COHESITY

Best Practices Guide for Protecting Microsoft SharePoint Server with Cohesity and Ontrack PowerControls for SharePoint

COHESITY

Abstract

This white paper outlines the best practices for using Cohesity to provide data protection and granular recovery for a Microsoft SharePoint environment

Table of Contents

About This Guide

Hyperconvergence is becoming a norm in data centers today. Companies adopting this next generation infrastructure have realized significant savings in TCO/ROI. These savings are the result of vastly simplified architectures, lower power and cooling needs, workload consolidation, smaller hardware footprint and "pay as you grow" consumption model.

SpanFS is a completely new file system designed specifically for secondary storage consolidation. At the topmost layer, SpanFS exposes industry-standard, globally distributed NFS, SMB, and S3 interfaces. Cohesity is unique in its ability to support unlimited, frequent snapshots with no performance degradation. SpanFS has QoS controls built at all layers of the stack to support various workloads and can replicate, archive and tier data to another Cohesity cluster or to the cloud.

What ties all these benefits together is the simplicity of managing these web scale platforms from a single UI. The design principals of distributed control and data planes that eliminate complexities in infrastructure and management make hyperconverged architectures attractive and bring overall value to end customers.

Cohesity along with Ontrack PowerControls can provide a robust, scalable, and simple to administer data protection solution while also allowing for seamless growth. Cohesity provides a globally deduplicated, scale-out data protection solution that is natively integrated with the public cloud. This document describes how to configure and use Cohesity to backup and perform granular recovery of MS SharePoint environments using Ontrack PowerControls. Granular recovery of MS SharePoint content, such as folders, libraries, documents, etc. is simple and intuitive.

Intended Audience

This paper is written for Microsoft SharePoint Administrators, Microsoft SQL Server Administrators, and/or IT Administrators planning to utilize Cohesity for data protection of Microsoft SharePoint environments when performing backups and restores.

Cohesity recommends having familiarity with the following:

- Cohesity DataPlatform
- Microsoft Windows Server
- Microsoft SQL Server
- Microsoft SharePoint

Terminology

Cohesity DataPlatform

Cohesity DataPlatform is the industry's only hyperconverged solution designed to simplify secondary storage by consolidating all secondary data on a single, web-scale platform.

Cohesity DataProtect

Cohesity DataProtect is an end-to-end data protection solution that is fully converged on Cohesity DataPlatform. DataProtect simplifies your data protection environment with a single unified solution for backup, recovery, replication, disaster recovery, target storage, and multi-cloud integration.

Cohesity Storage Domain

Cohesity Storage Domain (View Box) represent storage efficiency domains within the Cohesity cluster and can optionally associate a Storage Domain with a specific cloud tier. The administrator assigns deduplication, compression and encryption attributes when Storage Domains are created.

Cohesity View

Cohesity Views represent mount points into a specific Storage Domain (View Box).

To find out more about Cohesity components and terminology, click on Help or "?" at the top right hand corner of the Cohesity UI Dashboard.

Remote Blob Store (RBS)

Microsoft introduced the ability to store BLOBs outside of MS SQL Server database instances. RBS does not reduce the total space needed or used by SharePoint, rather it can offset performance impacts potentially caused by larger BLOBs. If using or considering using RBS, it's important to understand the benefits as well as limitations or RBS with SharePoint and MS SQL Server. See Deciding to use RBS in SharePoint 2013 [technet.microsoft.com]

Abbreviations

Abbreviation	Description Description
SMB	Server Message Block
BLOB	Binary Large Object
RBS	Remote BLOB Storage
CBT	Change Block Tracking
RPO	Recovery Point Objective
WSFC	Windows Server Failover Clustering
FCI	(MS SQL) Failover Cluster Instance
UI	User Interface

Solution Components

The following components were used for interoperability testing

MS SharePoint Farm	MS SQL Server	MS Windows Server
MS SharePoint 2010 SP2	MS SQL Server 2008R2 SP3	MS Windows 2008R2 SP1
MS SharePoint 2013 SP1	MS SQL Server 2012	MS Windows 2012R2
MS SharePoint 2016	MS SQL Server 2016 SP1	MS Windows 2016

Component	Component Version
Windows 2016 WSFC /w MS SQL Server 2016 SP1	
MS SQL Server 2016 (FILESTREAM) Remote BLOB Store	13.1.4001.0
Cohesity DataProtect	4.1.2, 5.0.0
Ontrack PowerControls	9.0.1

The following components were not validated

MS SQL Server 2008 (FILESTREAM) Remote BLOB Store

MS SQL Server 2012 (FILESTREAM) Remote BLOB Store

Cohesity Overview

Cohesity introduced the world's first scale-out data management platform to enable organizations to standardize secondary workflows on a unified and fully distributed solution. Cohesity's scale-out distributed file system SpanFS™ which was built from the ground up to ensure complete scalability to enable organizations to flexibly grow their environment by adding nodes to a cluster. With this scalability, organizations can eliminate the costs of data migrations and forklift upgrades, while benefiting from the simplicity of a homogenous solution. SpanFS also provides global, variable-length deduplication and unlimited snapshots and clones - making it the ideal storage target for enterprise environments.

Cohesity is well suited to provide data protection for Microsoft SharePoint because it provides:

- A single and unified interface for provisioning, managing, and monitoring backup jobs
- Native and application consistent Microsoft SQL Server database backups
- Ability to restore SharePoint folders, libraries, documents, etc. from any point in time backup
- Simple, easy to use and yet powerful granular recovery UI provided by Ontrack PowerControls
- Variable-length, post-process or in-line, global deduplication
- Non-disruptive Cohesity hardware refresh and expansion without downtime

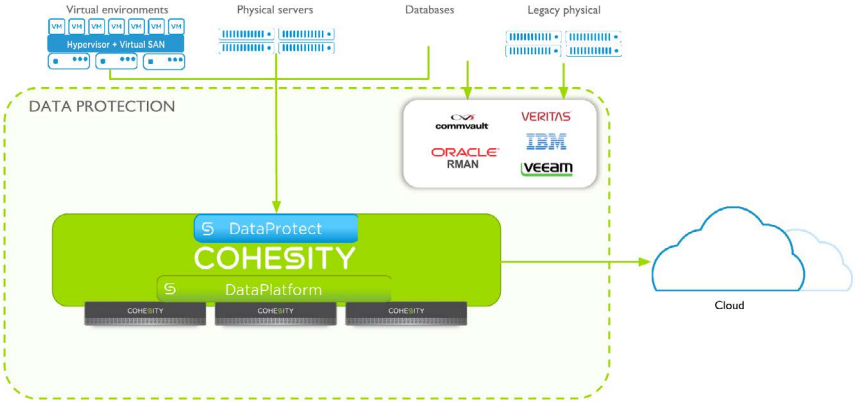

SharePoint Overview

SharePoint is a web-based platform for collaboration and document management. Microsoft SharePoint is very configurable and thus how it's deployed and used varies between organizations. Microsoft SharePoint can be deployed on premises or through the cloud. For the purpose of this document, we'll be focusing on protecting on premises instances of SharePoint.

Ontrack PowerControls Overview

Ontrack® PowerControls™ for Microsoft SharePoint Server, is a powerful, yet easy-to-use application for searching, recovering, restoring, and managing data at a granular in SharePoint Server environments. With Ontrack PowerControls, you can perform source and target functions on Microsoft Office SharePoint Server data. It also provides a robust search function on source items to quickly and easily find, recover, and restore content – such as documents, lists, libraries, and folders – or entire Microsoft Office SharePoint Server sites. Ontrack PowerControls works previous backup copies of your database allowing you to restore content from MDF files, NDF files, and LDF files and restore directly back to your Microsoft Office SharePoint Server target, or to a different server.[1]

Ontrack PowerControls Benefits

1. **Minimizes the time to restore an individual document, item, list, folder, site, or library.** Ontrack PowerControls can slash restore time, making it possible to granularly restore items from a previous full backup. Items from database backup can be restored individually, giving you restore flexibility and eliminating the need to perform to a full site restoration in order to find the items you need.
2. **Eliminates same site restore constraint.** By eliminating same site restore constraint, Ontrack PowerControls gives you the flexibility to restore to origination or alternate locations. With Ontrack PowerControls, you can restore directly into your production Microsoft Office SharePoint Server or a file system.

