

Back up to the Cloud using Equinix with Hitachi Data Protection Suite: VSP E1090 to Amazon Web Services

With Hitachi Cloud Connect for Equinix

Reference Architecture

August 2022

Contents

Notices and Disclaimer 3

Executive Summary..... 4

About This Guide..... 5

 Introduction 5

 Intended Audience 5

 Document Revisions..... 5

 References 5

 Comments 5

Solution Overview 6

 Benefits..... 6

 Key Components 6

Engineering Validation..... 7

 Validation Method 7

 High Level Diagram 8

 Hardware and Software..... 9

 Test Scenarios..... 10

Guidelines and Recommendations..... 12

Validation Results..... 13

 Test 1: Prepare the Environment..... 13

 Test 2: Create Storage Libraries and Backup Plans 18

 Test 3: Create a Full Backup of the SQL Database 22

 Test 4: Create an Incremental Backup..... 24

 Test 5: Perform an In-Place Restore..... 25

 Test 6: Perform an Out-of-Place Restore to Cloud 29

 Test 7: Manual Failover of CommServe 33

 Test 8: Automatic Failover of CommServe..... 35

 Test 9: Unplanned Outage with Data Return 36

Notices and Disclaimer

© 2022 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara (collectively, "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video, and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara at https://support.HitachiVantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara.

By using this software, you agree that you are responsible for:

- 1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
- 2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls: The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

EXPORT CONTROLS: Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, Flash Copy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, BCPII™ and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or web site are properties of their respective owners.

Executive Summary

Hitachi Data Protection Suite (HDPS) provides a seamless and efficient backup and restore of a production environment from most mainstream operating systems, databases, and applications. This reference architecture documents how to set up a backup solution with HDPS to back up an on-premises Microsoft SQL Server to cloud storage in Amazon Simple Storage Service (S3).

The environment used for the validation included a Virtual Storage Platform (VSP) E1090 used as the primary storage system that was located in an Equinix colocation data center in California. The Equinix colocation was chosen because it offered high-speed and low latency connections to the major hyperscalers, such as Amazon Web Services (AWS). In fact, Hitachi Vantara has collaborated with Equinix to create a new near-cloud hybrid solution called **Hitachi Cloud Connect for Equinix**.

This offering allows clients to locate Hitachi VSP enterprise-class storage at Equinix International Business Exchange™ (IBX) data centers worldwide and includes the option for customers to procure this solution through one agreement and invoice, greatly simplifying and accelerating their time to market. By using Equinix IBX data centers and Equinix Fabric™ to interconnect sources of data to applications, VSP storage systems enable organizations to locate their data next to clouds while still maintaining control by enabling applications such as data protection and back-up for hybrid- and multi-cloud data availability.

If you want to discuss options for hosting a disaster recovery solution at Equinix, contact your Hitachi Vantara sales team. You can also visit the Hitachi Cloud Connect for Equinix webpage for more information: <https://hitachivantara.com/en-us/products/storage/flash-storage/cloud-connect-for-equinix.html>.

The backup location used in this solution was a cloud-based object store in the form of an Amazon S3 bucket. This service offers industry-leading scalability, data availability, security, and performance. The cost of owning and operating a second on-premises data center for disaster recovery can be significant, so leasing a small footprint in the cloud is a cost-attractive alternative.

About This Guide

Introduction

This reference architecture documents how to set up a backup solution with Hitachi Data Protection Suite to back up an on-premises Microsoft SQL Server to cloud storage in Amazon S3. It also documents test procedures for validating the resiliency of the solution, which you can leverage for your own proof-of-concept before deploying the solution.

The environment used for the validation included a VSP E1090 used as the primary storage system located in an Equinix colocation data center in California and an Amazon S3 cloud object store used as the backup target.



Note: The results obtained in these tests were specific to the environment, workload application, and parameters used in this configuration. We recommend conducting a proof-of-concept to obtain acceptable results in a non-production, isolated test environment matching your production environment before implementing this solution.

Intended Audience

This document is intended for Hitachi Vantara staff and IT professionals of Hitachi Vantara customers and partners who are responsible for planning and deploying such as solution.

Document Revisions

Revision Number	Date	Author	Details
1.0	March 2022	Hitachi Vantara LLC	Initial Release
1.1	August 2022	Hitachi Vantara LLC	Update to Figure 2

References

- <https://documentation.commvault.com/hitachivantara/>
- <https://docs.aws.amazon.com/>

Comments

Send any comments on this document to GPSE-Docs-Feedback@hitachivantara.com. Include the document title, including the revision level, and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara Corporation.

Thank You.

Solution Overview

Hitachi Data Protection Suite is a software solution that provides backup, recovery and management of enterprise and application data. HDPS offers the reliability required by the world's largest organizations, while featuring the simplicity, cost-effectiveness and modern capabilities that are needed to remain agile and competitive.

For any backup implementation, storage space is a major consideration. Typically, the storage comes from a block storage device or a file system located in a secondary data center. The primary data center and secondary data center must be separated geographically to reduce the chance of a single disaster affecting both sites.

To avoid the real estate, economic, and technical overhead of a second data center and minimize hardware footprint, cloud-based object storage, such as Amazon S3, can be used instead. This architecture enables you to back up on-premises applications to Amazon S3 object storage and includes the option to maintain a secondary copy of the backup data on a second object store to make the solution even more robust.

Benefits

The following describes the benefits of a backup and restore solution with Hitachi Data Protection Suite:

- The solution allows business to minimize data loss and resume operations quickly when a disaster brings down the production application.
- Backing up to the cloud is cost attractive because it avoids the infrastructure cost and management overhead of operating a secondary data center for backups.
- HDPS further reduces cost with software-based deduplication that significantly reduces the amount of data that must be stored in the cloud.
- Cloud-based infrastructure, such as those provided by AWS, is highly resilient so backup data is highly available.

Key Components

The major solution components are provided in the following list. Specifications are provided in the [Hardware and Software](#) section.

- Storage System: A VSP E1090 hosted the production SQL database instance at the primary site.
- Hitachi Data Protection Suite: Backup and restore software with the following components:
 - CommServe: Command and control center of the software. Responsible for handling all activity between agents, such as initiating backup and recovery jobs.
 - MediaAgent: Oversees the transfer of data between backup targets and storage libraries. Each MediaAgent communicates locally to one or more storage libraries.
- Networking: Equipment at the primary site in the Equinix colocation data center was connected to AWS using the following:
 - Equinix Fabric: Uplink to AWS cloud.
 - Cisco Nexus 9000 Switch: Transferred data between the primary site and AWS.
- AWS Cloud: The following AWS resources in region us-west-1 were used:
 - Amazon EC2: Used to deploy MediaAgents and standby CommServe in the cloud. EC2s are virtual machines that provide scalable computing capacity, which makes them attractive for temporary business requirements.
 - Amazon S3: Provided storage for backups. Amazon S3 provides 99.999999999% durability for objects stored in buckets, making it a reliable choice for a data backup solution.
 - AWS Direct Connect (10 Gbps): Transferred data between AWS and production site.

Reference Architecture

Engineering Validation

This section describes the method, test environment, hardware and software, and test scenarios used in the validation.

Validation Method

To validate the solution, a 2-node Microsoft SQL Server cluster was used. New records were written to the SQL database prior for each backup operation test and then verified after the restore/failback to ensure data consistency.

The SQL cluster was located on a Microsoft Windows Failover Cluster running on ESXi virtual machines. Two volumes were mapped from the VSP E1090 storage system at the Equinix SV10 site and passed through to the virtual machines as raw device mappings. These volumes hosted the database files and logs. A third volume, configured as a VMFS datastore, hosted the virtual machine operating system drive.

Figure 1 illustrates the SQL setup.

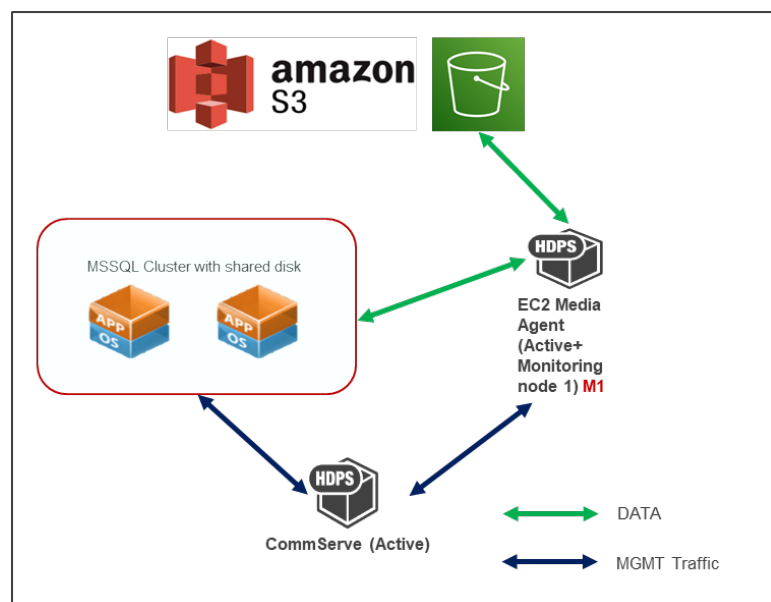


Figure 1. SQL Setup

High Level Diagram

Figure 2 shows the test environment used to run the validation.

