File selector wildcards
- Retrieving log files
- Archiving script
- Debug
- Login errors
- Add notes to failed jobs
- Host name processing - Filters and aliases
- Change the Portal server's IP address
- Changing the port for Portal server communication
- Determine report rendering statistics
- Clearing the inventory and report cache
- Troubleshoot Oracle performance issues in IT Analytics
- Troubleshoot test connection locale issue

Support tools

To help facilitate troubleshooting when working with Support, you can download data collector/portal log and configuration files for a specified time period.

To request a support package, select **Admin > Advanced > Support Tools**.

Use the view pane to track the status of your download and access the files as they become available. The grid is sorted by **Request Date**.

Table 26-1 Support tools

Field name	Description
Add	Click to create a support package request. You can create a portal server configuration files package or a data collector log package. See "Requesting a support package" on page 812.
Delete	Click Delete to remove or cancel a package creation.
Refresh	Click Refresh to update the status on the view pane. The pane does not refresh automatically.
Name	Displays the name assigned to the package.
Collector	The collector associated from which the log files were retrieved.
Type	Displays the package type: Portal or Data Collector (with collector name).
Package Content	Displays the content of the support package.
Status	The Support Package requests are listed with links to access the downloadable files. The Status of a request can be: ■ Queued ■ Processing ■ Available - Click the Available link to download the zip file. ■ Failed - Mouse-over a Failed status to view the details of the failure.
Request Date	Displays the creation date and time of the support request package. By default, the grid is sorted by descending request date.

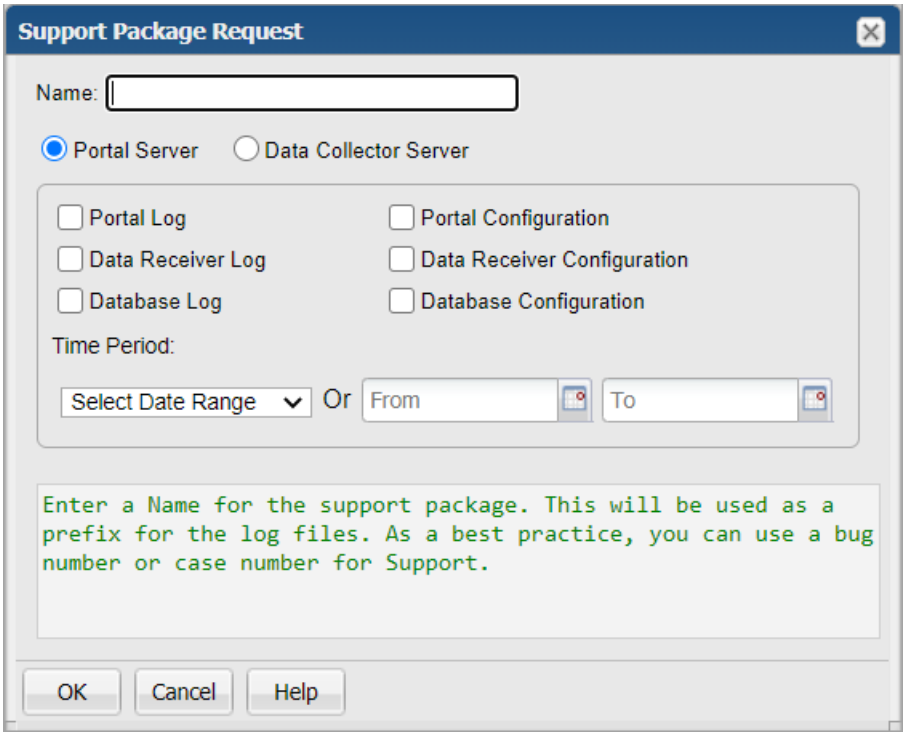
Requesting a support package

When creating a support package, you can choose to download data collector or portal files for a specified time period. Use the view pane to track the status of your request.

To Request a Support Package with Portal Server Files

1. Select **Admin > Advanced > Support Tools**.
2. Click **Add**. The **Support Package Request** dialog is displayed.
3. Enter a **Name** for the support package. This will be used as a prefix for the log files. As a best practice, you can use a bug number or case number for Support.

4 Choose **Portal Server**



5 Configure the dialog:

Portal Log	Select to download the Portal.log, upgrade.log, audit.log and relevant tomcat logs. This provides detailed logging for the portal servlet. Logs portal login requests, user impersonations, portal reports that are run - basically everything the user does in the Portal web browser window. Database problems show up as SQL exceptions and often list the associated Oracle error number (ORA nnn).
Portal Configuration	Select to download the httpd.conf, httpd-ssl.conf, portal.properties, server.xml, startup_agent.sh, startup_portal.sh, systemConfiguration.properties and web.xml.
Data Receiver Log	Select to download the datarcvr.log and relevant tomcat logs.
Data Receiver Configuration	Select to download all files under /opt/aptare/datarcvrconf.
Database Log	Select to download **scon.log, scon.err** and **aptare-trans.log**. The scon.log provides detailed logging from database stored procedures. Shows input and output values, and timing of queries. Oracle errors are prefixed with ORA-nnnn, where nnnn is the error number. On Windows, if the directory C:\tmp exists, the file will be written there. Otherwise, it will be written to **C:\opt\oracle\logs**. The scon.err provides error messages from the database stored procedures. Oracle errors are prefixed with ORA-nnnn, where nnnn is the error number. On Windows, if the directory **C:\tmp** exists, the file will be written there. Otherwise, it will be written to **C:\opt\oracle\logs**. The **aptare-trans.log** provides the transactions that have run and the time taken to run them. This is useful to troubleshoot performance issues and reports that take a very long time to run.
Database Configuration	Select to download the initscdb.ora file.
Time Period	Choose a date range such as last 4 hours, last 24 hours, last week or specify a time period.
	Check this box to send these files directly to Veritas Support for analysis.