3. **Minimizes the time to locate all matching specific criteria.** Ontrack PowerControls includes an Advanced Find feature that can search across all content databases in an archived Microsoft Office SharePoint Server file, rather than bringing an old backup back online for analysis. You can search by a variety of criteria, including keywords, subject, date, specific users, and file and attachment data.
4. **Maintains data integrity of the Microsoft Office SharePoint Server source.** Ontrack PowerControls does not change the contents or metadata of the source content database, and maintains data integrity of the source by performing read-only operation.
5. **Minimizes the costs involved in restoring lost items.** Ontrack PowerControls eliminates the need to have an expensive restoration server available.[1]

Kroll Ontrack PowerControls Rebranding

In early 2018, the Ontrack PowerControls parent company Kroll Discovery rebranded itself as KLDiscovery and along with it, rebranded Kroll Ontrack to simply Ontrack. The press release can be found here: https://www.ontrack.com/resources/press/details/65119/kroll-ontrack-announces-global-rebrand-to-ont/

Logical Flow

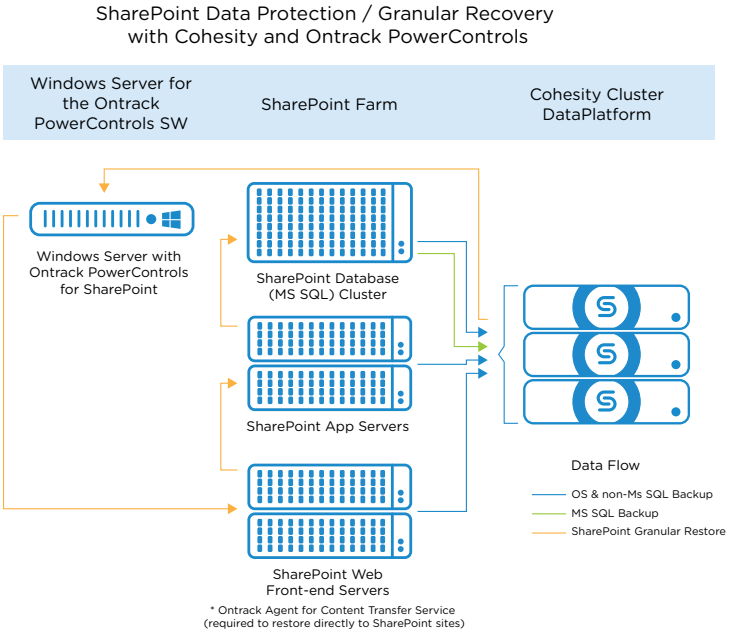

The above diagram shows the logical data flow and relationship between the SharePoint farm, the Cohesity Cluster, and Ontrack PowerControls.

Backup Data Flow

Application Consistent backups are taken of the SharePoint SQL Database using Cohesity's native and integrated MS SQL Data Protection, along with the SharePoint Web front-end and Application servers.

Restore Data Flow

The restore process involves presenting a specific Point-in-Time backup of the SharePoint SQL Database to a MS Windows server that has the Ontrack PowerControls application installed. This is done by leveraging Cohesity SnapTree technology. SnapTree provides fully hydrated snapshots to recover quickly to a certain point in time. A Cohesity SnapTree clone view of a previous MS SQL backup is a read/write copy that is presented to the original or different server. This is done in a quick and very space efficient manner. Ontrack PowerControls can connect to the offline database and live SharePoint site to facilitate restoring one or more items.

Best Practices

For the purpose of this documentation and examples, all servers will be physical. The actual workflow for performing granular recovery is identical for physical and virtual environments.

Creating the View Box

Create a suitable View Box, for the best de-dup performance, the same View Box should be used for multiple SharePoint servers and databases backup jobs.

Ontrack PowerControls Server

Creating or Identifying a Server for the Ontrack PowerControls Software

The Ontrack PowerControls software needs to be installed on a Windows server and will be the UI used for selecting and restoring objects from a SharePoint backup on Cohesity. This server is only used when doing granular level recovery and is not in the critical path for backups.

Multiple different versions of Windows operating systems are supported and the server does not necessary need to be dedicated only for launching and using the Ontrack PowerControls software. However, it may make sense and Cohesity would suggest dedicating a specific server just to keep things simple and to avoid conflicts with a server being used for several different purposes. This server does not need MS SQL or SharePoint installed. It should however be part of the same AD domain as your SharePoint environment.

The only requirement for the server used for launching and doing granular recoveries with Ontrack PowerControls and Cohesity DataProtect is it must be configured as same type of server as the SharePoint MS SQL Server within Cohesity e.g. Physical or Virtual and it's OS listed as supported in the Ontrack PowerControls ReadMe.

Source SharePoint MS SQL Server		Ontrack PowerControls Server
Virtual	→	Virtual
Physical	→	Physical

If the SharePoint MS SQL database is physical, say for example a physical SQL Cluster, but installing another physical server for Ontrack PowerControls is not desirable, it is possible to simply install the Cohesity physical agent on a Windows VM and register it within Cohesity as a physical server.

Refer to the Ontrack PowerControls ReadMe for the complete list of supported environments.

Installing Ontrack PowerControls Software

Locate and launch PC901.exe. This is a self-extracting ZIP file that will place installers and documentation in the location it's initiated from. It's a good idea to view the Ontrack PowerControls ReadMe and User Guide. The ReadMe will list out the limitations of what can and can't be done between different versions of SharePoint. We'll now install Ontrack PowerControls on our Windows 2012 R2 server we've dedicated for Ontrack PowerControls.

Click Next> To continue with the install

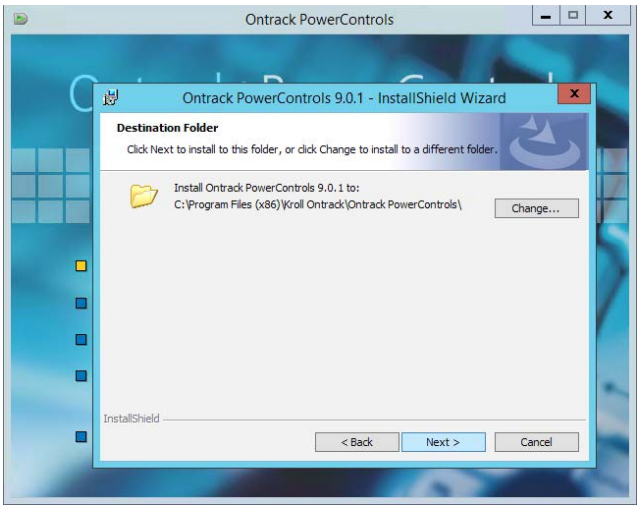

Accept the default or choose another location. The Ontrack PowerControls takes up less than 100MB of space.

Click Next >

You may choose to install all of the Ontrack PowerControls (SharePoint, Exchange, MS SQL, etc) or only install the applications you need/want. Click Next >

Click Install.

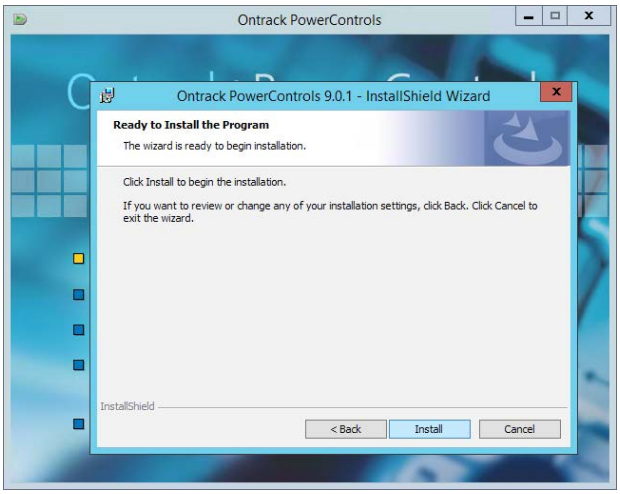

Click Finish.

Applying The Ontrack PowerControls License

Once it's first launched it will be running in eval mode. We'll need to apply the license to unlock all the features. The license file should be provided to you though your Cohesity SE.

Place the license.ini file somewhere where it's readable by anyone who will be using Ontrack PowerControls. In this example, we will place it under on the C: drive however it can be located anywhere that make sense.

Let's launch Ontrack PowerControls

Once Ontrack PowerControls for SharePoint is launched, you'll see the 30 Days Remaining Free Trial message. We'll go ahead and apply the license.

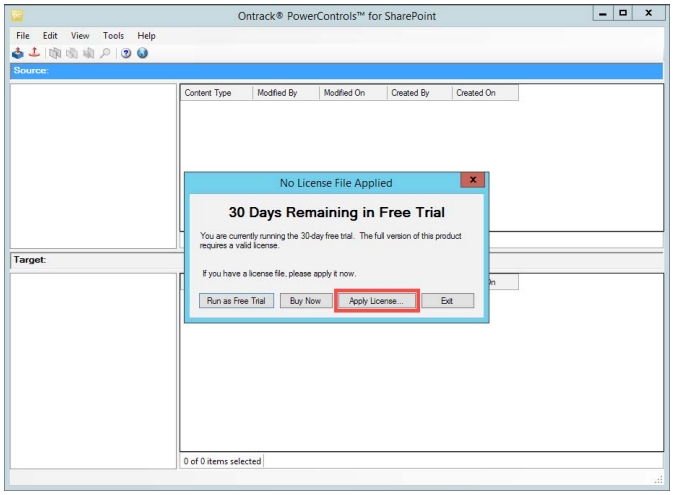

Click on Apply License...