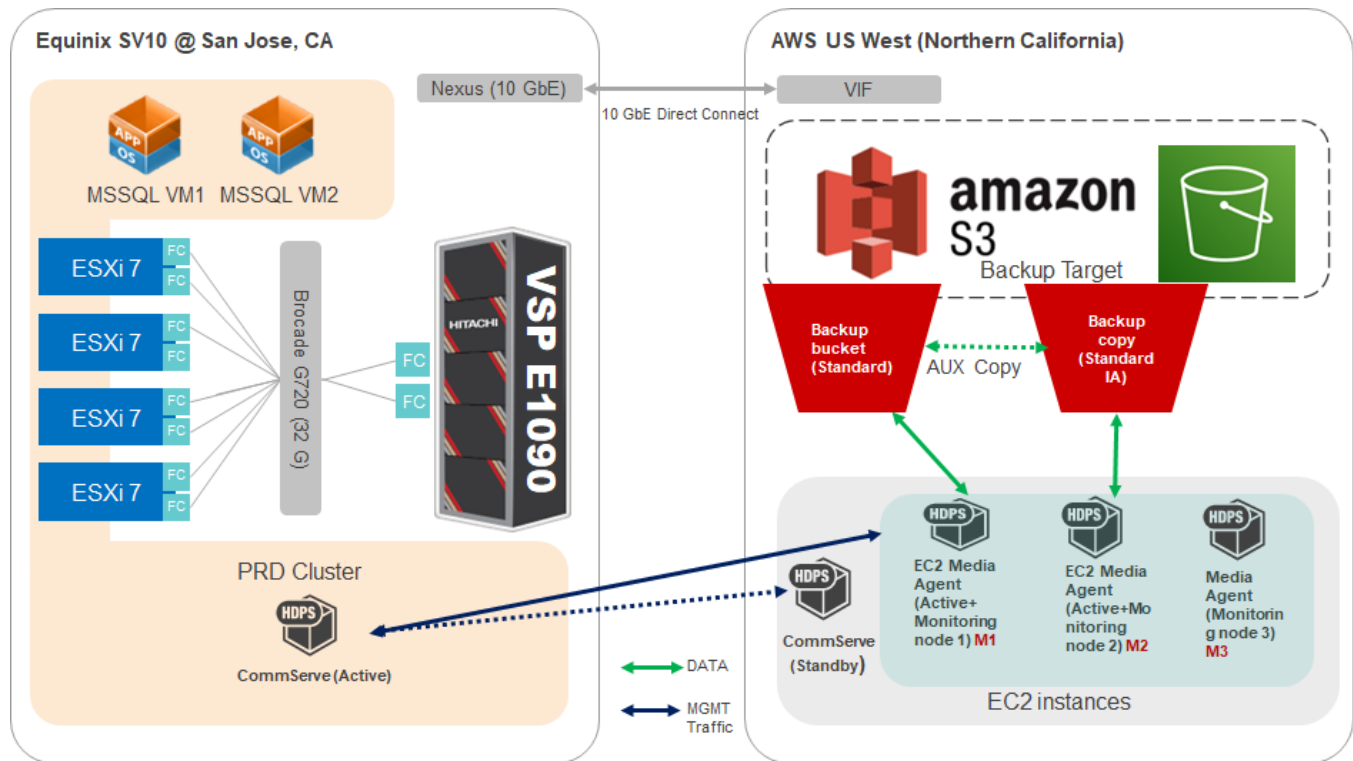


Figure 2. Test Environment

Hardware and Software

Table 1 provides the hardware specifications used in this validation.

Item	Description	Version	Function
Hitachi VSP E1090	1,024 GB cache (2) 32-core MPUs (3) RAID6 6D-2P parity groups (2) 32 Gbps FC ports	SVOS RF 9.8 93-06-21-80/00	Primary storage system
Hitachi Advanced Server DS220	(2) 18-core Intel Xeon Gold 6140 @ 2.3 GHz 128 GB cache (1) Emulex LPe32002 HBA (1) Intel Ethernet Network Adapter XXV710	BMC 4.70.06 BIOS S5BH3B22.H00	4-node primary vSphere cluster
CommServe (Active Node)	Virtual Machine: (6) Virtual CPUs 16 GB RAM 350 GB virtual disk	VM Version 19 Guest OS: Windows Server 2019 (64-bit)	CommServe node at Equinix data center. Managed all activity between MediaAgents and storage libraries.
Brocade 6510	16 Gbps capable Fiber Channel switch	FOS 8.2.1c	Provided FC connectivity between VSP E1090 and vSphere cluster.
Cisco Nexus C93180YC-FX	(48) 1/10/25-Gbps fiber ports (6) 40/100-Gbps QSFP28 ports	NXOS 9.3(4)	Provided network connectivity between equipment at Equinix and AWS cloud.
CommServe (Standby Node)	Instance type: x2.large	AMI: Windows_Server-2019-English-Full-Base-2022.01.21	CommServe node on AWS cloud, served as standby.
MediaAgents on AWS	Instance type: x2.large Each with (1) 300 GB and (1) 2 TB Elastic Block Storage (EBS) volumes	AMI: Windows_Server-2019-English-Full-Base-2022.01.21	Transferred data between backup clients and storage libraries. Also acted as Monitoring Agents for CommServe LiveSync.
Amazon S3	(1) Standard S3 bucket (1) Standard-IA S3 bucket	N/A	Standard bucket: storage for primary backup. Standard-IA bucket: storage for auxiliary backup copy.

Table 1. *Hardware Components*

Table 2 provides the software specifications used in this validation.

Item	Version	Function
Hitachi Data Protection Suite	11.24.29	Backup and restore software solution
VMware vSphere	7.0 U2 (17867351)	Hypervisor operating system at Equinix data center
VMware vCenter Server Appliance	7.0 U3 (18700403)	Management interface for vSphere cluster at Equinix data center
Microsoft Windows Server 2019 Datacenter	Microsoft Windows Server 2019 Datacenter	Guest operating system for all virtual machines at Equinix and AWS, including SQL database instances, CommServe nodes, and MediaAgents
Microsoft SQL Server Enterprise	15.0.2000.5	Database application used to validate data consistency

Table 2. Software Components

Test Scenarios

Table 3 lists the test scenarios performed in the validation.

Test	Description	Success Criteria
1	Prepare the environment: 1. Deploy production SQL database instance: a. Provision (2) 2 TB volumes from the VSP E1090 to use as datastores. b. Create two Windows Server 2019 virtual machines on the datastores. c. Install Microsoft SQL Server 2019 on the virtual machines. d. Provision (2) 2 TB volumes from the VSP E1090 to use as shared disks. e. Map the (2) 2 TB disks as raw device mappings to both virtual machines. f. Configure Windows Server Failover Cluster on the virtual machines. g. Configure SQL Server failover cluster on the virtual machines. h. Create a new database on the shared disks. 2. Deploy a Microsoft SQL Server 2019 as an EC2 virtual machine. 3. Deploy Hitachi Data Protection Suite: a. Deploy active CommServe node on the datastore created in step 1a. b. Deploy standby CommServe node and MediaAgent nodes as EC2 virtual machines. c. Attach (1) 2TB EBS volume to MediaAgent M1 and M2 for Deduplication Database. d. Register all MediaAgents on the active CommServe node. e. Create (2) S3 buckets.	Environment is set up per specifications.
2	Create storage libraries, backup plans, and register SQL database instances.	Objects created successfully.
3	Create a full backup of the production SQL database instance and store the data in S3.	Backup job is successful.
4	Create an incremental backup of the production SQL database instance and store the data in S3.	Only changed data is backed up.

Test	Description	Success Criteria
5	Restore backup data using an in-place restore to the production SQL instance at the primary site.	Restore job is successful. Data is consistent.
6	Restore backup data using Out of Place Restore to the standby SQL instance in AWS.	Restore job is successful. Data is consistent.
7	Manual failover of the active CommServe node: <ol style="list-style-type: none"> 1. Trigger a failover from the active CommServe node to the standby CommServe node. 2. Verify that all services shut down on the active node and started up on the standby node. 3. Verify that all backup plans and storage libraries are available on the standby node. 4. Trigger a failback from the standby node to the active node. 	Manual failover and failback are successful.
8	Automatic failover of the active CommServe node: <ol style="list-style-type: none"> 1. Shut down the active CommServe node. 2. Verify that all services start automatically on the standby node. 3. Verify that all backup plans and storage libraries are available on the standby node. 4. Trigger a failback from the standby node to the active node. 	Automatic failover is successful.
9	Unplanned outage with data return: <ol style="list-style-type: none"> 1. Abruptly power off the SQL virtual machines and active CommServe node at the primary site. 2. Wait about 30 minutes for CommServe to automatically fail over to the standby node. 3. Perform an Out of Place restore job to bring the production database online on the standby SQL virtual machine. 4. Make changes to the SQL database and then perform a backup job. 5. Power off the standby SQL virtual machine. 6. Power on the SQL virtual machines and the CommServe node at the primary site. 7. Perform an Out of Place restore job to bring the latest updates from the standby SQL instance to the production SQL instance. 	<p>Able to restore production database to SQL on AWS.</p> <p>Able to bring back data created in AWS to on-premises SQL.</p>

Table 3. *Test Scenarios*

Guidelines and Recommendations

This section describes the lessons learned from this validation, along with guidelines and recommendations.

- This validation focused on a crash-consistent backup solution. Because the client application, a Microsoft SQL Server, was in a cluster configuration with shared disks, application-consistent backup was not possible. For your reference, the difference between the two options is:
 - Crash-consistent: The backup client is not aware of the backup job. Any pending I/O operations or data in memory is not captured by the backup job. Extra work might be required to restore an application backed up using this method.
 - Application-consistent: The backup client is aware of the backup job. The client is quiesced, the memory flushed to disk, and then the backup job is initiated. No additional work is required to restore the backed-up application.
- Host MediaAgents on EC2 virtual machines instead of on-premises virtual machines for better fault tolerance. The AWS infrastructure will most likely have higher availability than an on-premises data center.
- Ensure the various components involved, including the backup source, CommServe, MediaAgents, and S3 buckets, can resolve hostnames through a common DNS server.
- Ensure network ports required by HDPS are open on the firewall. See the following for a list of port requirements:
https://documentation.commvault.com/hitachivantara/v11/essential/7102_port_requirements_for_commvault.html.
- Even if your MediaAgents and standby CommServe are deployed on EC2 virtual machines in the same Virtual Private Cloud, you must explicitly allow traffic between the virtual machines by editing the security group.
- Do not include an underscore character in the name of the CommServe node. It prevents you from accessing the web-based Command Center interface on the standby CommServe after a LiveSync failover. This issue is due to a limitation in the Apache Tomcat application that runs on the CommServe node.
- After completing an Out of Place Restore job, remember to stop any scheduled database replication tasks. Otherwise, they will corrupt the newly restored database.