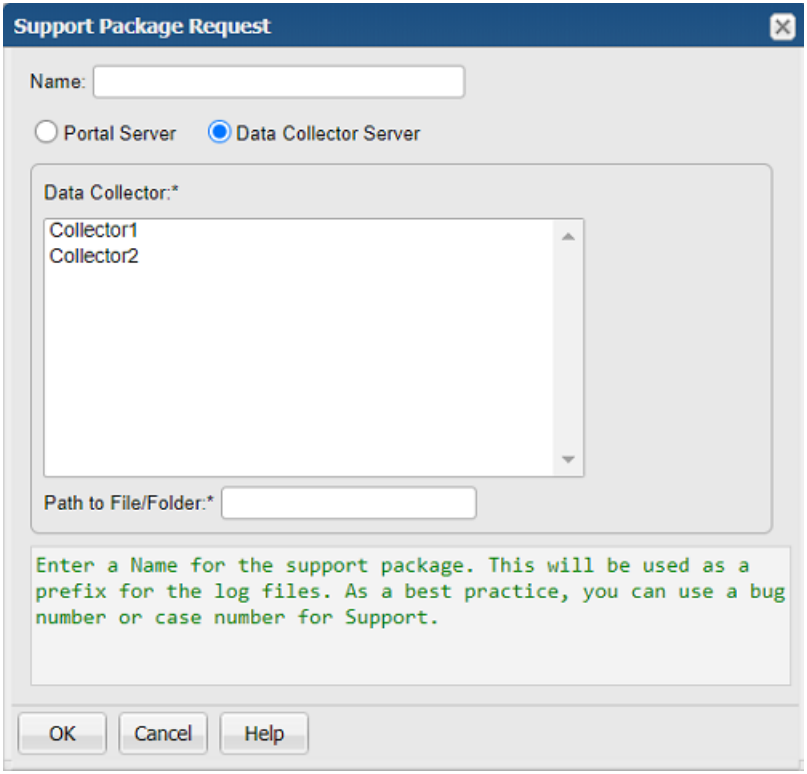
To Request a Support Package with Data Collector Server Files

1 Select Admin >Advanced > Support Tools.

2 Click Add. The Support Package Request dialog is displayed.

3 Enter a Name for the support package. This will be used as a prefix for the log files. As a best practice, you can use a bug number or case number for Support.

4 Choose Data Collector Server.



5 Configure the dialog:

Data Collector*	Select the Data Collector(s) from which the log files will be retrieved. You can multi-select the data collectors. This enables you to create multiple support packages at once. Each collector will be listed as an individual support package. The associated collector is displayed as a column in the same row as the support package so you can easily identify the origin of the log.
Path to File/Folder*	Enter one or more paths separated by commas. You may use /* to select multiple files in the directory, or /** to also include files in subdirectories. The paths will always be relative to the <APTARE_HOME> folder - do not include the explicit <APTARE_HOME> portion in the path. For informational and troubleshooting purposes only, know that <APTARE_HOME> is typically configured to be: /opt/aptare on Linux and C:\Program Files\Aptare on Windows. Moving up the directory hierarchy using .. is not supported. In addition to log files, other files useful for troubleshooting can be downloaded; for example, collectorconfig.xml or updateconfig.sh. The most common location from which to download files is: /mbs/logs/**

File selector wildcards

When specifying file paths in the Support Tools, wildcards may be used to select multiple files. All matching is case-insensitive, regardless of platform. Multiple paths can be specified by separating the paths by commas (with optional white space) for example, mbs/logs/, mbs/rawdata/. If a path contains a comma then you should enclose the entire path in [square brackets]. Note that doing so will preclude the use of some of the following wildcard options.

The following table displays potential wildcards for file selection:.

Table 26-2 File selector wildcards

Value	Description
`<file path>`	The exact file will be retrieved for example, mbs/conf/collectorconfig.xml
`<directory path> or <directory path>/`	All contents of the directory will be retrieved for example mbs/logs/
`*`	Matches any part of the filename within the same directory. For example mbs/logs/*.log will get all .log files that are in the mbs/logs directory (but not in any sub-directories).
`**`	Matches any part of the path, including crossing directory boundaries. For example mbs/rawdata/**192.168.0.1** will retrieve any file in the mbs/rawdata directory or subdirectories that contains 192.168.0.1 in the filename.
`/**/ (or **/ at the start of the path)`	Matches zero or more directories. For example mbs/rawdata/**/192.168.0.1/** will retrieve any file under the mbs/rawdata directory that is also under a sub-directory named 192.168.0.1. **Note:** /**/ (or **/ at the start of the path) may only be present once.
`?`	Matches any single character. For example mbs/rawdata/192.168.0.? would match either of mbs/rawdata/192.168.0.1 or mbs/rawdata/192.168.0.2.
`{a,b}`	Specifies alternative values, exactly one of which must be matched. For example mbs/logs/{metadata,eventcollector}.log will retrieve either of mbs/logs/metadata.log or mbs/logs/eventcollector.log.
`[characters]`	Matches any single character that is specified within the square brackets. Either individual characters or a range may be specified. For example 192.168.0.[125-7] would match any of 192.168.0.1, 192.168.0.2, 192.168.0.5, 192.168.0.6, 192.168.0.7. **Note:** This wildcard cannot be used if the path is enclosed in square brackets.
`[!characters]`	Matches any single character that is not specified within the square brackets. Either individual characters or a range may be specified. For example 192.168.0.[!125-7] would match any of 192.168.0.3, 192.168.0.4, 192.168.0.8, but would not match 192.168.0.1. **Note:** This wildcard cannot be used if the path is enclosed in square brackets.

Retrieving log files

Portal and data collector files retrieved using the **Support Tools** function are archived as .zip.xz. This format is used to minimize the disk space required on the portal and collector plus minimize network transfer time.

The format .zip.xz is not natively supported by most platforms, but there are standard tools available to process them. On most platforms it will be necessary to first extract the uncompressed .zip from the .xz, and then extract the files from the .zip. This means that the system on which the files are being extracted will need free disk space equal to approximately twice the size of the uncompressed files.