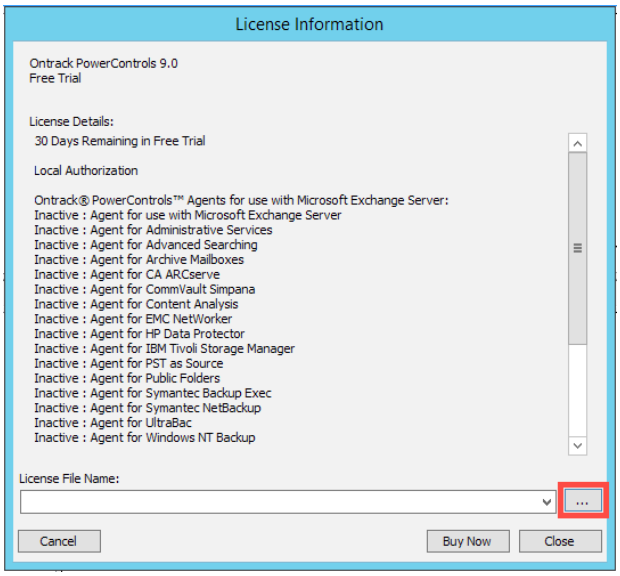

Click on the … box/button and navigate to where we put the license.ini file.

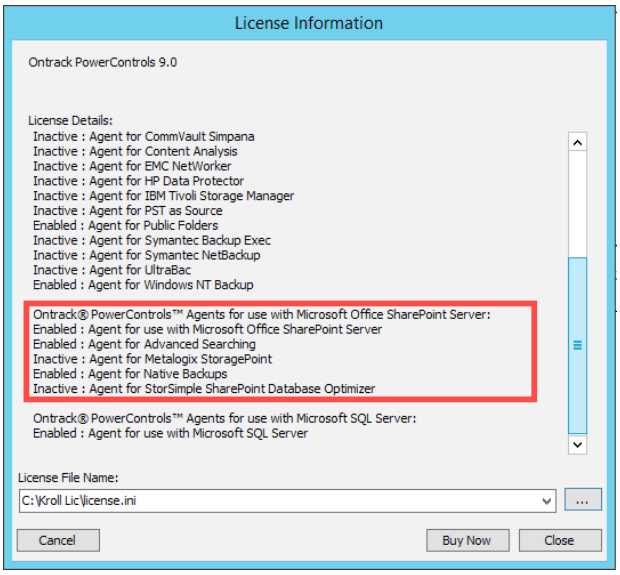

Scroll down and verify the appropriate features/agents are enabled under “Ontrack PowerControls Agent for use with Microsoft Office SharePoint Server” as seen in the screenshot above.

Click Close.

If another Admin logs in with a different Windows account to this server and launches the Ontrack PowerControls application, they will need to also apply the same license.

Note: The windows server will need access to the Internet as Ontrack PowerControls does reach out to its license server to verify the license. Direct Internet access is not needed as this license check can also happen if a web proxy is used and configured within IE.

At this point the last thing we need to do it register this server to the Cohesity DataProtect cluster.

Install the Cohesity Agent - Physical Servers

Login into the Cohesity cluster

Click on Protection > Sources.

Download the Physical Agent.

Launch the Physical Agent installer.

Click Next > to continue.

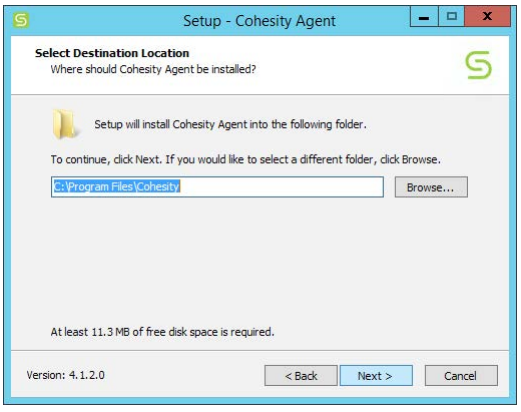

Choose where you'd like the agent installed, or just click Next > to take the default.

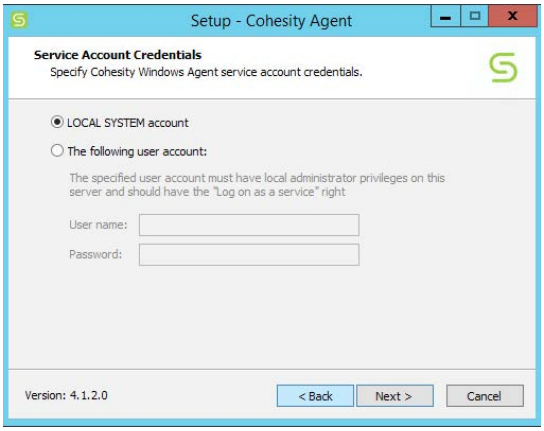

Choose what user you'd like the agent to run as and click Next > to continue.

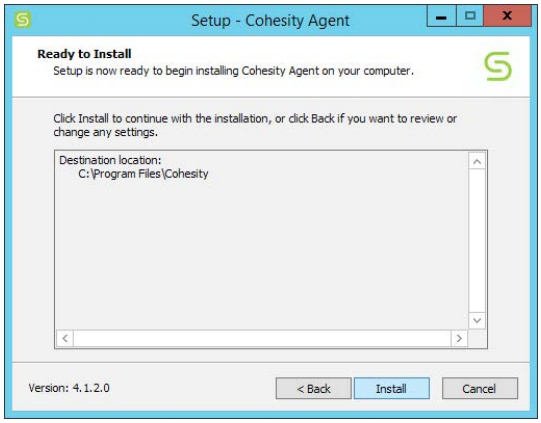

Click Install to complete the install.

If the server is not being used, you can go ahead and reboot to enable the CBT driver. Otherwise it can be done during a scheduled maintenance window. You can register and backup the server without it first being rebooted to enable the CBT driver, however until a reboot is done, each backup will be a full vs incremental. As the Cohesity DataProtect cluster will de-dup, this won't adversely affect the raw space used on the Cohesity cluster, only the amount of time required for the backup job to run will be longer without the CBT driver being active.

We can now register this physical Windows server to the Cohesity DataProtect cluster.

Register the Ontrack PowerControls Server

Click on Protection > Sources

Click on the Register Source button

Click on the Physical Server button

Enter the (DNS) name of the Physical server and click Register Server

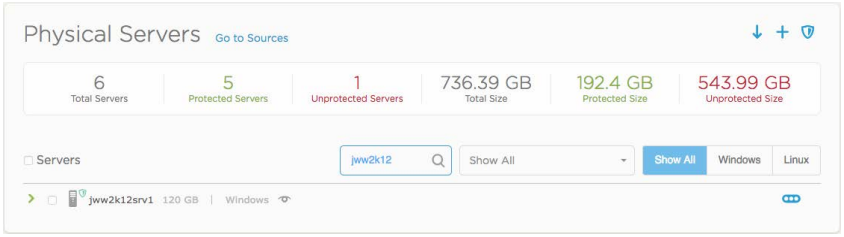

We can see the server in the list of Physical Servers.

Create a Cohesity Job and Add the Ontrack PowerControls Windows Server

Although it's not necessary to add the Ontrack PowerControls server to a Cohesity backup job, it's always a good idea in case any OS or Application files need to be recovered. It's not needed to do granular recovery of SharePoint.

Ontrack PowerControls Agent for Content Transfer Service

Ontrack PowerControls includes the Ontrack PowerControls Agent for Content Transfer Service, or ACTS, to support the authentication and communication requirements for restoring objects to a live SharePoint environment. ACTS is installed on the Web front-end of each site where a restore may be directed. If multiple Web front-ends are load balanced for high availability, the load balancer should either be bypassed or be configured for persistent (sticky) sessions.

Note: ACTS does not require additional licensing.

On each SharePoint Web front-end server, locate and launch the Ontrack PowerControls installer, in this example PC901.exe.

Click on the Install Ontrack PowerControls Agent for Content Transfer Service.

Click Next > to continue.

Read the license agreement and click Yes to accept and continue.

Enter your username and company name and click Next >.

Click Browse to change the destination folder, or simply click Next > to accept the default.

As indicated, enter in an AD account for the SharePoint Farm Admin you'd like used for this service.

Leave the port number as is or change it if desired and click Next > to continue.

Click OK to continue with the installation.

Click Finish. This should be completed for each SharePoint web server in the farm.

At this point, everything is in place for doing granular recovery from backups.

Backups - Physical Servers

MS SQL Backups

Backups of the SharePoint MS SQL databases is crucial for any comprehensive SharePoint data protection solution. Ontrack PowerControls will use a copy of these backups to do restores.

The MS SQL Backups of the SharePoint databases are the same as any other MS SQL database backup job on Cohesity. Follow the same procedure for installing the Cohesity agent as done on the Windows physical server where the Ontrack PowerControls.

After the MS SQL Server is registered, indicate it's a MS SQL server so that Cohesity knows it's a MS SQL server and it can be added to a new or existing MS SQL job.