Validation Results

This section contains specific steps and screenshots for each test scenario.

Test 1: Prepare the Environment

This test case describes the configuration of the various components used in the validation.



Note: It is important to ensure the various components involved, including the backup source, CommServe, MediaAgents, and S3 buckets, can resolve hostnames through a common DNS server, and necessary network ports are open on the firewall.

- Four volumes were provisioned from the VSP E1090 storage system at the Equinix SV10 site to the primary vSphere cluster as follows: Descriptions of the volumes are given below.

Pool Volumes Virtual Volumes TI Root Volumes							
Create LDEVs Add LUN Paths Expand V-VOLs More Actions							
Filter ON OFF Select All Pages Column Settings							
	LDEV ID	LDEV Name	Status	Capacity			
				Total	Reserved	Used	Used (%)
<input type="checkbox"/>	00:00:00	Juno-A1	Normal	2150.40 GB	0.00 GB	185.47 GB	8
<input type="checkbox"/>	00:00:01	Juno-A1	Normal	2150.40 GB	0.00 GB	41.75 GB	1
<input type="checkbox"/>	00:00:10	Juno-A1-SQL-0	Normal	2048.00 GB	0.00 GB	198.06 GB	9
<input type="checkbox"/>	00:00:11	Juno-A1-SQL-1	Normal	2048.00 GB	0.00 GB	390.96 GB	19
<input type="checkbox"/>	00:00:12	Juno-A1-SQL-2	Normal	1.00 GB	0.00 GB	0.20 GB	20

The following provides a description of the volumes:

- Volume 00:00:00 and 00:00:01: Used for datastores to host the SQL virtual machines.
 - Volume 00:00:10 and 00:00:11: Used as shared cluster disks for SQL.
 - Volume: 00:00:12: Used for Windows Failover Cluster Quorum Disk.
- S3 buckets were used for this validation and public access was blocked for the S3 buckets as shown in the following screenshot of the AWS Management Console:

The first bucket, a1-e1090-backup, was created for the storage library.

a1-e1090-backup Info		
Objects	Properties	Permissions Metrics Management Access Points
Bucket overview		
AWS Region US West (N. California) us-west-1	Amazon Resource Name (ARN) arn:aws:s3::a1-e1090-backup	Creation date January 24, 2022, 11:01:41 (UTC+05:30)

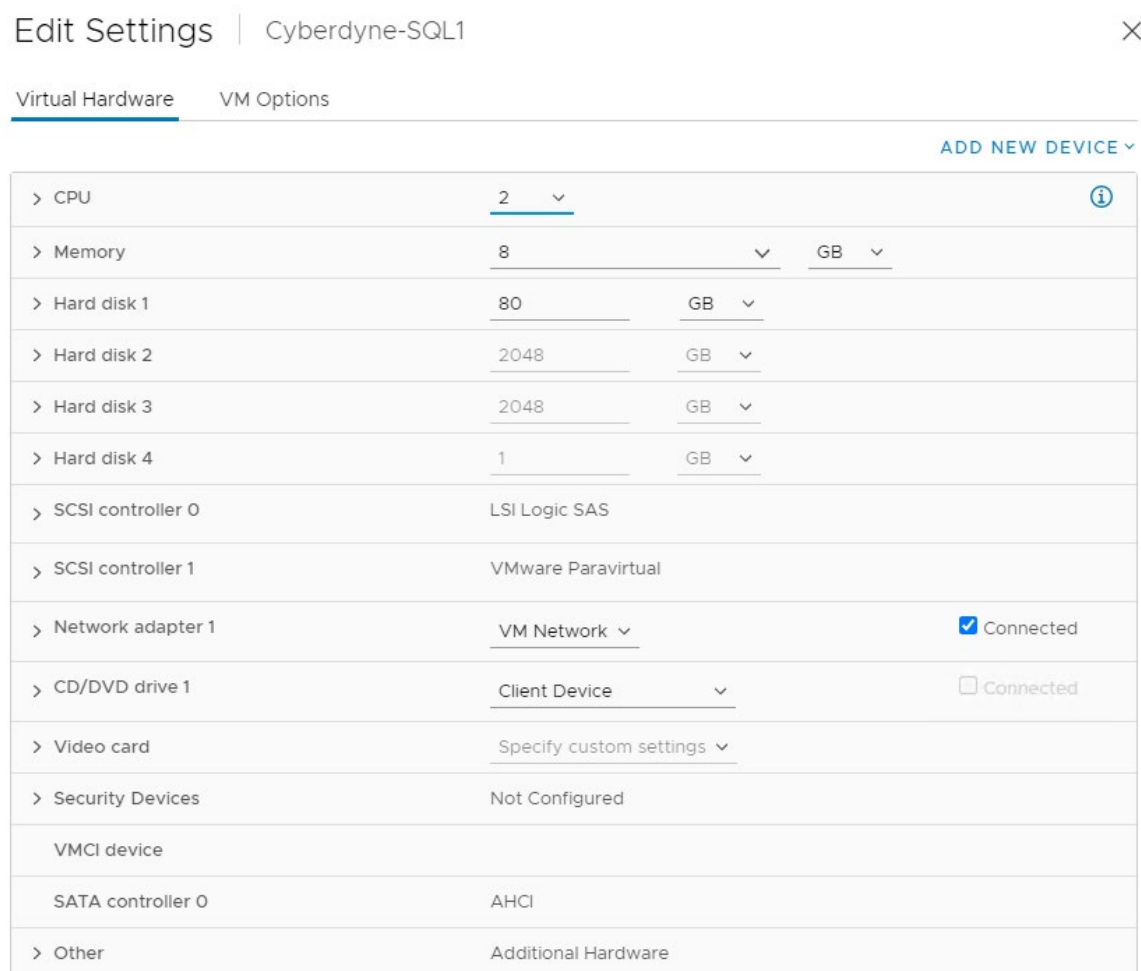
The second bucket, a1-e1090-backup2, was created for a secondary backup location.

a1-e1090-backup2 Info		
Objects	Properties	Permissions Metrics Management Access Points
Bucket overview		
AWS Region US West (N. California) us-west-1	Amazon Resource Name (ARN) arn:aws:s3::a1-e1090-backup2	Creation date February 14, 2022, 12:11:10 (UTC+05:30)

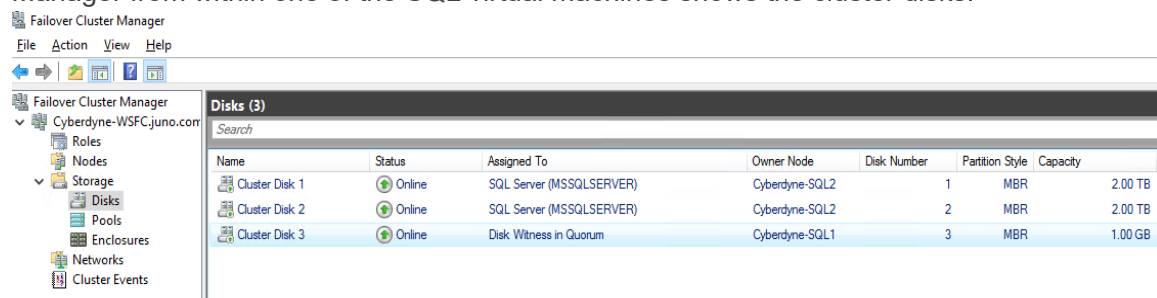


Note: HDPS recommends having a second copy of the backup data. Therefore, we created a second S3 bucket, with a cheaper storage class. The primary copy was kept in a1-e1090-backup which used Standard S3 storage class, while a1-e1090-backup2 used Standard-IA S3 storage class.

3. The SQL virtual machine was configured with four hard disks: one 80 GB virtual disk, two 2,048 GB RDMs, and one 1 GB RDM from the VSP E1090 storage system. The second SQL virtual machine was configured identically, but instead of three RDMs, it used “Existing Hard Disks” of the first SQL virtual machine. The settings as shown in the following screenshot:



4. This is the screenshot of Windows Failover Cluster Manager from within one of the SQL virtual machines. The Windows Failover Cluster 'Cyberdyne-WFSC' owns two 2 TB cluster disks and the two cluster disks have been assigned to the SQL Server role. The following Windows Failover Cluster Manager from within one of the SQL virtual machines shows the cluster disks:

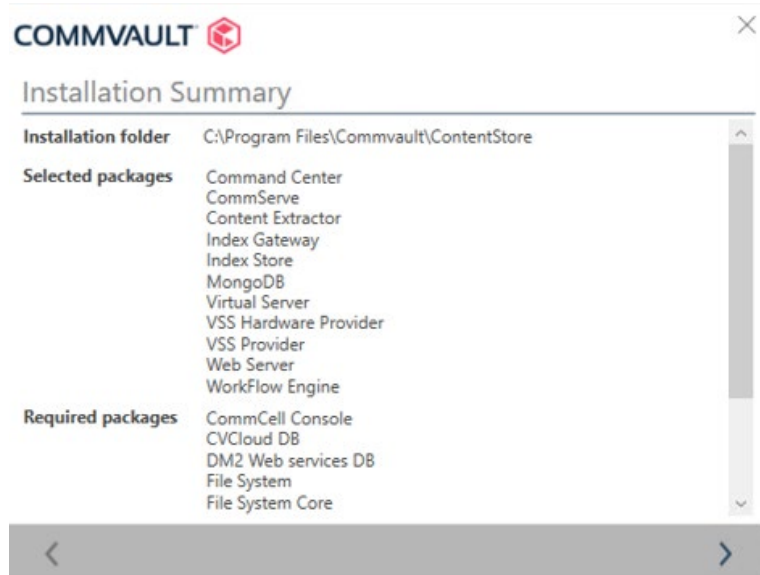


5. An SQL query was used to locate the database files registered to the SQL instance. Entries 18 and 19, fenced in red, make up our test database. The primary database file, PRD-A1-Test.mdf, is located on the first of the two cluster disks. The database log file, PRD-A1-Test.ldf, is located on the second of the two cluster disks.

```
USE master;
SELECT
    name 'Logical Name',
    physical_name 'File Location'
FROM sys.master_files
```