On a Windows platform

Windows requires a third-party tool to process the archives. 7-zip (http://www.7-zip.org/) and its derivatives (such as Easy 7-zip (http://www.7-zip.org/) are suitable. Other extraction tools may also provide .xz and .zip support.

To extract using 7-zip is a two-step process:

1. Extract the .zip from the .xz
2. Extract the files out of the .zip

On a Linux platform

Most Linux distributions come with command-line support for .xz and .zip archives, or such support can be easily added.

To extract using the standard command-line tools is a two-step process, as the standard unzip tool cannot take a zip archive via standard input (cannot pipe into unzip).

```
unxz archive.zip.xz
```

```
unzip archive.zip
```

Archiving script

There is a script available: mbs/bin/archive.sh/mbs/bin.archive.bat, for archiving logs, rawdata and configuration files. The script is essentially a collector-side version of the portal Support Tools, to be used when the portal's Support Tools are not available (for example, the collector service is not running).

When executed with no arguments, the script will archive configuration files (not mbs/conf/subsystem), logs (including upgrade and WMI logs) and rawdata to a time-stamped archive in mbs/tmp/archive.

Additional usage options are available via archive.sh -- help for example, changing the output archive or specifying different patterns to archive. All patterns available for the Support Tools in the portal are available.

Debug

If you experience system problems, Cohesity Support wants to help you troubleshoot your problem and interpret information in the log files. To speed up troubleshooting, you may be asked to provide the appropriate log files. These log files are often specific to your operating system. Use the **Support Tools** function to create a support package that can be sent to Cohesity Support.

See "Support tools" on page 810.

A debug log will provide the most verbose level of messages which includes record database operations, system processes, and errors that occur when executing a transaction. Use the **Set Logging Level** dialog to enable Debugging. Additional information is available about Log Files. See the *System Administrator Guide*.

Debugging Level Scope

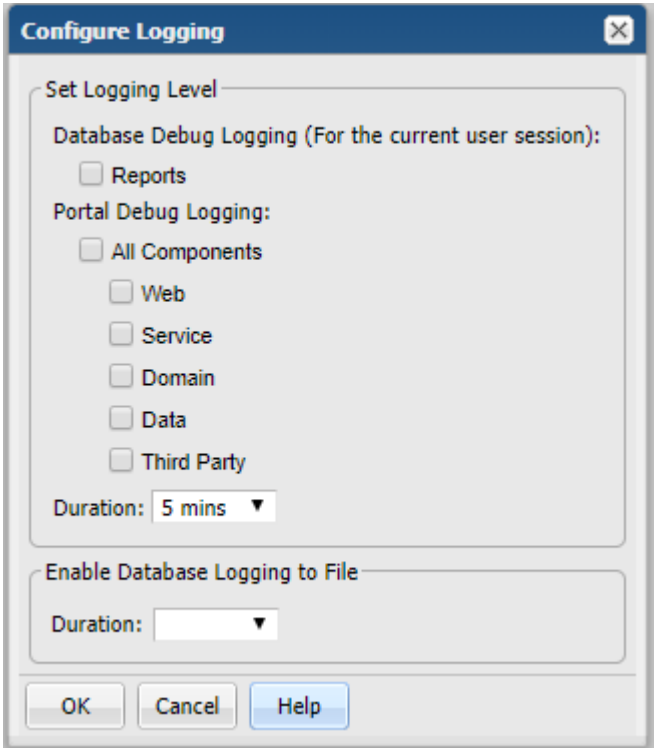
You can set debugging for an individual user session and/or on a system wide scope for all users. If you select system level logging, you can specify the Portal components to monitor.

To Access Debugging

- In the Portal, within a report window, enter the following key combination:

  Ctrl+Alt+D

Note: Do not activate the Debug logging without first consulting with Cohesity Support.



To Set Debugging for an Individual User Session

You can set debugging for an individual user session.

- Click **Reports** to enable/disable the database logging for the current user session.

To Set Debugging for System Wide Components

Enable/disable the specific logs you want to collect across the Portal.

- Select **All** unless you have been asked to only choose specific components to collect the related logs:

- Web
- Service
- Domain
- Data
- Third Party

To Reset the Logging Level

Because changing the log level can impact resources and ultimately system performance, a time period for resetting the level is required.

- Choose a time period for the logging to reset. This selection applies to both session-based logging and system wide. Once selected, the logging level will automatically reset to the default level when the time period is complete. The default is 5 minutes.

  Note: Session-based logging is automatically terminated when the user logs out as well.

To Enable Logging Database Messages to a File

The database message logging framework writes to a table by default, but you can write to a log file for a specified duration. This file can then be made available and downloaded using the Support Tools.

- Select a duration for all database messages to be written to a log file.

To Reset the Logging Level

Because changing the log level can impact resources and ultimately system performance, a time period for resetting the level is required.

- Choose a time period for the logging to reset. This selection applies to both session-based logging and system wide.

  Once selected, the logging level will automatically reset to the default level when the time period is complete. The default is 5 minutes.

  Note: Session-based logging is automatically terminated when the user logs out as well.

Login errors

Portal login errors

- File system is out of disk space.
- Fully qualified URL incorrectly set up
- URL not in the local hosts file or in DNS.
- Domain incorrectly specified
- Values in Tomcat Application Server and Apache Web Server do not match what's in DNS.

Cannot log on to Portal

When a user can't log in to the IT Analytics, the reasons are usually one of the following:

- User forgot password. Change the user's password as outlined in the Portal Online Help.
- LDAP service is not running.
- There is a port conflict. Another program is listening on port 80, resulting in a port conflict. Apache Web Server needs port 80. Use the **netstat** command to determine the other application that's listening on port 80, then assign that application a different port.

Troubleshooting recommendations

The following list suggests actions you can take to determine what is causing the issue.

- Review the portal log file: /opt/tomcat/logs/portal.log. Check for an "Exception" or "ERROR."

Add notes to failed jobs

As you monitor backup jobs, keep in mind that you can add notes to a job. You can use notes to:

- Specify the cause of a failed job, and how you or a third party resolved the problem, providing an excellent training tool.
- Provide an audit trail.