Use the to the right of the SharePoint MS SQL Server to Register as MS SQL Server.

A new MS SQL job can then be created with the SharePoint MS SQL Server, or it can be added to an existing MS SQL job.

In this example, the SharePoint MS SQL server is part of a MS SQL job where it's setup to do a full backup every hour. Your RPO requirements will dictate the settings for your Job Policy and thus the frequency of SharePoint Database backups.

The full backups will be used for doing granular recovery restores. Log backups are not used for doing restores with Ontrack PowerControls, however they are still important and should be backed up by Cohesity, if full logging (Recovery model: Full) is enabled for the database. Although not covered in this document, the SQL logs could come into play if a full SharePoint content database restore is needed.

It's always good to configure the Cohesity cluster to send alerts in the event backup jobs fail and/or to periodically review the job run history to verify there are no issues.

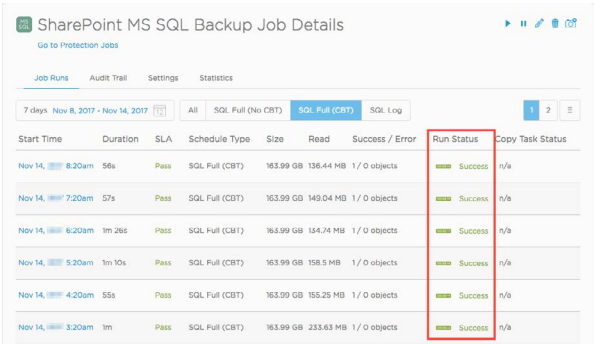

We can see the job runs are all successful.

If we dig into one of the job runs, we can verify all the SharePoint databases are seen and will be backed up.

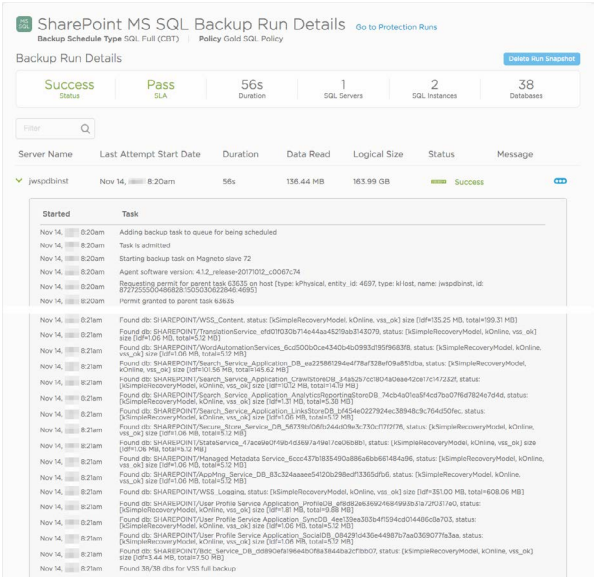

We can see that Cohesity was able to query MS SQL and will be backing up WSS_Content and the other SharePoint specific databases.

SharePoint Application and Web Server Backups

Although not as critical as SharePoint MS SQL database backups, the SharePoint web and application servers should be backed up as well to allow restoring critical files in the event of corruption or other data/file loss.

Granular Recovery / Restore

There are a number of options when it comes to restoring SharePoint content. For example, documents from a document library can be restored back to the original document library, to a different document library, or even exported from any previous backup to a file system.

The granular restore involves only a few steps:

- Use the Instant Mount Point Recovery to present a copy of a full MS SQL backup to a Windows server
- Use Ontrack PowerControls to open the mdf/ldf/ndf [source] files and connect to the [target] SharePoint Environment
- Locate objects to be restored from the [source] and the destination [target] where they are to be restored to
- Drag and Drop

The underlying process of the Instant Mount Recovery presents a snapshot / clone from a selected backup to the chosen target server where the Ontrack PowerControls application is installed. Any changes made to the files presented on this clone will have no effect on the original backup. Cohesity's SnapTree technology allows near instant cloning that is also space efficient. So even if a 10TB MS SQL database is cloned for the purpose of doing an Instant Mount Recovery and granular recovery, almost no net additional raw space is consumed. If some changes take place on this clone that result in space usage, that space will be reclaimed once the clone is removed as part of the tear down process and background garbage collection process.

Let's run through a possible data loss situation. We'll use this scenario: Some invoices were mistakenly deleted yesterday, the SharePoint recycle bin was emptied, and we need to be able to recover these files back to their original location.

This is what our Invoices list looked like yesterday, before the last 3 items in the list were deleted.

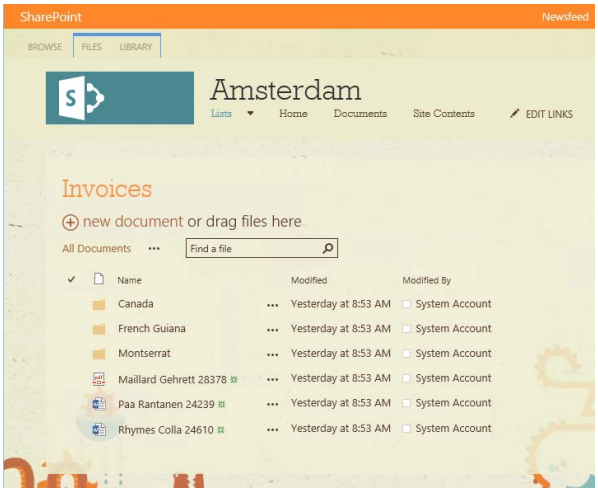

Now we're missing those last three files.

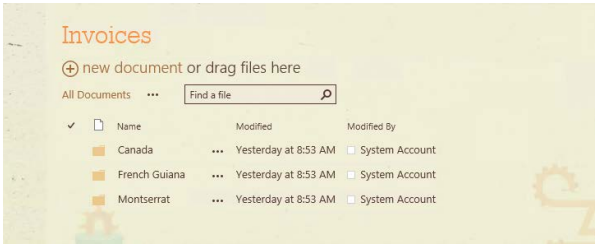

After checking the recycle bin, we do find the files have been permanently removed and can't be recovered through the recycle bin.

The first step to recovering these files is to determine which backup we will use as our restore point. We know the files where their yesterday morning and that they were deleted last night. So we'll use the backup that happened at 3:20pm yesterday.

We can log into our Cohesity cluster and use the Protection > Recovery > Recover > Instant Volume Mount.

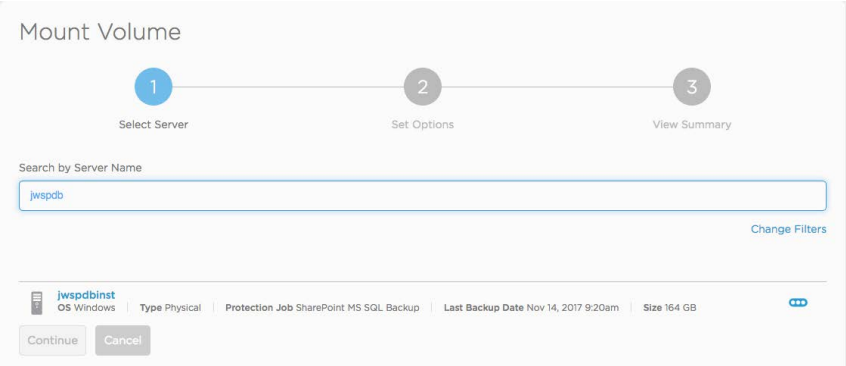

We'll search for our MS SQL SharePoint DB Instance, we don't need the full name, just part of the name will work. Click on the server/instance name.

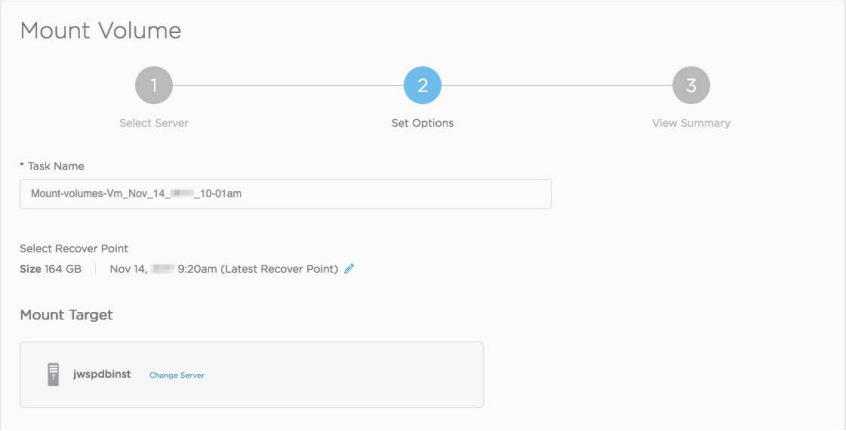

By default, Cohesity will default to the latest snapshot and the target will be the original server from which the backup happened. We'll want to change this and pick our specific backup point in time and use the Windows server which we've dedicated for Ontrack PowerControls. The Task Name can be left as is.