	Logical Name	File Location
1	master	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\master.mdf
2	mastlog	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\mastlog.ldf
3	tempdev	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\tempdb.mdf
4	templog	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\templog.ldf
5	temp2	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\tempdb_mssql_2....
6	modeldev	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\model.mdf
7	modellog	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\modellog.ldf
8	MSDBData	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\MSDBData.mdf
9	MSDBLog	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\MSDBLog.ldf
10	DWDiagnostics	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\DWDiagnostics....
11	DWDiagnostics_log	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\DWDiagnostics_...
12	DWConfiguration	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\DWConfiguration...
13	DWConfiguration_log	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\DWConfiguration...
14	DWQueue	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\DWQueue.mdf
15	DWQueue_log	S:\MSSQL15.MSSQLSERVER\MSSQL\DATA\DWQueue_log.ldf
16	CYBERDYNE	S:\CYBERDYNE-DB\CYBERDYNE.mdf
17	CYBERDYNE_log	T:\CYBERDYNE-log\CYBERDYNE_log.ldf
18	PRD-A1-Test	S:\PRD-A1-DB\PRD-A1-Test.mdf
19	PRD-A1-Test_log	T:\PRD-A1-log\PRD-A1-Test_log.ldf
20	Test-PRD	S:\Test-A1-DB\Test-PRD.mdf
21	Test-PRD_log	T:\Test-A1-log\Test-PRD_log.ldf

6. The active CommServe node was installed as shown in the following screenshot:



7. The standby CommServe node was deployed on an EC2 virtual machine as shown in the following screenshots:

COMMVAULT

Select Roles

- ☐ Virtualization
- ☐ File server
- ☐ Office365
- ☐ Cloud apps
- ☐ Databases
- ☐ Laptop
- ☐ Backup node
- ☐ MediaAgent
- ☐ Big data
- ☒ **Commserve failover**

[Learn more about roles](#)

COMMVAULT

Failover Install Option

☐ This is an active node.
☒ This is a standby node.

Note: failover instance will be installed first on the standby node before installing the commserve instance.


8. Three MediaAgents were registered with the active CommServe node at the primary site. The first agent, a1-aws-ma , was used by the primary storage library. The second agent, a1-aws-m3, was used by the secondary storage library. The third agent, a1-aws-m4, was used by the storage library dedicated for CommServe LiveSync.

In addition to these roles, all three MediaAgents also served as Monitoring Nodes for the purpose of failing over the active CommServe node to the standby CommServe node during a disaster. The following screenshot shows the status as Monitoring Nodes.

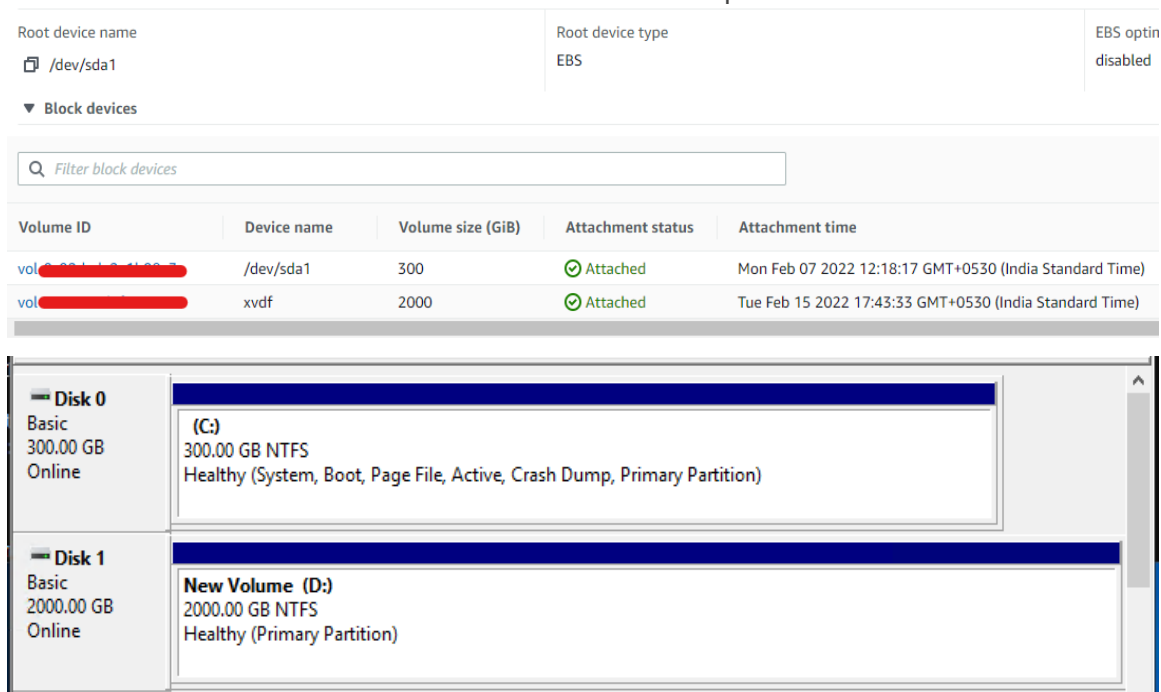
Monitoring Nodes Add

Nodes	Status	Last sync time	Action
a1-aws-ma	Active	Feb 20, 8:36 PM	
a1-aws-ma3	Active	Feb 20, 8:35 PM	
a1-aws-ma4	Active	Feb 20, 8:35 PM	

9. A standby SQL database instance, named SQLDR_EC2 was installed in an EC2 virtual machine as shown in the following screenshot. In case the production SQL database is corrupted or becomes unavailable, we can use CommServe to restore backed up data to this standby node and resume operations with it.



10. The MediaAgents a1-aws-ma and a1-aws-ma3, which are associated with the primary storage library and secondary storage library, were provisioned one 2 TB EBS volume each as shown in the following screenshot. These EBS volumes are needed for the Deduplication Database.

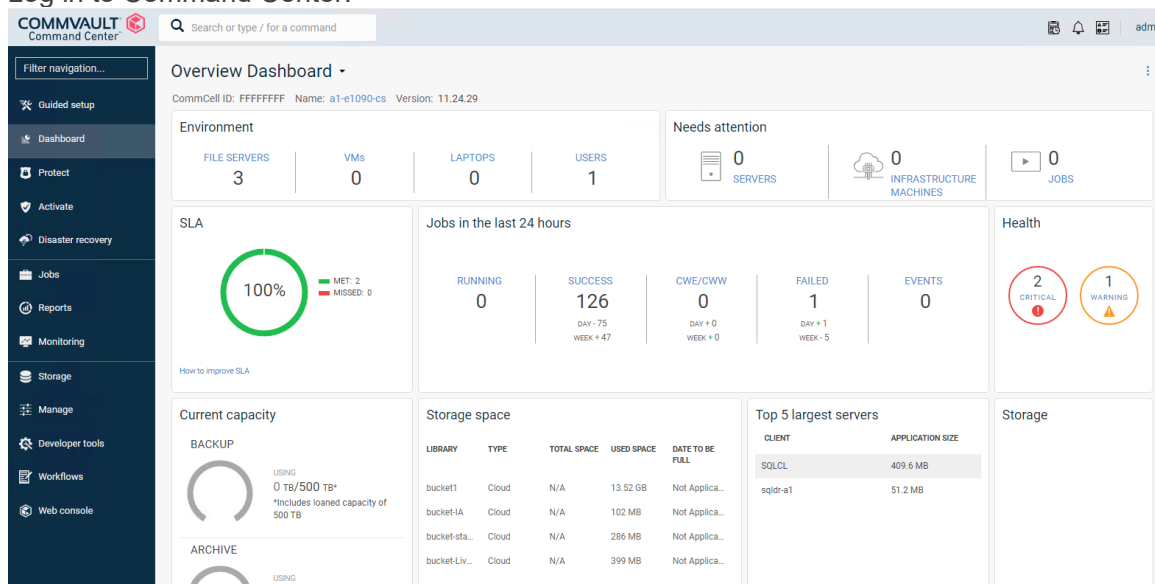


Test 2: Create Storage Libraries and Backup Plans

This test case walks you through the process of creating storage libraries and backup plans and registering an SQL database as a backup target.

1. To create the storage libraries on the active CommServe node, complete the following steps:

a. Log in to Command Center.



b. From the left navigation panel, select **Storage > Cloud > Add > Cloud Storage**. The following screenshot shows the task for S3 bucket a1-e1090-backup. MediaAgent a1-aws-ma was selected. The Storage class selected was **Standard** for the primary backup. The 2 TB EBS volume already attached to this virtual machine was designated as the Deduplication DB location.

Storage

Type: Amazon S3

MediaAgent: a1-aws-ma

Service host: s3.us-west-1.amazonaws.com

Authentication: Access and secret keys

Credentials: cred

Bucket: a1-e1090-backup

Storage class: Standard

☒ Use deduplication

Deduplication DB location

MediaAgent ↑	DDB Location
a1-aws-ma	D:\DDB-logs

Equivalent API

Cancel Save

- c. Repeat the same task to create the secondary storage library. Notice a different storage class, **Standard - Infrequent access**, was selected. Also, a different MediaAgent, a1-aws-ma3, was selected. The 2 TB EBS volume already attached to this virtual machine was designated as the Deduplication DB location.