To add a note to a failed job, refer to the following.

Host name processing - Filters and aliases

A two-pass approach is taken to process hosts. See Load Host Alias in the technical documentation for the description of the utility to set up host aliases.

Pass 1: Filtering out characters

Specific characters can be filtered from the host name. Character sequences are stored in the database table: **apt_domain_alias_filter**

During Pass 1, the host name is compared to the character strings in the filter table.

- Only the suffix can be filtered--that is, trailing characters
- The filter must be at least 2 characters in length

Pass 2: Checking for aliases

After filtering, the **apt_object_alias** table is checked for aliases, looking for a direct match. Update this table with your host name aliases.

- Wildcards are not supported
- This is a one-for-one check. A lookup is performed with the host name and if an alias is found, the alias is substituted for the host name.

Note: Even if filtering is not required, this second pass still looks for aliases.

Example 1:

The following example demonstrates how host processing and filtering works:

1. Pass 1 - Host name (**myserver_nbu.mycompany.com**) is compared to filter string: **_nbu**
2. After filtering, the host name that remains is: **myserver.mycompany.com**
3. Pass 2 - If there was an alias set up for **myserver.mycompany.com** as **mymachine**, then mymachine is used as the host name.

Example 2:

Filter = nb1

Input host name = anb1bcdnb1

Output host name = anb1bcd

Change the Portal server’s IP address

If you want to test your Portal Server in a test environment, you’ll need to change its IP address when you migrate the Portal Server to production.

To change the Portal server's IP address

1. Update the DNS entries to point to the new IP address.
2. If any of the following files have the Portal Server IP address hard-coded, update the files with the new IP address or the localhost that would serve for both systems.

```
/opt/aptare/portalconf/portal.properties
            <URL>jdbc:oracle:thin@localhost:1521:scdb</URL>
/opt/aptare/datarcvrconf/datrarcvrproperties.xml
            <URL>jdbc:oracle:thin@localhost:1521:scdb</URL>
/opt/aptare/policyEngine/conf/systemconfig.properties
            <URL>jdbc:oracle:thin@localhost:1521:scdb</URL>
/opt/aptare/policyEngine/conf/policyproperties.xml
            <URL>jdbc:oracle:thin@localhost:1521:scdb</URL>
/opt/aptare/rptNotifyEngine/conf/rptnotifyproperties.properties
            <URL>jdbc:oracle:thin@localhost:1521:scdb</URL>
/opt/aptare/rptNotifyEngine/conf/rptnotifyproperties.xml
            <URL>jdbc:oracle:thin@localhost:1521:scdb</URL>
```

3. Change the IP address of the actual Portal Server.

Changing the port for Portal server communication

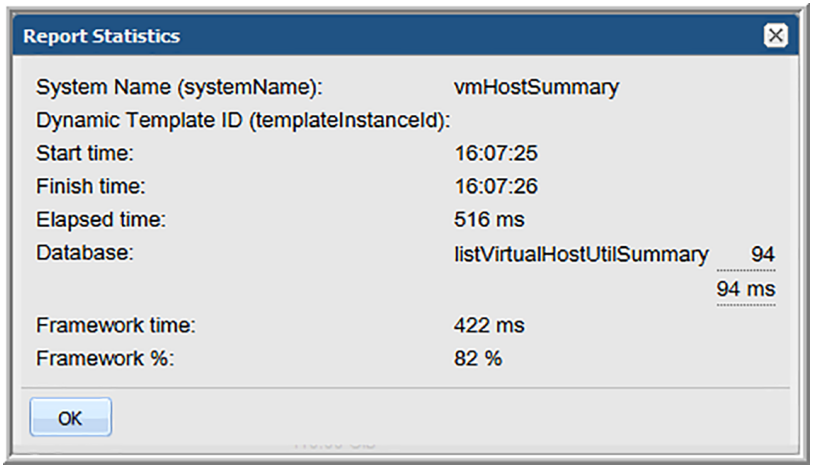
Apache HTTP Web Server is listening on one port on the Portal Server for incoming requests--port 80, by default. You can configure Apache HTTP Web Server to use an alternative port. However, all incoming requests--that is, requests from all agents, browsers, etc.--pass through the new port. The best practice is to choose one port for everything, but this isn't always feasible. If you want to have some agents use one port, and everything else use the default port 80, there is a solution.

To configure the Apache HTTP Web Server to use a different port on the Portal Server

1. Create a separate instance of Apache HTTP Web Server on a separate server using a standard install.
2. Update the Apache HTTP Web Server configuration to listen on the new port.
3. Configure that separate server to communicate with the Reporting Database.

Determine report rendering statistics

If a report seems to take a long time to render the results, use **CTRL+ALT+T** in the active browser window of the rendered report. These elapsed time and database statistics may be useful to someone troubleshooting a report.



Clearing the inventory and report cache

The Clear Cache function on the User Account menu is only displayed to those user's with a Super User role. This function is used by Cohesity Support for debugging purposes. It clears all cached Inventory objects and cached reports in the memory. After clearing the cache, the Inventory must fetch objects from the database. Reports must also fetch data from the database to render. This operation will slow down performance.

See "How reports and caching work together?" on page 839.

Troubleshoot Oracle performance issues in IT Analytics

IT Analytics Portal installer deploys a diagnostic script during the installation that helps you diagnose and analyze Oracle database performance issues in your environment. The script generates an output file, the contents of which you can analyze to diagnose the issues.

This section provides the commands to run the diagnostic script and provides the guidelines to use the script output to diagnose the Oracle database performance issues.

Script location

The diagnostic script is installed with the IT Analytics Portal at the following OS-specific locations.

- On Linux: `/opt/aptare/database/tools`
- On Windows: `c:\opt\oracle\database\tools`

Run the diagnostic script on Linux

Run this Shell script on your Linux portal server as Oracle user. Permissions available to a IT Analytics Portal user with admin privileges are adequate to run this script.

```
$ dbinfo.sh
```

Enter Oracle Database Service Name (for example SCDB), Database Username and Password when prompted by the script.