In the following step, choose the backup that was taken before the files were deleted from SharePoint.

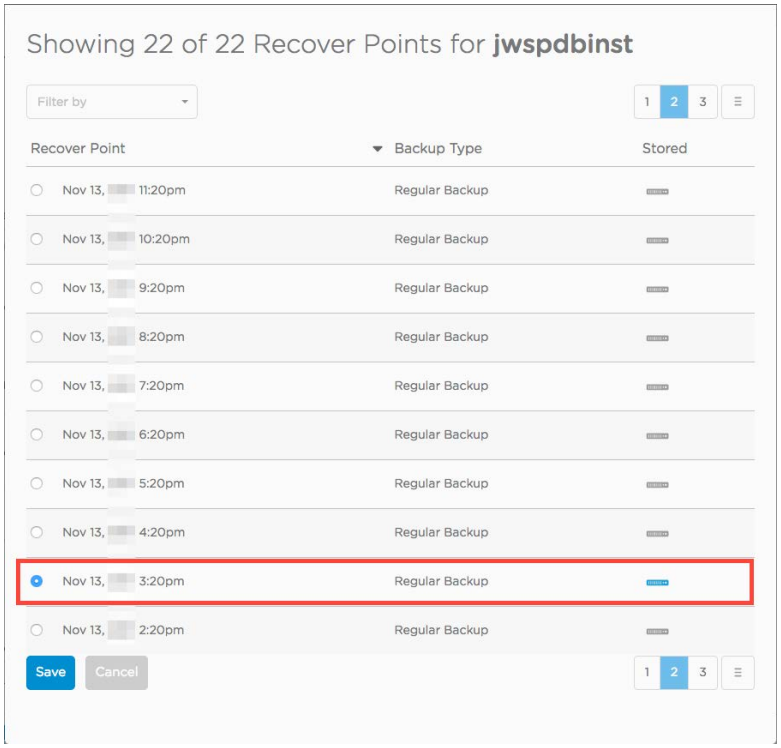

We are choosing yesterday's backup that happened at 3:20pm. Click Save to continue.

Next we need to choose our target server which has Ontrack PowerControls installed.

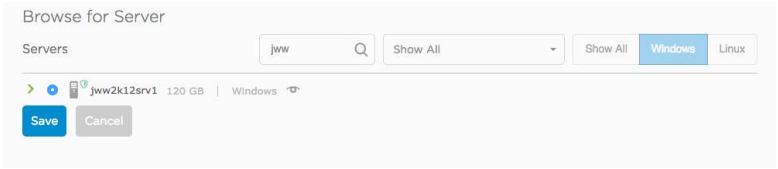

Just like when we searched for the source SharePoint MS SQL server, we can use the search to locate our Windows server where we have Ontrack PowerControls installed. Select the server and click Save to continue.

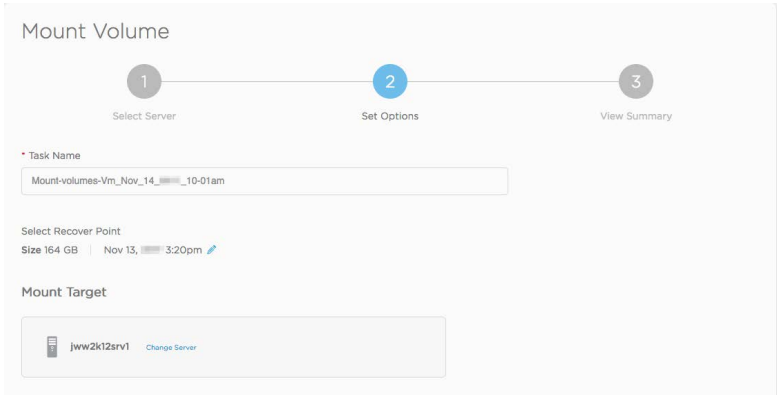

We can see and verify the correct date/time and mount target server have been selected. We need to select which original drive letter we want from the backup to be mounted to the Windows server where we have Ontrack PowerControls installed. We happen to know our database files are on the D: drive. If you don't know, you can select all the drives and then use file explorer to look, or look on the original SharePoint MS SQL server.

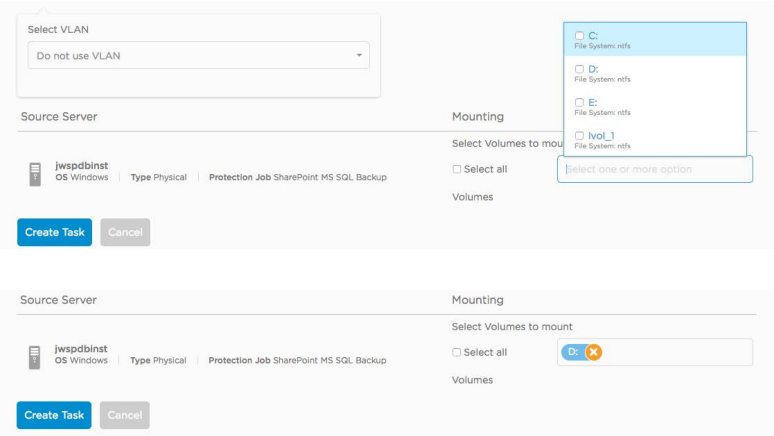

We can now create the task by clicking on the Create Task button.

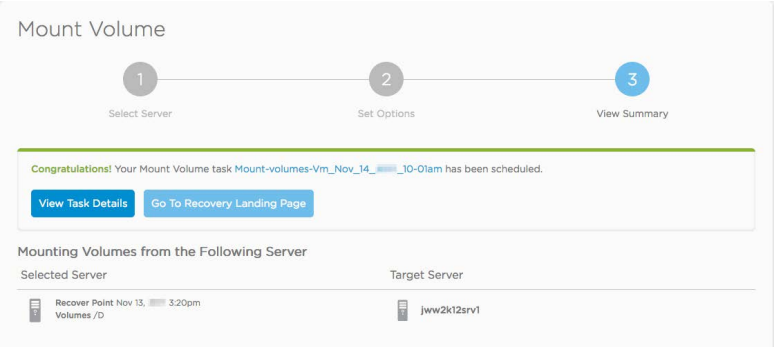

From the task page we can see the Mount Volume task has been scheduled. We can click on the View Task Details to view the details and status of the task.

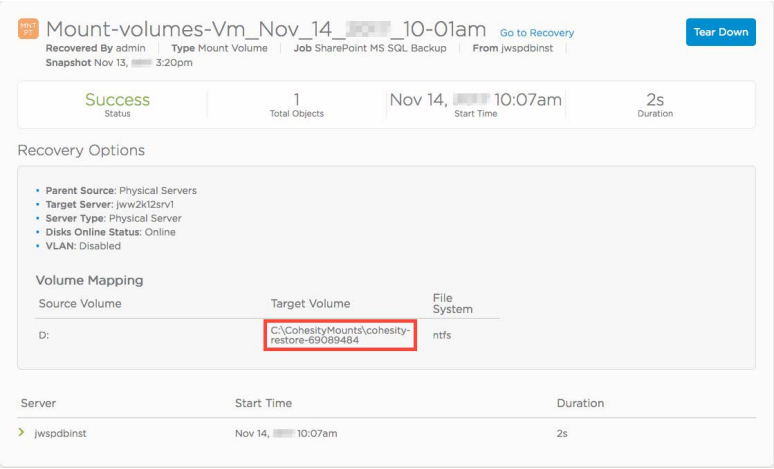

We can see where the D: drive from the source backup was mounted on the target Windows server, we'll use that path with Ontrack PowerControls. We can also see this took 2 seconds.

If we want, we can use File Explorer on verify the mdf and ldf files are present. This is optional and can be skipped.

Let's launch Ontrack PowerControls for SharePoint. By default it will prompt to select the source path to the mdf, ndf and ldf files. We'll add the mdf file first by copying and pasting the mount path from the Cohesity task details page and then drilling down to where the WSS_Content mdf file is. We will do the same thing for the WSS_Content ldf file as well.

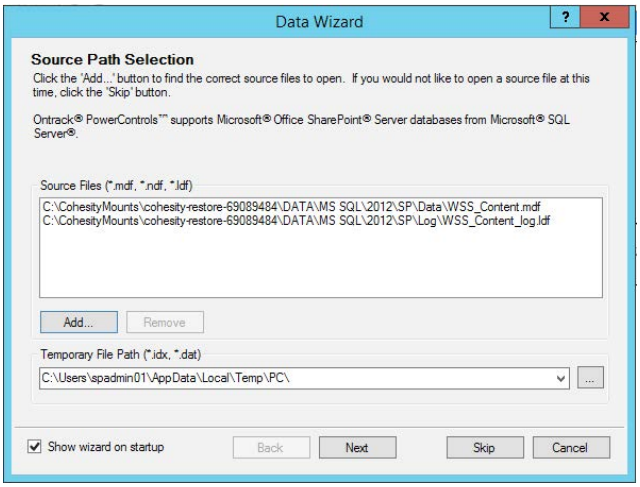

The Temporary File Path can be left as is or changed if desired.