Storage

Type	Amazon S3
MediaAgent	a1-aws-ma3 +
Service host	s3.us-west-1.amazonaws.com
Authentication	Access and secret keys
Credentials	cred + ✎
Bucket	a1-e1090-backup2
Storage class	Standard - Infrequent access

☒ Use deduplication

Deduplication DB location [Add](#)

MediaAgent ↑	DDB Location
a1-aws-ma3	D:\DDB-logs

- d. A third storage library was created for the CommServe LiveSync feature. This feature enables automatic and manual failover of the CommServe application. This storage library was pointed to the S3 bucket a1-e1090-backup (a dedicated bucket may be used for this purpose) and MediaAgent a1-aws-ma4.

Access Proxy Advanced

Name	bucket-LiveSync
Type	Amazon S3
MediaAgent	a1-aws-ma4
Service host	s3.us-west-1.amazonaws.com
Authentication	Access and secret keys
Credentials	cred + ✎
Bucket	a1-e1090-backup
Storage class	Standard

The following screenshot shows that three storage libraries were created:

Storage /

Cloud [Add](#) ⚙

Name ↑	Status	Capacity	Free space	Actions
bucket-IA	Online	N/A	N/A	⋮
bucket-LiveSync	Online	N/A	N/A	⋮
bucket-standard	Online	N/A	N/A	⋮

2. To create a backup plan on the active CommServe node, complete the following steps:
- From the left navigation panel, select **Manage > Plans > Create Plan**.

- b. Select the **Primary** and **SecondaryCopy** backup destinations. We selected the previously created bucket-standard for the Primary backup destination and bucket-IA for the SecondaryCopy backup destination. There are other backup options, such as **Backup Frequency**, **Snapshot options**, and **Database options** that you can set on this screen, as shown in the following screenshot:

Create server backup plan

Plan name: Backup-Plan-AUXCopy

Backup destinations [Add copy](#)

Name ↑	Storage	Retention period
Primary	bucket-standard	1 Month
SecondaryCopy	bucket-IA	1 Month

RPO ⓘ

Backup frequency: Runs every 1 Day(s)

Start time: 09 : 00 PM

☐ Add full backup

Backup window: Monday through Sunday : All day

Full backup window: Monday through Sunday : All day

Folders to backup ⓘ

Snapshot options

Database options ⓘ

Override restrictions

3. The next step is to register the on-premises, primary SQL instance and cloud-based, standby SQL instance.
- a. From the left navigation panel, select **Protect > Databases > Add SQL Server**. Enter the information for the on-premises SQL virtual machines and previously created backup plan. HDPS will install agent software on the virtual machines and discover the SQL database.

Add SQL server

OS Type: ☒ Windows ☐ Unix and Linux

Database server name: cyberdyne-sql1.juno.com cyberdyne-sql2.juno.com

User name: Juno\Administrator

Password:

Plan: Backup-Plan-AUXCopy

Installation location: Enter the installation location

SQL server account: Optional

- b. Repeat the previous step to register the standby SQL virtual machine running in AWS EC2. Attach the same backup plan as before.

Add SQL server

OS Type
☒ Windows
☐ Unix and Linux

Database server name

SQLDR-A1.juno.com

User name
Juno\Administrator

Password
.....

Plan
Backup-Plan-AUXCopy

Installation location
Optional
Enter the installation location

SQL server account
Optional



Note: For additional steps on how to set up a SQL client from CommServe Command Center, refer to the following video:
<https://www.youtube.com/watch?v=txjf829kb4U>.

4. Create a second backup plan to facilitate CommServe LiveSync, similar to the backup plan created in step 2. Use the third bucket, bucket-LiveSync, for this plan. This backup plan ensures that the metadata from the active CommServe node is available if we ever need to fail over to the standby CommServe node.

Plans View All							
Plan name ↑	Plan type	Associated entities	RPO	Number of copies	Status	Action	
Backup-Plan-AUXCopy	Server	2	1 days	3	Enabled		
CommServe-LS	Server	0	1 days	2	Enabled		

Test 3: Create a Full Backup of the SQL Database

This test case demonstrates how to create a full backup of an on-premises, clustered SQL application to the cloud.

- From Command Center, select **Protect > Databases**. This displays the database instances registered in CommServe. In the following screenshot, CYBERDYNE-SQL is the production database instance and SQLDR-EC2 is the secondary database instance.

Name	Server	Database engine	Status	Last backup	Application s...	SLA status	Plan
A1-AWS-CS\COMMMVAULT	a1-aws-cs	SQL Server	Not Ready	Feb 17, 3:41 AM	0 B	Excluded	Not Assigned
A1-E1090-CS\COMMMVAULT	a1-e1090-cs_2	SQL Server	Not Ready	Feb 20, 10:00 PM	0 B	Excluded	Not Assigned
CYBERDYNE-SQL	SQLCL	SQL Server	Ready	Feb 20, 9:29 PM	412 MB	Met	Backup-Plan-AUXC...
SQLDR-A1\SQLDR_EC2	sqlldr-a1	SQL Server	Ready	Feb 20, 9:17 PM	49 MB	Met	Backup-Plan-AUXC...

- Click on the dots under the **Actions** column and click **Back up**.

Name	Server	Database engine	Status	Last backup	Application s...	SLA status	Plan	Actions
A1-AWS-CS\COMMMVAULT	a1-aws-cs	SQL Server	Not Ready	Feb 17, 3:41 AM	0 B	Excluded	Not Assigned	⋮
A1-E1090-CS\COMMMVAULT	a1-e1090-cs_2	SQL Server	Not Ready	Feb 20, 10:00 PM	0 B	Excluded	Not Assigned	⋮
CYBERDYNE-SQL	SQLCL	SQL Server	Ready	Feb 20, 9:29 PM	412 MB	Met	Backup-Plan-AUXC...	⋮
SQLDR-A1\SQLDR_EC2	sqlldr-a1	SQL Server	Ready	Feb 20, 9:17 PM	49 MB	Met	Backup-Plan-AUXC...	⋮

- For **Backup level**, select **Full** and click **OK**.

Select backup level

Backup level

☒ Full

☐ Transaction log

☐ Differential

- Under **Jobs**, follow the backup job progress.

COMMVAULT Command Center

Search or type / for a command

Filter navigation...

Active jobs | Job history

Active jobs View: All 1 Total | 1 Running | 0 Pending | 0 Waiting | 0 Queued | 0 Suspended

Job Id	Operation	St...	Server	Agent t...	Subclie...
16363	Backup	Running	SQLCL	SQL Server	default

16363 - [Backup]

View job details | Suspend | Kill | Mon

Job summary

Type	Backup type
Backup	Full
Current phase	Status
Database Backup	Running
Progress	Elapsed time
10%	15 sec
Source client computer	Instance
SQLCL	CYBERDYNE-SQL
Subclient	Last update time
default	Feb 20, 2022 10:19:57 PM
Start time	Job started by
Feb 20, 2022 10:19:35 PM	admin (Started interactively)
Storage policy	Encryption enabled
Backup-Plan-AUXCopy	No
Total objects	No of objects
9	1

View more

Events | Database status

No events found

Reference Architecture

5. Upon completion of the backup job, select **Jobs > Job History** to get the full details of the task.

16363 - [Backup] ×

[View job details](#)
[Resubmit](#)
[Restore](#)
[More actions ▾](#)

Job summary

Type	Backup type
Backup	Full
Status	Source client computer
Completed	SQLCL
Instance	Subclient
CYBERDYNE-SQL	default
Start time	End time
Feb 20, 2022 10:19:35 PM	Feb 20, 2022 10:20:29 PM
Job started by	Storage policy
admin (Started interactively)	Backup-Plan-AUXCopy
Encryption enabled	No of objects
No	9
Agent	Size of application
SQL Server	412.25 MB
Data written	Savings percentage
10.13 MB	97.54 %
Data transferred on network	Transfer time
10.12 MB	26 sec
Failures	Average throughput
0 Folders, 0 Files	55.74 GB/hr
Load	
(DDB lookup: 100%)	

Test 4: Create an Incremental Backup

This test case demonstrates how to create an incremental backup of an on-premises, clustered SQL application to the cloud. Incremental backups contain only data that is new or was changed since the last backup operation. By doing this, we reduce the amount of data that needs to be sent across the network and stored on disks or in the cloud.

1. From Command Center, click **Protect > Databases**. Select the database for incremental backup and click **Back up**.

Name	Server	Database engine	Status	Last backup	Application s...	SLA status	Plan	Actions
A1-AWS-CS\COMMMVAULT	a1-aws-cs	SQL Server	Not Ready	Feb 17, 3:41 AM	0 B	Excluded	Not Assigned	⋮
A1-E1090-CS\COMMMVAULT	a1-e1090-cs_2	SQL Server	Not Ready	Feb 20, 10:00 PM	0 B	Excluded	Not Assigned	⋮
CYBERDYNE-SQL	SQLCL	SQL Server	Ready	Feb 20, 9:29 PM	412 MB	Met	Backup-Plan-AUXC...	⋮
SQLDR-A1\SQLDR_EC2	sqlldr-a1	SQL Server	Ready	Feb 20, 9:17 PM	49 MB	Met	Backup-Plan-AUXC...	⋮

2. Select **Differential** and click **OK**.

Select backup level

Backup level

☐ Full
☐ Transaction log
☒ Differential

3. Review the summary of the completed differential backup task. Notice the backup type is Differential and the backup Savings percentage was 97.55% as follows:

Job summary	
Type	Backup type
Backup	Differential
Status	Source client computer
Completed	SQLCL
Instance	Subclient
CYBERDYNE-SQL	default
Start time	End time
Feb 21, 2022 1:03:31 AM	Feb 21, 2022 1:04:46 AM
Job started by	Storage policy
admin (Started interactively)	Backup-Plan-AUXCopy
Encryption enabled	No of objects
No	9
Agent	Size of application
SQL Server	88.50 MB
Data written	Savings percentage
2.17 MB	97.55 %
Data transferred on network	Transfer time
2.17 MB	43 sec
Failures	Average throughput
0 Folders, 0 Files	7.24 GB/hr
Load	
(Write: 23.07%, DDB lookup: 76.92%)	