Run the diagnostic script on Windows

Run this batch script as an administrator on your Windows portal server.

```
dbinfo.bat
```

Enter Oracle Database Service Name (for example SCDB), Database Username and Password when prompted by the script.

Script output file and its contents

The diagnostic Shell or batch script produces a `dbperformanceinfo_<date>.txt` file, where <date> is a timestamp. The output file is a created at the same location as the script. The file contains the following important sections:

- Top 10 queries by highest execution time.
- Top 10 queries with full table scans.
- Queries with large I/O data buffers.
- Report of indexes that require a rebuild.

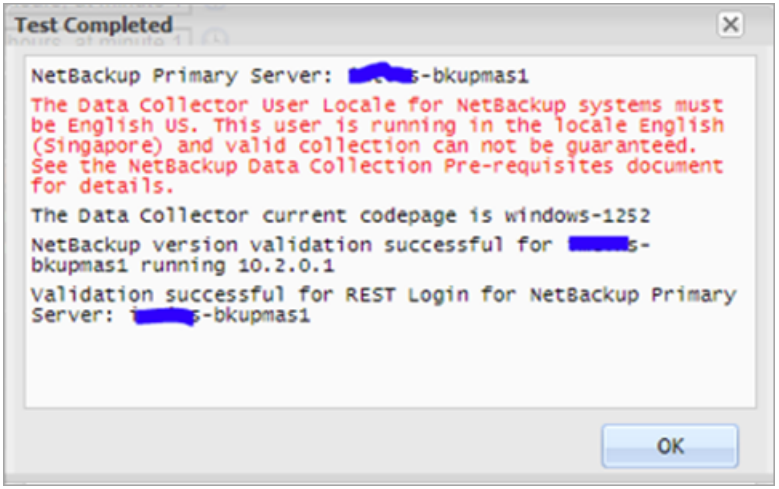
Diagnose performance issues

Look for the following information in the `dbperformanceinfo_<date>.txt` file to identify the causes of the performance issues:

1. Look for the most I/O intensive queries within the top 10 queries by execution time. Identify queries taking more than the anticipated time for investigation. Identify ad-hoc queries and tune the queries appropriately. From the MODULE column, verify whether the query was executing simultaneously with the Performance Script run. Also check rows processed and identify the rows with long process duration.
2. Analyze findings indicated by 'Waits' or 'Blocking' in the output file.
3. Identify queries doing Full Table Scans. Suppress or avoid running queries performing Full Table Scans.
4. Identify queries with large I/O buffers.

Troubleshoot test connection locale issue

Use these steps to resolve failure in test connection with the below error:



1 Create a new user with administrative permissions.

Reason for this new user creation is to copy English (United States) language and regional format to local system and local service user.

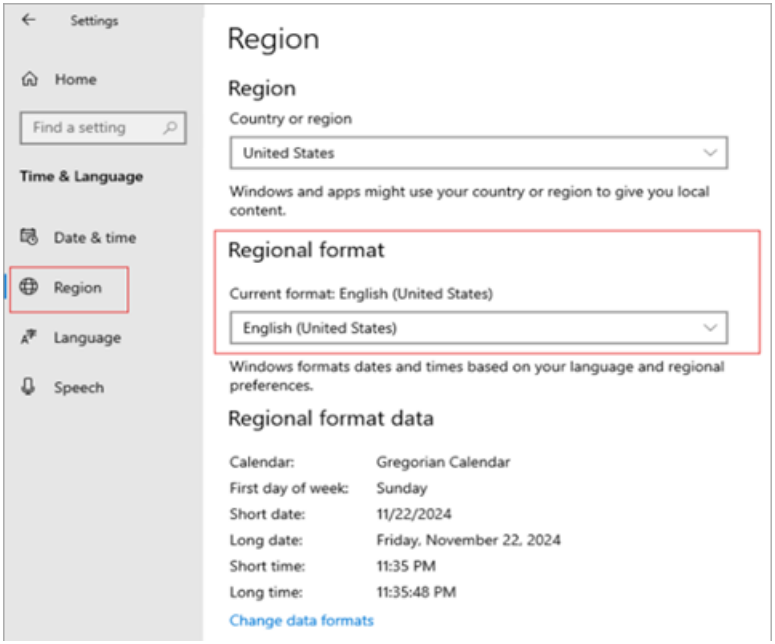
Sample command to create a new user and assign administrative permissions-

- `net user /add username password`
- `net localgroup administrators username /add`

2 Login to data collector machine with the newly created user

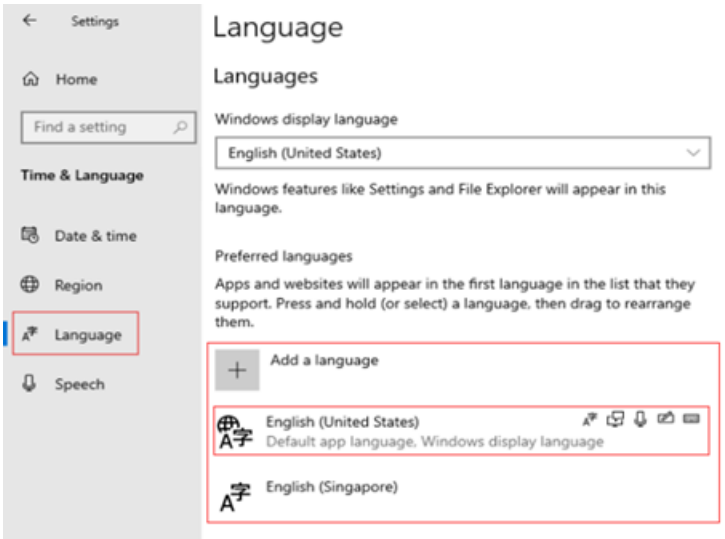
3 Go to **Windows** > **Settings** > **Time & Language** > **Region**.