With the source mdf and ldf files selected for our content database, we can click Next.

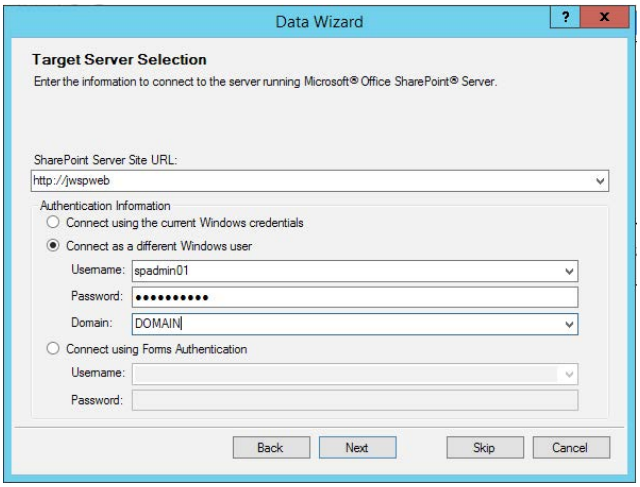

We now need to enter the URL to our SharePoint Server Farm and supply the SharePoint Farm admin account credentials. Click Next to continue.

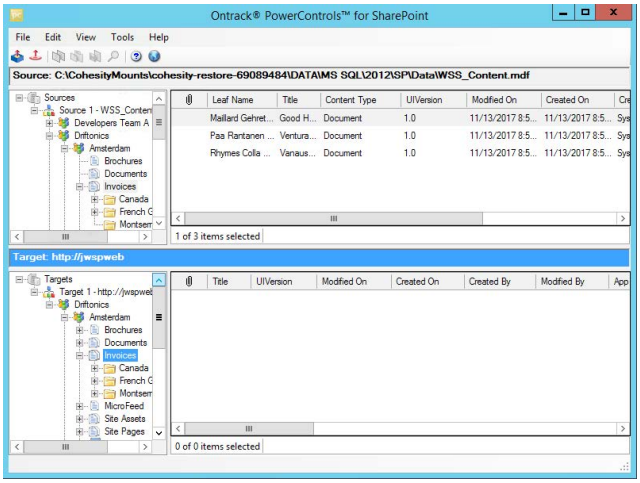

This will now open both the source and target. We can drill down to Driftonics > Amsterdam > Invoices. We can see our Target does not show any documents, but we can see them from the source (our backup). All that is needed to restore these is to select the three files and drag them to the Invoices list on the left hand side of the Target panel. In this example we are just recovering and restoring the three deleted files.

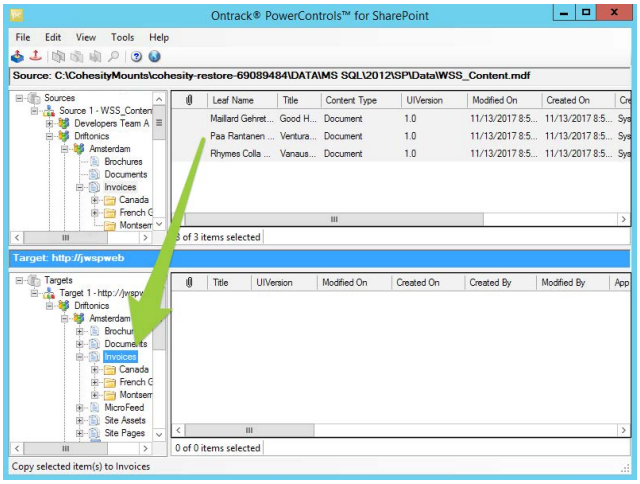

After dragging the files to the Invoices list, the restore will begin. At the completion there will be status of the restore / copy process.

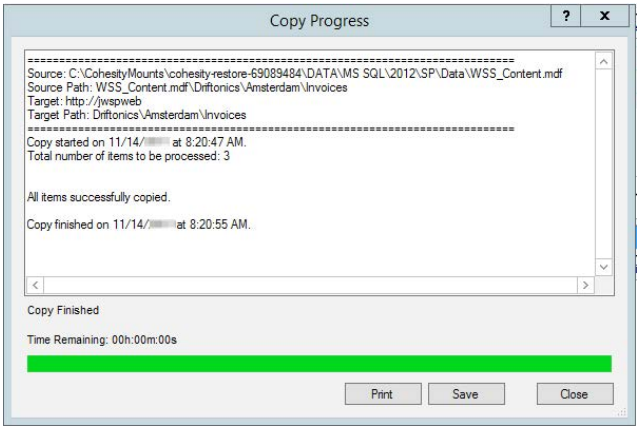

You can Print or Save the output from the copy progress, or simply click Close. We'll now see the files in the Target panel and can verify within SharePoint that all deleted items are restored.

If we refresh the SharePoint page, we'll see the documents are now listed and accessible.

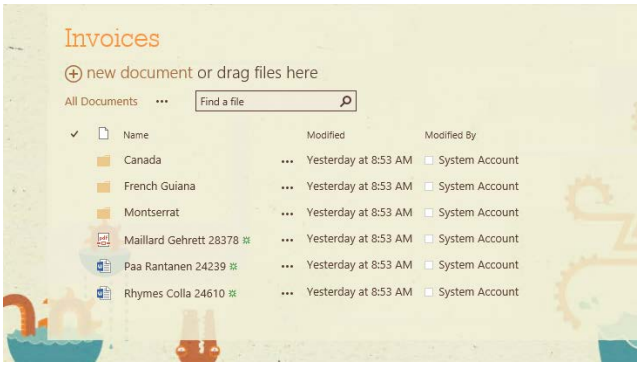

If we are done with granular restores using Ontrack PowerControls, we can exit from the application. The final step is to tear down the restore mount point.

Navigate to the Cohesity cluster and the job status page for the Instant Mount Recovery. The Tear Down button is in the upper right hand side of the Instant Mount Volume job status page. It will ask you to verify you really want to tear down.

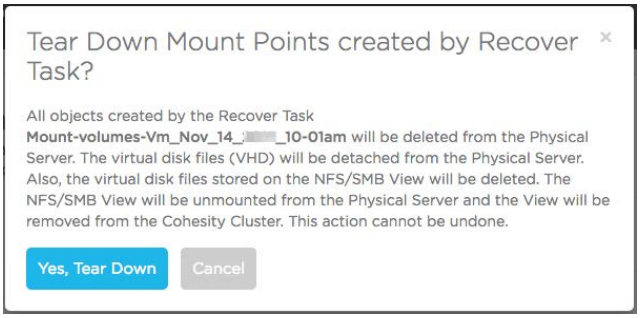

After confirming and clicking on the Yes, Tear Down. At this point the restore mount point will be removed from the Windows server and the cloned snapshot backup deleted on Cohesity. The original backup will remain unchanged.

Granular Recovery / Restore when using FILESTREAM RBS

Ontrack PowerControls supports a few RBS Providers for doing restores. We'll look at recovering a corrupted file stored to SharePoint when Microsoft's FILESTREAM RBS provider is being used.

We will simulate corrupting a file by deleting some of the shredded chunk files written to an iSCSI disk/folder through file explorer. If SharePoint can't locate and reconstruct the original file from RBS, it can't return the original file and in this case/example the pdf viewer won't be able to render the file.

Under Documents we have a few PDF's, we'll simulate corrupting the one called rbs_Reprint_sample_2.

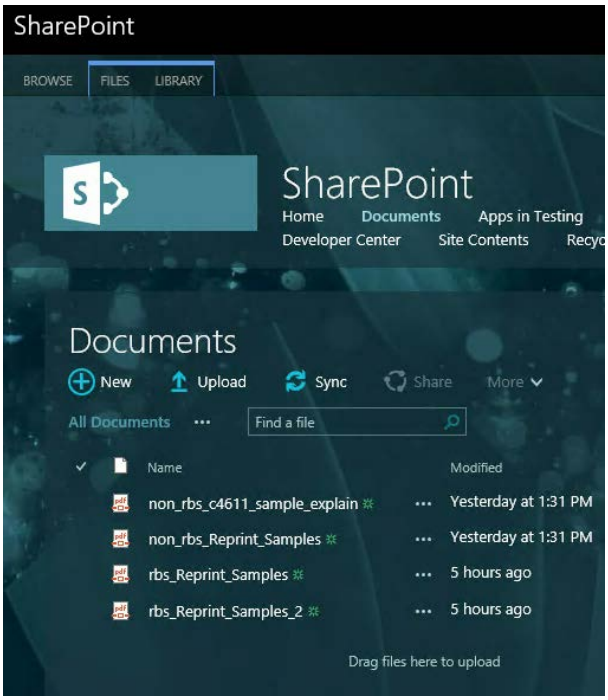

We can see the chunk files in file explorer, the ones from 7:12AM are related to rbs_Reprint_sample_2.