Test 5: Perform an In-Place Restore

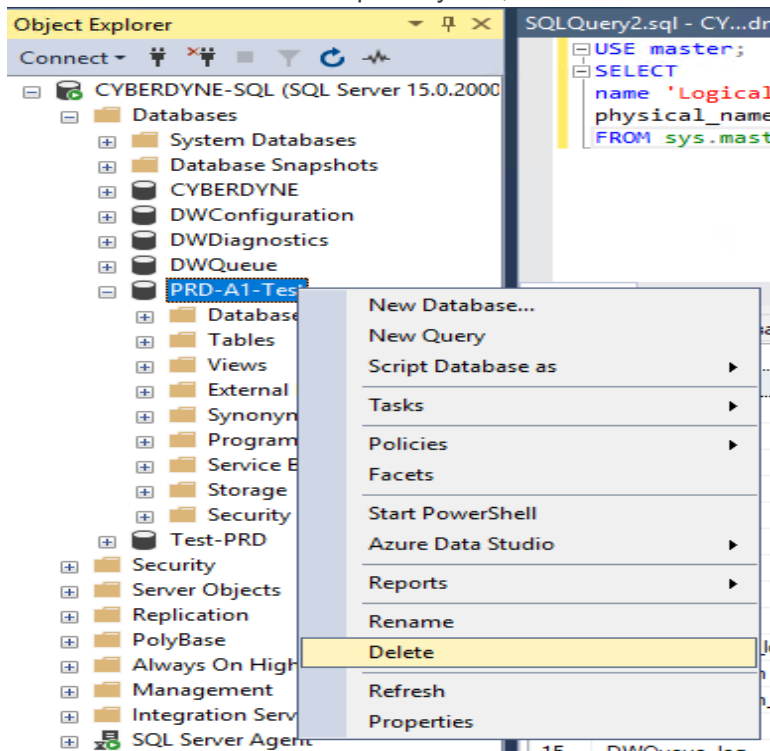
This test case demonstrates how a database can be restored from cloud backups to the primary site to restart operations after the production SQL database is deleted or corrupted. It also proves that the data is consistent from the last backup point.

1. We created a sample database called PRD-A1-TEST and populated a table called StudentDetails with data as follows:

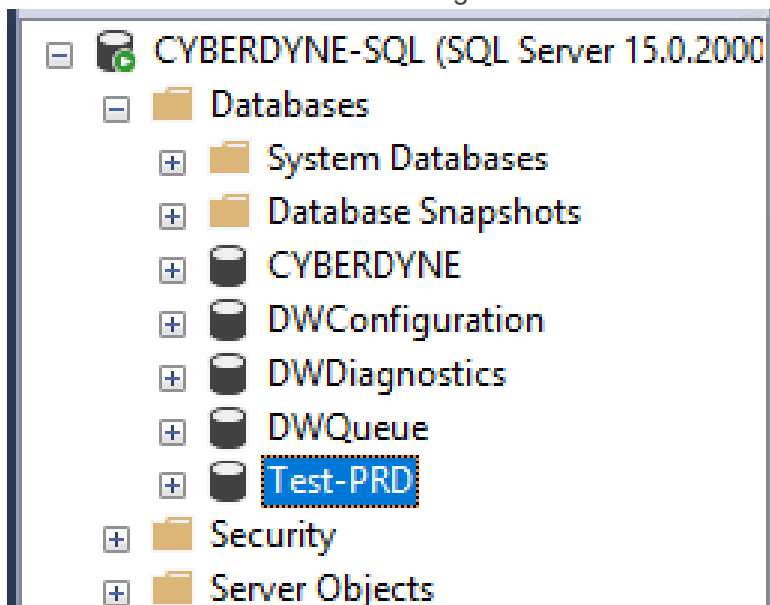
The screenshot shows the SQL Server Enterprise Manager interface. On the left, the Object Explorer displays the hierarchy of the CYBERDYNE-SQL (SQL Server 15.0.2) instance. The 'PRD-A1-Test' database is selected under the 'Databases' folder. The right pane shows the SQL Query Editor with the query `SELECT * from StudentDetails`. Below the query editor, the 'Results' tab displays the data from the 'StudentDetails' table.

	Student Name	Roll#
1	Letha241	00A8
2	Mollie1975	48KWUQ74
3	Quintin5	T3
4	Kathem49	N
5	Felisa221	O
6	Shenita2027	40
7	Amita2004	J
8	Woodrow2009	E
9	Chantel318	NULL
10	Aisha2010	X3LQ8
11	Christopher1976	SI15
12	Stephan1976	W4575

2. To simulate data loss at the primary site, we deleted the database PRD-A1-TEST.



The database PRD-A1-TEST no longer exists.



3. To restore the deleted database from a backup, complete the following steps:

- a. Under Command Center, click **Protect > Databases**. The deleted database was part of the primary SQL database instance CYBERDYNE-SQL, so we selected it and clicked **Restore**.

COMMVAULT Command Center

Search or type / for a command

Instances Databases Instant clones

Filter navigation...

Guided setup

Dashboard

Protect

Virtualization

File servers

Databases

Laptops

> Applications

View: All 2 Met 0 Missed 2 Excluded 0 May miss SLA

Name	Server	Database engine	Status	Last backup	Application s...	SLA status	Plan	Actions
A1-AWS-CS\COMMVAULT	a1-aws-cs	SQL Server	Not Ready	Feb 17, 3:41 AM	0 B	Excluded	Not Assigned	⋮
A1-E1090-CS\COMMVAULT	a1-e1090-cs_2	SQL Server	Not Ready	Feb 20, 11:00 PM	0 B	Excluded	Not Assigned	⋮
CYBERDYNE-SQL	SQLCL	SQL Server	Ready	Feb 20, 10:19 PM	412 MB	Met	Backup-Plan-AUXC...	⋮
SQLDR-A1\SQLDR_EC2	sqlldr-a1	SQL Server	Ready	Feb 20, 9:17 PM	49 MB	Met	Backup-Plan-AUXC...	⋮

Back up

Restore

Configure replication

- b. From the list of databases that were part of the selected SQL instance, we selected the deleted database, PRD-A1-TEST, and clicked **Restore**.

Instances / CYBERDYNE-SQL /

Select database to restore from CYBERDYNE-SQL

Restore

	Database	Server	Recovery model	Latest backup time	Size
<input type="checkbox"/>	CYBERDYNE	SQLCL	FULL	Feb 20, 10:19:57 PM	16 MB
<input type="checkbox"/>	DWConfiguration	SQLCL	FULL	Feb 20, 10:20:01 PM	16 MB
<input type="checkbox"/>	DWDiagnostics	SQLCL	SIMPLE	Feb 20, 10:20:02 PM	1.05 GB
<input type="checkbox"/>	DWQueue	SQLCL	SIMPLE	Feb 20, 10:20:04 PM	16 MB
<input checked="" type="checkbox"/>	PRD-A1-Test	SQLCL	FULL	Feb 20, 10:20:10 PM	16 MB
<input type="checkbox"/>	Test-PRD	SQLCL	FULL	Feb 20, 10:20:15 PM	683.59 GB

- c. For the Restore option (**In place**, **Out of place** and **Restore to disk**), we selected **In place** to restore the database to the original location at the primary site and clicked **Submit**.

Restore of PRD-A1-Test

In place Out of place Restore to disk

Restore database to the original location and instance.

Destination server SQLCL

Destination instance CYBERDYNE-SQL

Destination database PRD-A1-Test

☐ When the job completes, notify me via email

Warning: Production databases PRD-A1-Test will be overwritten on CYBERDYNE-SQL.



Note: An **In place** restore overwrites the source database, if it still exists, at the original location.

- d. Review the restore job summary in the Job screen as follows:

Job summary

Type	Restore
Status	Completed
Source server	SQLCL
Source instance	CYBERDYNE-SQL
Destination server	SQLCL
Destination instance	CYBERDYNE-SQL
Start time	Feb 20, 2022 11:29:40 PM
End time	Feb 20, 2022 11:29:51 PM
Job started by	admin (Started interactively)
Total objects	1
No. of successes	1

Job details

Total no. of files to be restored	1
Media agent	a1-aws-ma,a1-aws-ma4
No of files restored	1
Size of application	9.38 MB
Transfer time	11 sec
Average throughput	3.00 GB/hr
Failures	0 Folders, 0 Files

4. Review the database PRD-A1-TEST after it was restored. We can see that the information in table StudentDetails matches what was shown in step 1 as follows;

Object Explorer

- CYBERDYNE-SQL (SQL Server 15.0.2)
 - Databases
 - System Databases
 - Database Snapshots
 - CYBERDYNE
 - DWConfiguration
 - DWDiagnostics
 - DWQueue
 - PRD-A1-Test**
 - Database Diagrams
 - Tables
 - System Tables
 - FileTables
 - External Tables
 - Graph Tables
 - dbo.StudentDetails
 - Columns
 - Keys
 - Constraints
 - Triggers
 - Indexes
 - Statistics
 - Views
 - External Resources
 - Synonyms
 - Programmability
 - Service Broker
 - Storage
 - Security
 - Test-PRD
 - Security
 - Server Objects
 - Replication

SQLQuery3.sql - CY...dministrator (130))*

```
SELECT * from StudentDetails
```

Results

	Student Name	Roll#
1	Letha241	00A8
2	Mollie1975	48KWUQ74
3	Quintin5	T3
4	Kathem49	N
5	Felisa221	O
6	Shenita2027	40
7	Amita2004	J
8	Woodrow2009	E
9	Chantel318	NULL
10	Aisha2010	X3LQ8
11	Christopher1976	SI15
12	Stephen1976	W4575

Test 6: Perform an Out-of-Place Restore to Cloud

This test case demonstrates the Out-of-Place Restore feature. It allows us to restore backed up SQL databases to any other database instance. For example, we can restore a database to a standby instance in the cloud, perform development and testing on the replica, and not impact the production instance.