4 Set **Regional Format** > **English (United States)**.



5 Go to **Windows** > **Settings** > **Time & Language** > **Language**.

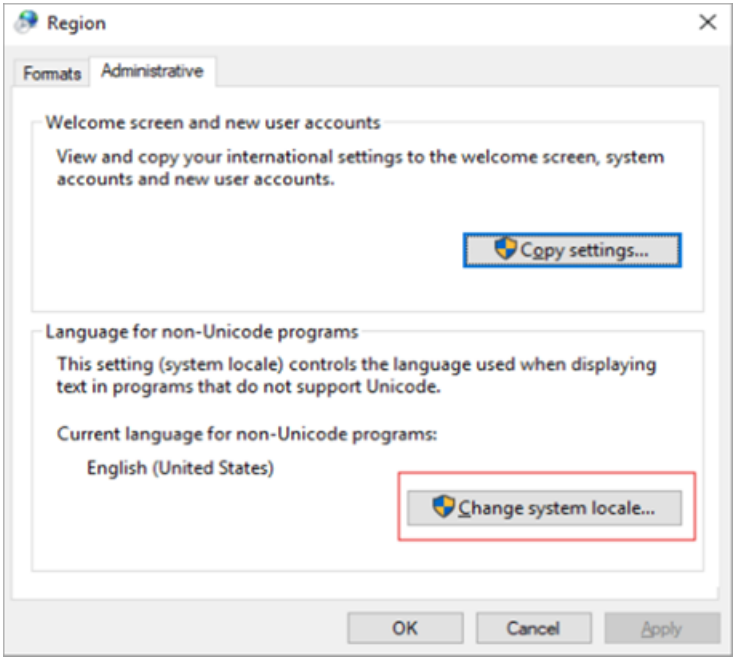
6 Keep **English (United States)** at the top of the list in case there are multiple languages. (Default language)



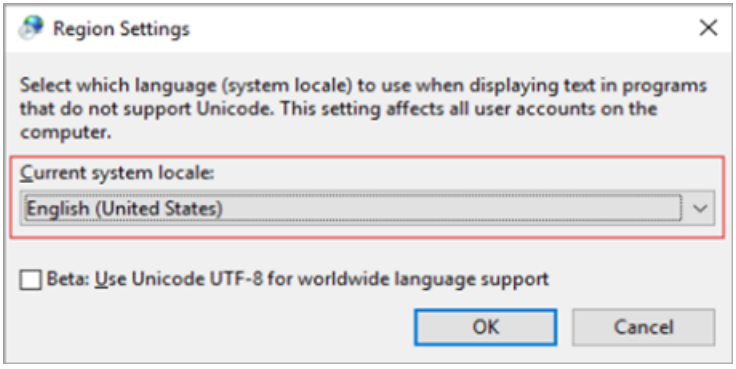
7 Click **Administrative language settings** on the same page.



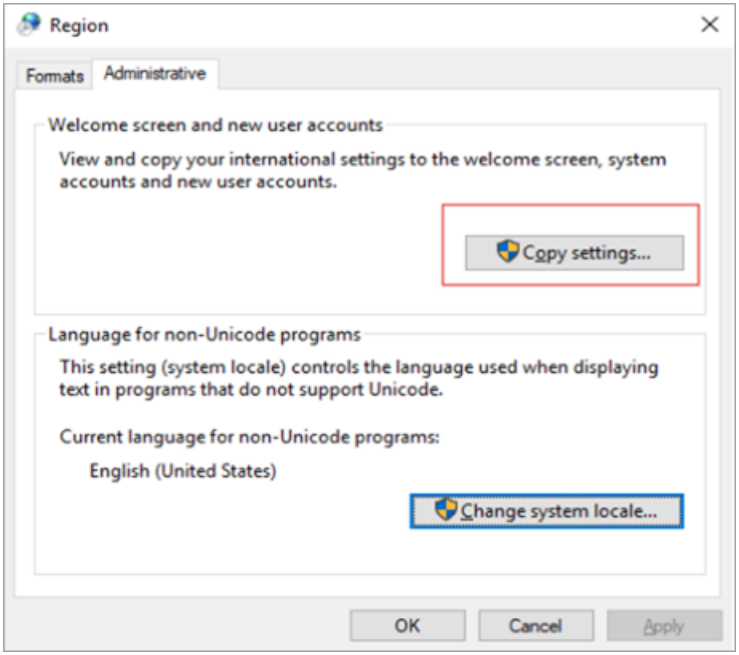
- Inside **Administrative** tab, **Section** > **Language for non-Unicode programs**, click **Change System Locale**.



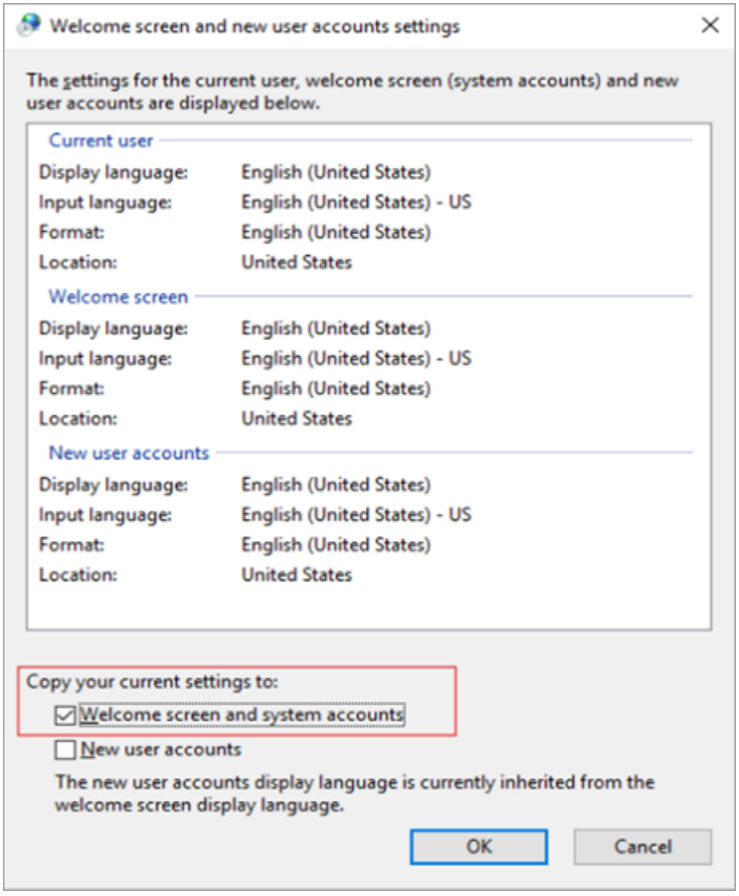
- Set **Current system locale** to **English (United States)** and click **OK**.