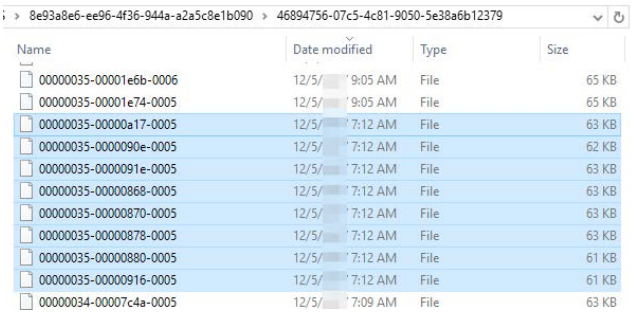

We'll delete a handful to simulate some file/file system corruption. It's not suggested this be done in a production environment. Previously we had 8 chunk files, after deleting we have 4.

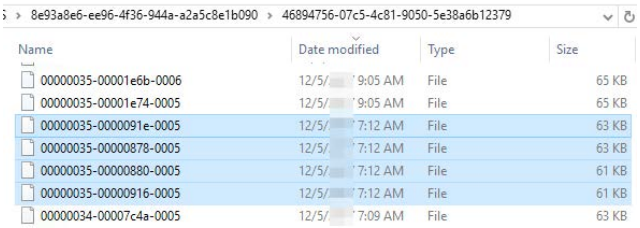

Let's try and view the rbs_Reprint_sample_2 pdf. Given enough time, the web browser will give up trying to open/download the file from SharePoint. In this case the web browser loaded the pdf viewer but was unable to render the pdf.

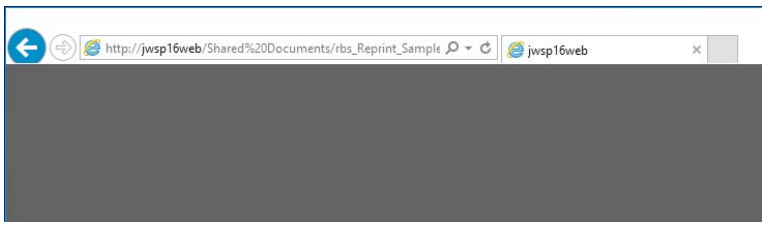

The restore process is nearly identical as that of doing a granular recovery without RBS, except we must present the backup (point in time snapshot) of the RBS file system to Ontrack PowerControls along with the mdf and ldf files.

In our case our FILESTREAM RBS data is stored under the H: drive. G: is where our SharePoint content database mdf and ldf files are.

Following the same steps via the Cohesity UI we will perform an Instant Volume Mount.

We need to include the H: drive along with the G: drive.

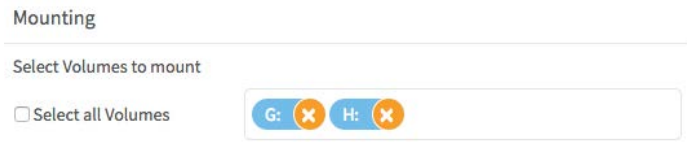

Once the Instant Volume Mount is complete you can find where each drive from the previous backup is mounted.

Launch Ontrack PowerControls

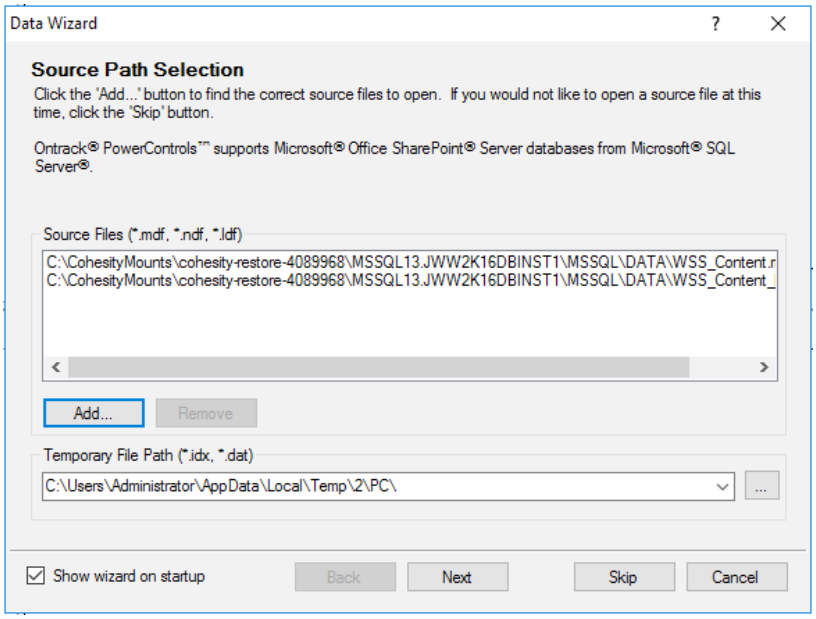

As before, we'll use the Add... button and select the appropriate .mdf and .ldf file for the content database. Then click Next.

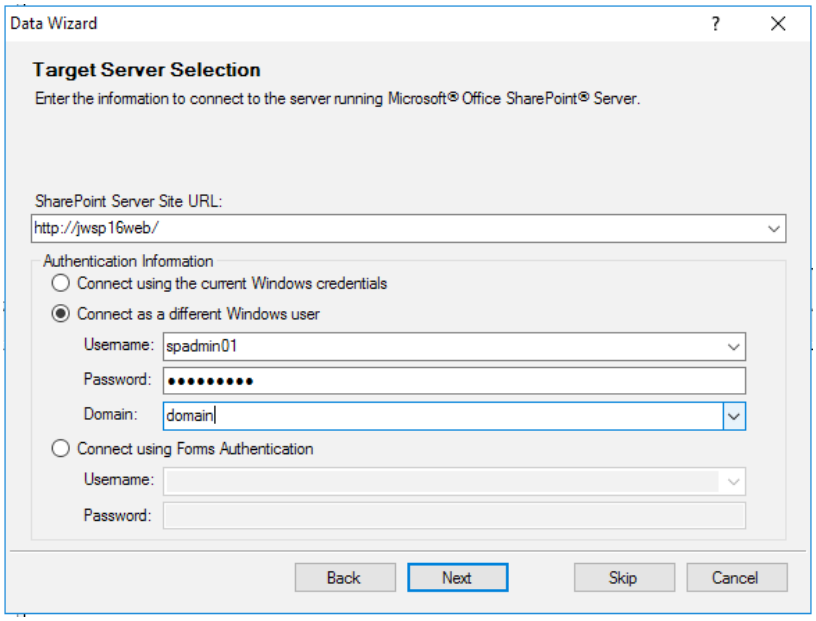

Type in the SharePoint Server Site URL and login details. Click Next.

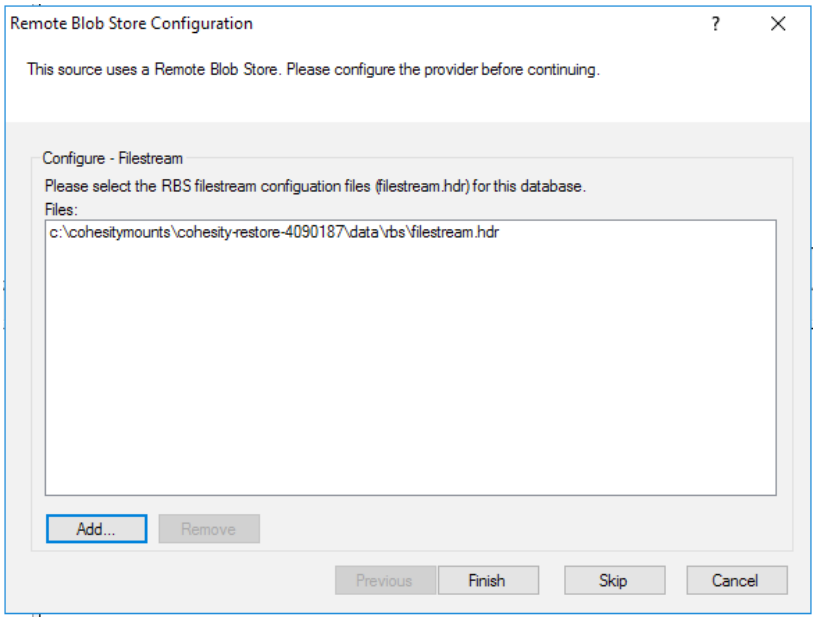

This screen is new and specific to RBS. Ontrack PowerControls will know if the RBS FILESTREAM provider is configured after opening the SharePoint Content database.

Use the Add... button to locate the filesteam.hdr file and select it. Click Finish.