1. Prepare the standby EC2 SQL instance by provisioning two volumes to host the database replica.

Details

Security

Networking

Storage

Status checks

Monitoring

Tags

▼ Root device details

Root device name

/dev/sda1

Root device type

EBS

EBS optimization

disabled


▼ Block devices

Q

Filter block devices

Volume ID	Device name	Volume size (GiB)	Attachment status	Attachment time	Encrypted	KMS key ID	Delete on termination
vol- <div></div>	/dev/sda1	50	<div>Attached</div>	Thu Feb 03 2022 15:17:40 ...	No	-	Yes
vol- <div></div>	xvdf	2000	<div>Attached</div>	Thu Feb 17 2022 08:33:00 ...	No	-	No
vol- <div></div>	xvda	2000	<div>Attached</div>	Thu Feb 17 2022 08:34:05 ...	No	-	No

Within the guest standby instance, we formatted the two volumes and configured them as drives S and T.



2. Establish a replication relationship between the production database instance and the standby database instance as follows:
 - a. Select **Protect > Databases > Instances** and select the production database instance. Under **Actions**, click **Configure replication**.

COMMVAULT
Command Center

Filter navigation...

Guided setup

Dashboard

Protect

Virtualization

File servers

Databases

Laptops

Applications

Search or type / for a command

Instances

Databases

Instant clones

Add instance

Add server

View

All

2 Met

0 Missed

2 Excluded

0 May miss SLA

Name	Server	Database engine	Status	Last backup	Application s...	SLA status	Plan	Actions
A1-AWS-CS/COMMVAULT	a1-aws-cs	SQL Server	Not Ready	Feb 17, 3:41 AM	0 B	Excluded	Not Assigned	
A1-E1090-CS/COMMVAULT	a1-e1090-cs_2	SQL Server	Not Ready	Feb 21, 7:30 PM	0 B	Excluded	Not Assigned	
CYBERDYNE-SQL	SQLCL	SQL Server	Ready	Feb 21, 5:01 PM	412 MB	Met	Backup-Plan-AUXC...	
SQLDR-A1\SQLDR_EC2	sqlldr-a1	SQL Server	Ready	Feb 21, 5:03 PM	49 MB	Met	Backup-Plan-AUX	

Back up

Restore

Configure replication

- b. On the Source screen, enter a name for the new replication relationship. Select the source Servers and Instances. Command Center displays a list of databases from the specified database instance for replication. Select the databases that you want to replicate and then click

Next.

COMMVAULT Command Center

Search or type / for a command

Filter navigation...

Guided setup

Dashboard

Protect

Activate

Disaster recovery

Jobs

Reports

Monitoring

Storage

Manage

Developer tools

Workflows

Web console

Source

Target

Override options

Source

Name: ReplicationGroup-Test

Servers: SQLCL

Instances: CYBERDYNE-SQL

Databases

<input type="checkbox"/>	Name	Latest Backup Time	Size
<input type="checkbox"/>	CYBERDYNE	Feb 21, 5:01:21 PM	16 MB
<input type="checkbox"/>	DWConfiguration	Feb 21, 5:01:22 PM	16 MB
<input checked="" type="checkbox"/>	PRD-A1-Test	Feb 21, 5:01:24 PM	16 MB
<input type="checkbox"/>	Test-PRD	Feb 21, 5:01:26 PM	683.59 GB

Next

- c. On the Target screen, select the target Servers and Instances and click **Next**.

COMMVAULT Command Center

Search or type / for a command

Filter navigation...

Guided setup

Dashboard

Protect

Activate

Disaster recovery

Jobs

Reports

Monitoring

Storage

Manage

Developer tools

Workflows

Web console

SQL Server /

Configure replication group - SQL Server

Source

Target

Override options

Target

Servers: sqlldr-a1

Instances: SQLDR-A1\SQLDR_EC2

Back

Next

- d. On the **Override Options** screen, specify where to restore the database files on the standby database instance and click **Submit**.

COMMVAULT Command Center

Search or type / for a command

Filter navigation...

Guided setup

Dashboard

Protect

Activate

Disaster recovery

Jobs

Reports

Monitoring

Storage

Manage

Developer tools

Workflows

Web console

Source

Target

Override options

Override options

Redirect options

PRD-A1-Test

Database name: PRD-A1-Test

Data file path: S:\PRD-A1-DB

Log file path: T:\PRD-A1-log

Advanced options

Sync delay: 0 hours

Recovery Type: ☒ No recovery ☐ Stand by

Equivalent API

Back

Submit

Reference Architecture

- e. To view the new group, select **Disaster Recovery > Replication > Replication groups** to view the new group.

Summary

Source	SQLCL
Source instance	CYBERDYNE-SQL
State	Enabled
Destination	sqlldr-a1
Destination instance	SQLDR-A1\SQLDR_EC2

Advanced options

Replication window: Monday through Sunday : All day

Replication monitor

Source	Source instance	Destination instance	Last sync time	Status
PRD-A1-Test	CYBERDYNE-SQL	SQLDR-A1\SQLDR_EC2	Feb 21, 5:02:18 PM	In Sync

3. Now we restore a database from the production instance to the standby instance in the cloud as follows:
- Select **Protect > Databases**, select the production database instance, and then under **Actions**, click **Restore**.

Databases

Name	Server	Database engine	Status	Last backup	Application s...	SLA status	Plan	Actions
A1-AWS-CS\COMMVAULT	a1-aws-cs	SQL Server	Not Ready	Feb 17, 3:41 AM	0 B	Excluded	Not Assigned	...
A1-E1090-CS\COMMVAULT	a1-e1090-cs_2	SQL Server	Not Ready	Feb 21, 8:00 PM	0 B	Excluded	Not Assigned	...
CYBERDYNE-SQL	SQLCL	SQL Server	Ready	Feb 21, 5:01 PM	412 MB	Met	Backup-Plan-AUXC...	Back up Restore Configure replication
SQLDR-A1\SQLDR_EC2	sqlldr-a1	SQL Server	Ready	Feb 21, 5:03 PM	49 MB	Met	Backup-Plan-AUXC...	

- In the next screen, select the database to restore and then click **Restore**.
- In the **Restore** screen, click the **Out of place** tab. Enter a name for the destination database and specify the restore locations. We chose **No Recovery** for **Recovery types** so that the database replica does not come online automatically to avoid conflicting with the production database instance. Click **Submit**.

Select database to restore from CYBERDYNE-SQL

Database	Server	Recovery model
<input type="checkbox"/> CYBERDYNE	SQLCL	FULL
<input type="checkbox"/> DWConfiguration	SQLCL	FULL
<input type="checkbox"/> DWDiagnosics	SQLCL	SIMPLE
<input type="checkbox"/> DWQueue	SQLCL	SIMPLE
<input checked="" type="checkbox"/> PRD-A1-Test	SQLCL	FULL
<input type="checkbox"/> Test-PRD	SQLCL	FULL

Out of place

Restore database using a different name, file paths or instance.

Destination server: Select an instance

Destination instance: Select an instance

Data files:

Verify that the data file path and log file path exist on the destination server:

Destination database	Data file path	Log file path
PRD-A1-Test	S:\PRD-A1-DB	T:\PRD-A1-log

Additional options:

Recovery types:

- ☐ Recovery
- ☒ No recovery
- ☐ Stand by



Note: Under **Additional options**, if **Recovery** is selected, the database replica is created in a usable state. If **No recovery** is selected, the database replica is created in a Restoring state, and the database administrator must bring it online manually.

- d. Click **Jobs** to verify that the restore task completed successfully.

Jobs / Job 16477 - [Restore] Resubmit More actions


Job summary

Type	Restore
Status	Completed
Source server	SQLCL
Source instance	CYBERDYNE-SQL
Destination server	sqlldr-a1
Destination instance	SQLDR-A1\SQLDR_EC2
Start time	Feb 21, 2022 8:16:43 PM
End time	Feb 21, 2022 8:18:25 PM
Job started by	admin (Started interactively)
Total objects	1
No. of successes	1

Job details

Total no. of files to be restored	1
Media agent	a1-aws-ma,a1-aws-ma4
No of files restored	1
Size of application	19.88 MB
Transfer time	20 sec
Average throughput	3.49 GB/hr
Failures	0 Folders, 2 Files


4. Review the database PRD-A1-Test to verify that it was restored to the standby database instance. Note that we chose the **No Recovery** option, which is why the database is marked as Restoring.




Test 7: Manual Failover of CommServe

This test case demonstrates the manual failover capability of CommServe. This feature facilitates business continuity of the backup application during a planned outage or site disaster.

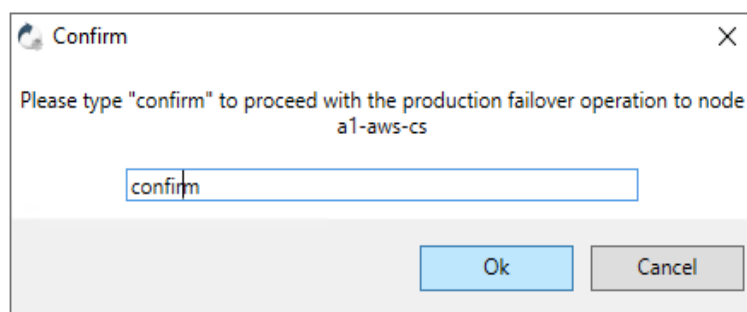
1. In Command Center of the active CommServe node, select **Manage > System > CommServe LiveSync** and then toggle **Status** and **Automatic failover** to on.



2. On the active CommServe node, select **Windows Start > Commvault > Process Manager (Failover Assistant)**.
3. After the application opens, click the **Failover Assistant** tab. Verify that the node names listed under Production Node and Passive Node are as expected.



4. Click **Initiate Failover** to start the manual failover process. Type **confirm** and click **OK**.



5. Monitor the manual failover progress as shown in the following screenshot:

Process Manager

General Plugin Processes Services Logging Troubleshooting Failover Assistant

Node Information

Production Node	Node Name
Passive Node	a1-e1090-cs_2
	a1-aws-cs

Failover Details


Failover To	a1-aws-cs
Failover Type	Production

Production failover to a1-aws-cs is in progress ...