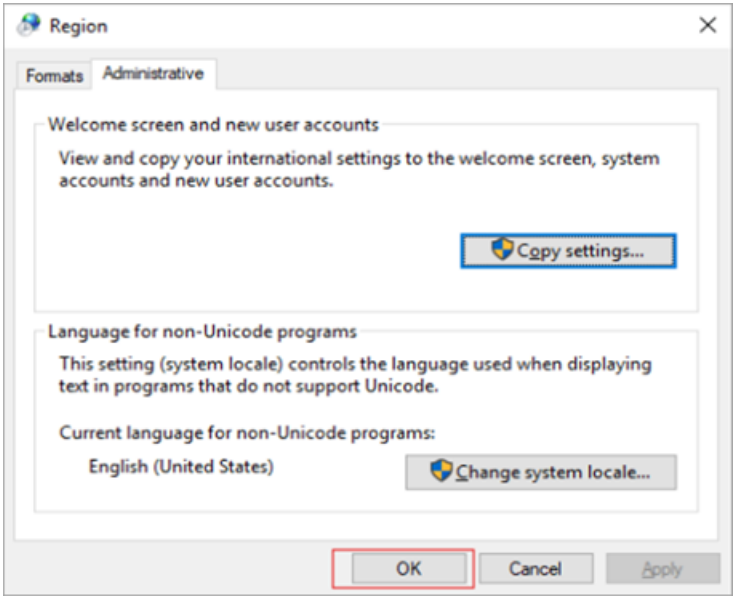
- On the **Administrative** tab, under **Welcome screen and new user accounts** section, click **Copy settings**.



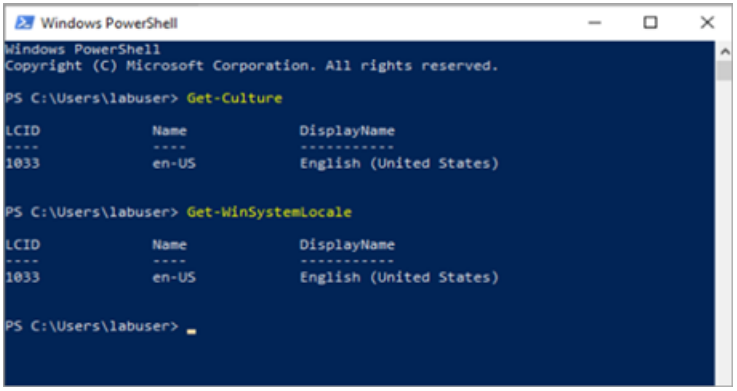
- Under **Copy your current settings to**, select **Welcome screen and system accounts** and click **OK**.



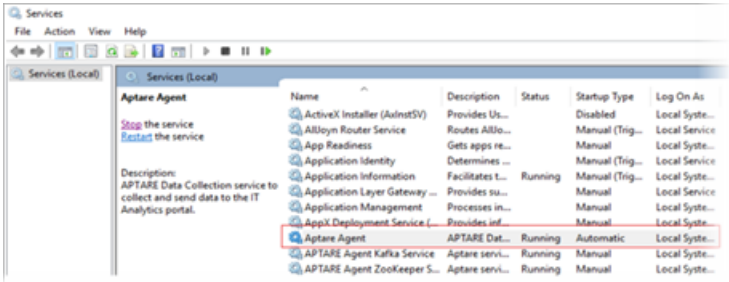
- Click **OK** again on the **Region** dialog box.



8 Open PowerShell command prompt and check responses for `Get-Culture` and `Get-WinSystemLocale`.



9 Stop and Start the Aptare Agent service from the Services window.



10 Delete the newly created user in step 2.

Chapter 27

Attribute inheritance overrides

This chapter includes the following topics:

- Overview
- Override inherited attribute values

Overview

Attribute inheritance is relevant primarily in multi-tenancy environments where domains are used to partition the database to maintain security controls. In this configuration, a hierarchical structure provides a parent-child relationship that controls access to data and also a structure for inheriting configurations from parents. For example, a Managed Services Partner (MSP) supports many client companies, each with its own domain.

See “Override inherited attribute values” on page 837.

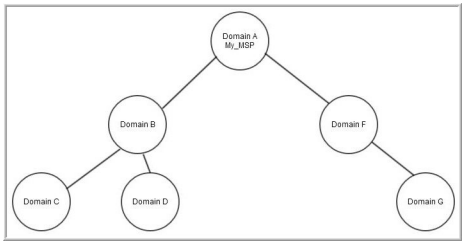
In the following diagram, attributes created at the MSP domain are inherited by the child domains. In this case, the child domains may want to override inherited attribute values because it is likely that each company would have its own values for an attribute.

See “Override inherited attribute values” on page 837.

Likewise, an enterprise may simply want to use domains to segment the data within its environment. For example, subsidiaries or divisions within an organization have different business assets and therefore different reporting needs. In this case, an enterprise might configure a root domain, with sub-domains for each of its subsidiaries. An attribute defined for the parent company’s domain will be inherited by its subsidiaries.

Use the diagram to visualize a hierarchy that supports attribute inheritance.