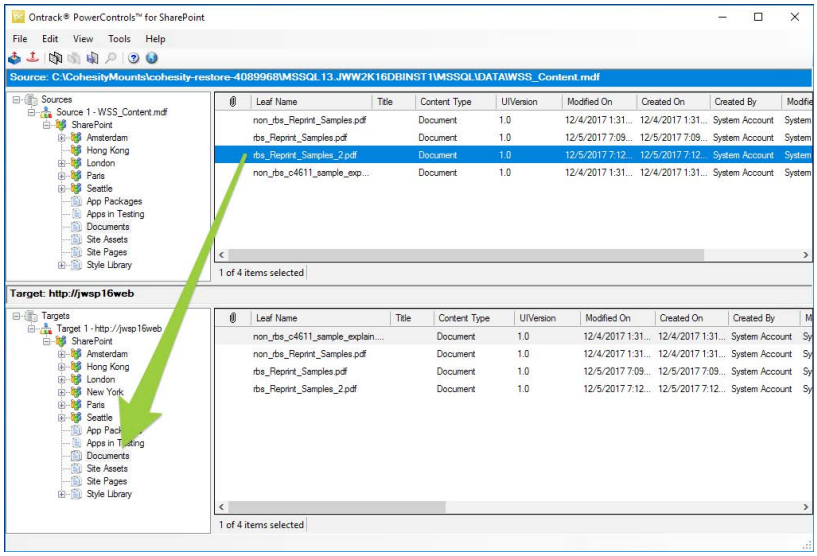

As with the non-RBS restore, we simply drag and drop the item to be restored. In this case the original document is still there, just corrupted. We will simply overwrite the corrupted pdf.

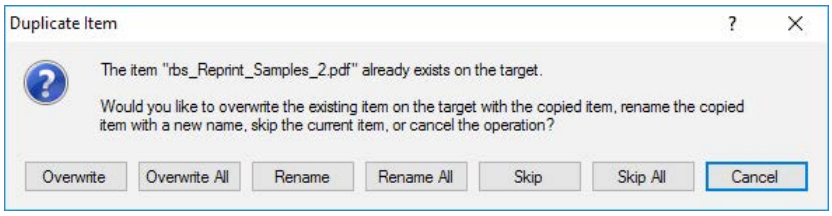

We'll click Overwrite as we know we want to replace the current pdf.

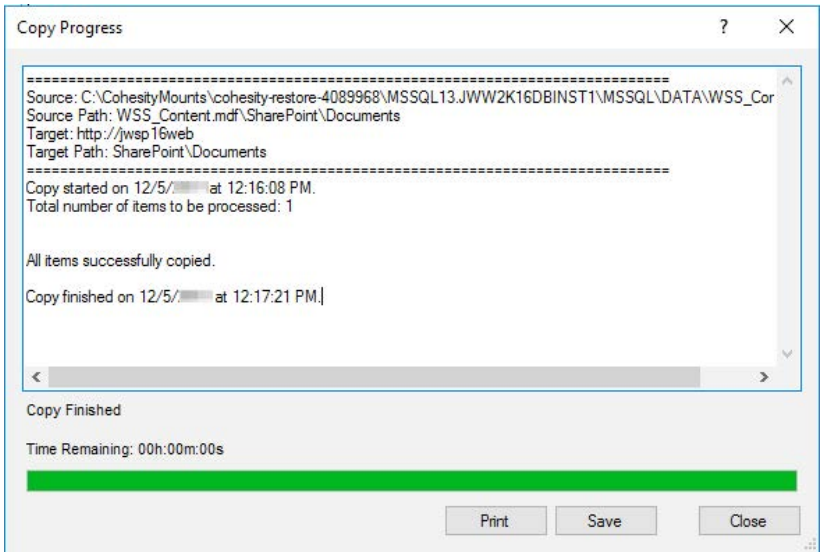

If we try and open the document in SharePoint, it succeeds.

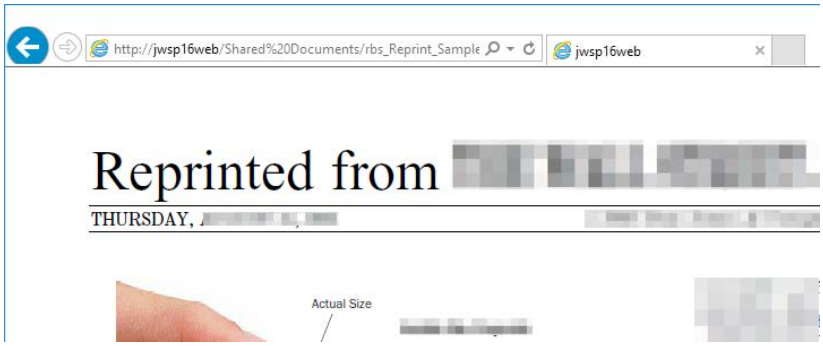

Microsoft SQL Server Clustering

Clustered MS SQL instances can be backed up with Cohesity starting with Cohesity version 5.0.

Backing up a Clustered MS SQL Instance (FCI) is very similar to backing up a stand-alone SQL Server. Review the Cohesity User Guide and/or Cohesity MS SQL Clustering Best Practice Guide for more details.

If the FILESTREAM RBS provider is utilized, you'll want to be sure it's being backed up along with the SQL FCI.

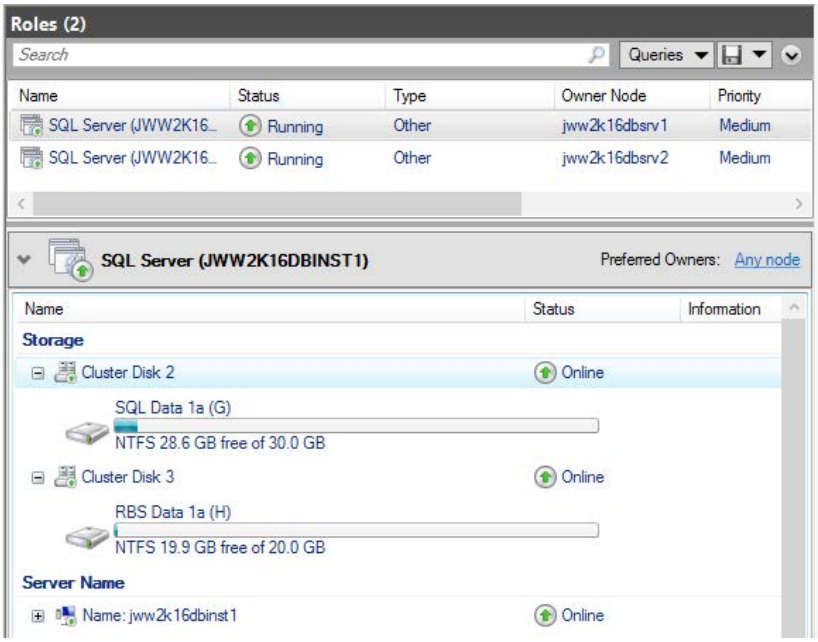

In our example, our fail-over MS SQL drive is G: and our RBS drive is H:. We can see both are being backed up if we review the job run details.

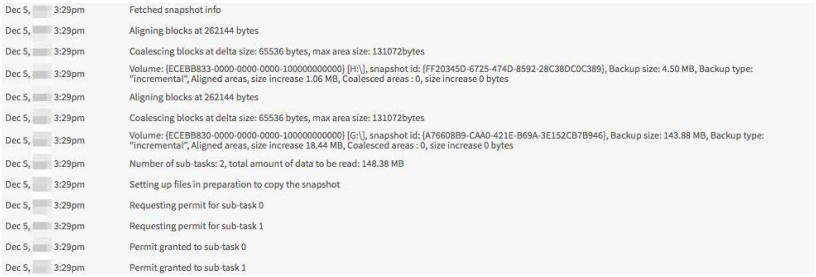

In the event the Cohesity cluster is not backing up the RBS fail-over drive because it does not think it's being used by the MS SQL instance, a simple work-around is to create an empty database within the instance that is located on the RBS fail-over drive. This will ensure that the RBS fail-over drive is backed up along with the MS SQL instance drive.

With the MS SQL data drive and RBS FILESTREAM drive being backed up, we can perform granular recovery of objects within SharePoint, whether they are stored within the MS SQL database, or stored external of the MS SQL database via the FILESTREAM RBS provider. Follow the same recovery steps for non-Clustered SQL Instances to recover SharePoint that's using SQL Clustering.

Version History

Version	Date	Document Version History
Version 1.0	May 2018	Original Document

About the Author

Justin Willoughby is 20-year IT veteran, currently working for Cohesity as a Solution Engineer. In this role, Justin architects, builds, tests, and validates business-critical applications, databases, and virtualization solutions with Cohesity's DataProtect.

About Cohesity

Cohesity delivers the industry's first solution for secondary storage consolidation. Cohesity enables companies of all sizes to bring order to their data chaos by converging storage workloads, including file services, data protection, Test/Dev, and in-place analytics, onto an infinitely scalable, intelligent data platform.

With Cohesity, customers can manage and protect data seamlessly, use it efficiently, and learn from it instantly. Cohesity is headquartered in the heart of Silicon Valley, California with a global presence across the Americas, EMEA, and APAC.

Trademarks

References

[1] Ontrack PowerControls User Guide [pdf], April 2017, Page 1 / Printed Page 8

152018