Figure 27-1 Hierarchical Domain Structure for Attribute Inheritance



Attributes have the following inheritance characteristics, based on a domain hierarchy:

- An attribute at the parent level is visible to all of its children.
  Example: An attribute created at Domain A will be visible to all other domains.
- Children inherit attributes and values from any domain that is higher in its hierarchy path.
  Example: An attribute created at Domain B will be visible only to Domains C and D. Likewise, an attribute created at Domain A will be visible to both Domains F and G.
- Duplicate attribute names are not allowed in a single hierarchy path. When creating an attribute, the system checks for duplicate attribute names and it will not create the attribute if it already exists in either a child or parent domain.
  Example: Duplicate attribute names cannot exist in the A-F-G Domain path.
- Duplicate attribute names are allowed in sibling hierarchies.
  Example: Domain B could have an attribute that is a duplicate of a name in Domain F.

Override inherited attribute values

Often, attributes inherited from a parent to a child provide necessary report filtering. However, the parent's attribute values may not be relevant for the child company. In this case, the child can override the list of values inherited from its parent.

Note: Overriding attributes is relevant only in multi-domain/multi-tenancy environments.

Use case 1

Consider a Managed Services Partner (MSP) that supports many client companies. In this scenario, the MSP might have an attribute named Client with values that list all of its client company names. The MSP does not want this list to be publicly available to all of its clients. In this case, the MSP can choose one of the following options:

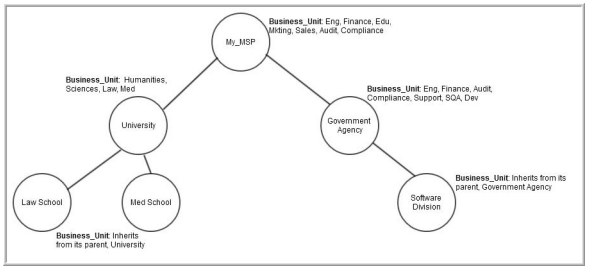
- Uncheck the Allow Inheritance box so that sub-domains cannot see the attribute in the Inventory or in a report scope.
- Override the values in each of the client domains so that they cannot see the root domain's values. An administrator at the parent domain or the client domain can override the values for the attribute at the client's domain.
  See " Manage attributes " on page 540.

Use case 2

Likewise, as shown in the diagram, the Business Units listed for the root domain may not be relevant for the child domains. In this case, the child domains need to have an override of the list of values.

See Figure 27-2 on page 838.

Figure 27-2 Use case for Overrides of Attribute Values



Note: For inherited attributes, you are not permitted to modify attribute names in child domains. However, you can modify/override the list of values and you can enable/disable inheritance for domains that are lower in the hierarchy.

See " Edit or rename attributes " on page 544.

Chapter 28

Understanding report data caching

This chapter includes the following topics:

- Understand report data caching
- How reports and caching work together?
- Store and purge data
- Disable optimistic caching with portal.properties

Understand report data caching

When working with large amounts of data, system performance can be an issue. In some environments, your Portal server may not be able to view the information in a timely manner.

Performance issues are addressed when working with large amounts of data by employing different solutions - one is caching reports. Cached data is used to improve Portal performance. Caching allows the system to present the same report without having to build it from scratch (building a new report includes making a call to the database which can impact the time it requires to render).

The caching mechanism is implemented on the Portal server. Refer to the following topics for details:

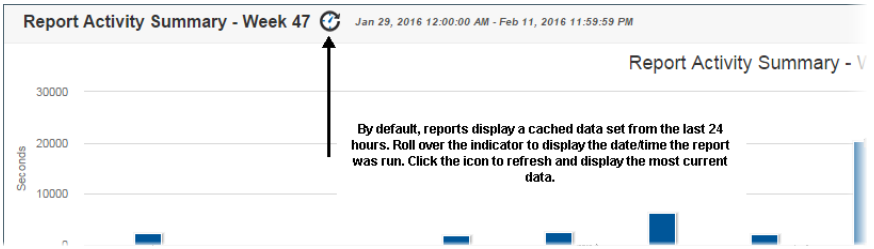
How reports and caching work together?

By default, reports display a cached data set from the last 24 hours.

When you run a report, the scope of the report is read, and is checked if the cache contains the same report, for the same scope. If it does, the results are displayed from the cache. If the combination does not exist, the report is run from the database, saved in the cache, and then sent to the user interface. Cached reports are shared across users who belong to the same home host group.

When a report is served from the cache, an indicator icon is displayed on reports and dashboards. You can roll over the indicator to show the age of the report from the cache. Click the icon to purge the old report from the cache and rerun the report from the database. You can manually refresh at any time.

See "Getting started with the Inventory navigator" on page 61.



When a report is run, and it's pulled from the cache, it does not show in the Report Activity Summary or the Report Activity Detail reports. These reports only show reports that ran and actually executed against the database.

Note: Scheduled, emailed and exported reports are not derived from the cache. These events are run in real-time, so current data is always used. Caching is only applicable for reports run in the browser.

Store and purge data

The cache can retain up to 0.5 GB of reporting data and if it reaches capacity, it frees up space for new reports by purging the data for the least frequently used reports. You can change the retention value by revising the portal.properties parameter: portal.reports.cache.maxSizeInMemory.

Purging also occurs when:

- Portal Tomcat service is stopped
- A cached report is more than 24 hours old

You can change the time period by revising the port.properties parameter: portal.reports.cache.timeOut.

The cache can be stored in memory. By default, it is located in:

Linux:

/opt/tomcat/aptare_instances/portal/temp

Windows:

C:\opt\tomcat\temp

Optimistic caching is used to pre-generate reports related to Inventory objects a user may not have selected. For example, if you select a category in the Inventory with 12 VM Servers, and choose the first VM Server, the next nine servers displayed in the Hierarchy Panel are queued up to have their inventory reports run and populated in the cache. This enables a quick return of reports and almost instant results.

Note, optimistic caching is designed to limit the load to the CPU of the Portal server. This means, that in the case of a large number of objects with a large number of associated reports, the caching speed is throttled down.

The mechanism can be fine tuned using guidelines provided by Cohesity Support.

See “Clearing the inventory and report cache” on page 826.

Disable optimistic caching with portal.properties

Administrators can disable optimistic caching using a portal.properties parameter: portal.ocn.optimisticReportCachingOn=false.