Hide Details

02/21/2022 21:33:48	Checking node [a1-e1090-cs_2] is reachable...
02/21/2022 21:33:48	Sending the request to make node [a1-e1090-cs_2] passive...
02/21/2022 21:33:49	Disabling all activities... Done
02/21/2022 21:33:49	Disabling the scheduling activity... Done
02/21/2022 21:33:49	Comparing the database and binaries version... Done
02/21/2022 21:33:50	Confirming no CommServe SQL LiveSync jobs are running... Done
02/21/2022 21:33:50	Stopping and disabling all the services...

6. When prompted that the failover is complete, click **OK**.



7. Review the services to verify that all CommServe services are now running on a1-aws-cs (left side) while all services are stopped on a1-e1090-cs_2 (right side).

a1-aws-cs (New Active Node)

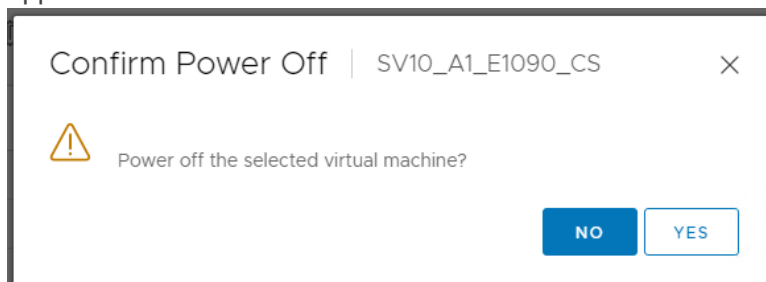
a1-e1090-cs_2 (Original Active Node)

8. To swap the CommServe nodes back to the original roles, click **Initiate Failover** again.

Test 8: Automatic Failover of CommServe

This test case demonstrates the automatic failover capability of CommServe. The feature facilitates business continuity of the backup application during an unplanned outage. During such an event, the three MediaAgents, configured as Monitoring Nodes, trigger an automatic failover when they are unable to reach the active CommServe node for 30 minutes.

1. To simulate a failure with the active CommServe node, we powered off the virtual machine where the application was installed.



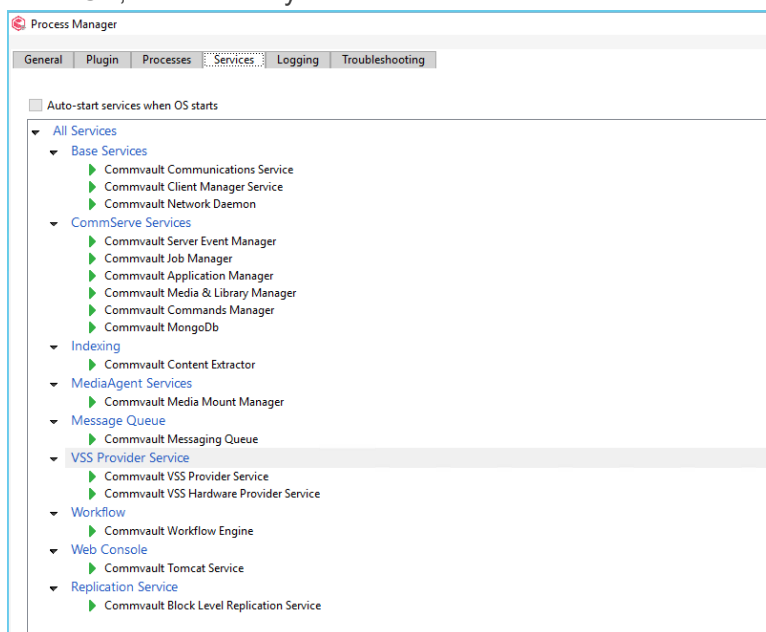
2. On the standby CommServe node, review the file **C:\Program Files\Commvault\ContentStore2\Log Files\CommServeLiveSyncMonitoring.log** to see that there are errors about losing contact with the active node.

```

- Failed to send beat message to [A1-E1090-CS.juno.com*a1-e1090-cs_2*8405]: 0xE8090010:{cvhb::sendBeatTo(1075)} + {CS
CS is not active
FailoverHandler invoked via Client A1-AWS-MD-juno.com

```

3. Review the services after about 30 minutes, and you can see that all services on the standby node, A1-AWS-CS, automatically started.



4. To swap the CommServe nodes back to the original roles, click **Initiate Failover** again.



Note: After an automatic failover, CommServe does not automatically fail back even if the original CommServe node is online again. You must initiate a failback by clicking **Initiate Failover**.


Test 9: Unplanned Outage with Data Return

This test case demonstrates the recovery procedure to resume business operations in the cloud after an unplanned outage at the primary site. It relies on LiveSync to automatically fail over CommServe to the cloud. Then, an Out-of-Place Restore is used to restore backed up data to the standby SQL EC2 virtual machine.




Note: This test requires an established database replication relationship between the production database instance and standby database instance. This procedure is documented in section [Test 6](#).

1. We simulated an unplanned outage at the primary site by abruptly powering off the SQL virtual machines and active CommServe node.
2. Review the services after about 30 minutes, and you can see that LiveSync automatically brought up CommServe services on the standby CommServe node as shown in the following screenshot:



3. The next step is restoring the latest backed up data to the standby SQL EC2 virtual machine.
 - a. Log in to Command Center on the standby CommServe node.
 - b. Select **Protect > Databases > Restore**. Perform an Out of Place restore job of the production database, similar to the following screenshot, and click **Submit**.



Reference Architecture

- c. Review the job summary to see that the restore job completed successfully.

Active jobs / Job 21659 - [Restore] Resubmit More actions ▾

Job summary

Type	Restore
Status	Completed
Source server	SQLCL
Source instance	CYBERDYNE-SQL
Destination server	sqlldr-a1
Destination instance	SQLDR-A1\SQLDR_EC2
Start time	Feb 23, 2022 3:29:59 AM
Job started by	admin (Started interactively)
Total objects	1
No. of successes	1

Job details

Total no. of files to be restored	1
Media agent	a1-aws-ma,a1-aws-ma4
No of files restored	1
Current throughput	N/A
Average throughput	N/A

- d. Open the restored database and verified that the data is accessible.

Object Explorer SQLQuery3.sql - SQL...Administrator (64)**

Connect ▾

SQLDR-A1\SQLDR_EC2 (SQL Server 15.0.2000.5 - JUN 2019)

- Databases
 - System Databases
 - Database Snapshots
 - PRD-A1-Test
 - Database Diagrams
 - Tables
 - System Tables
 - FileTables
 - External Tables
 - Graph Tables
 - dbo.StudentDetails
 - Columns
 - Keys
 - Constraints
 - Triggers
 - Indexes
 - Statistics
 - Views
 - External Resources
 - Synonyms
 - Programmability
 - Service Broker
 - Storage
 - Security
 - Server Objects
 - Replication
 - PolyBase
 - Always On High Availability
 - Management
 - Integration Services Catalogs
 - SQL Server Agent
 - XEvent Profiler

100 %

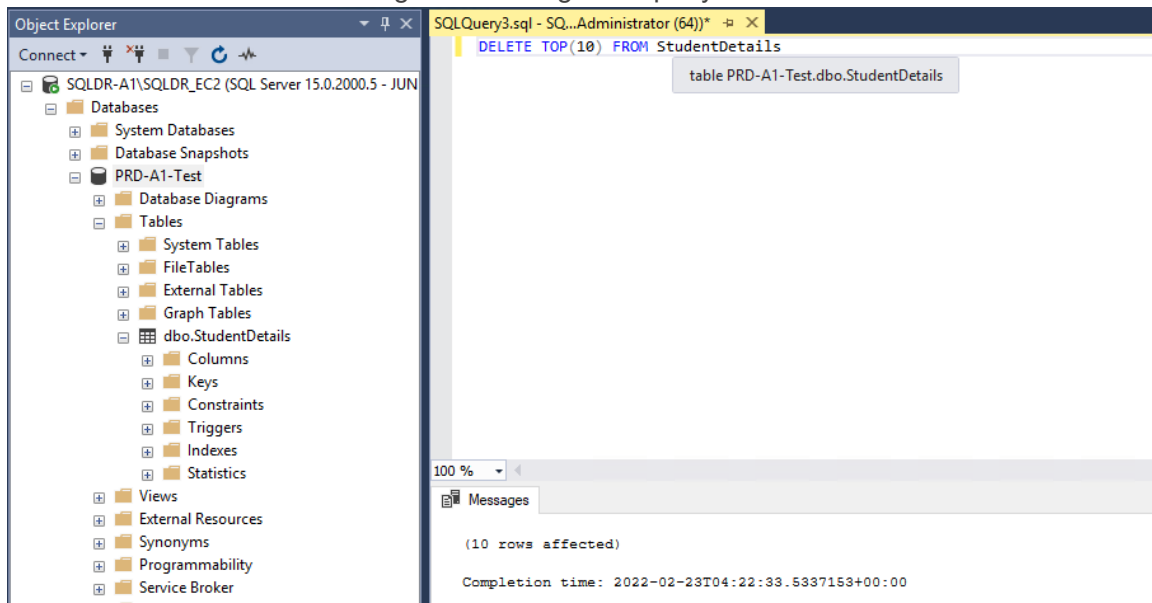
Results Messages

	Student Name	Roll#
1	Letha241	00A8
2	Mollie1975	48KWUQ74
3	Quintin5	T3
4	Kathem49	N
5	Felisa221	O
6	Shenita2027	40
7	Amita2004	J
8	Woodrow2009	E
9	Chantel318	NULL
10	Aisha2010	X3LQ8
11	Christopher1976	SI15
12	Stephan1976	W4575
13	Brinson972	I
14	Carpenter1994	52
15	David2	77

Query executed successfully.

4. The next step is to make updates to the database by deleting the first 10 records from the table StudentDetails. After, we will create a backup of the database.

- a. Delete the first 10 records using the following SQL query:



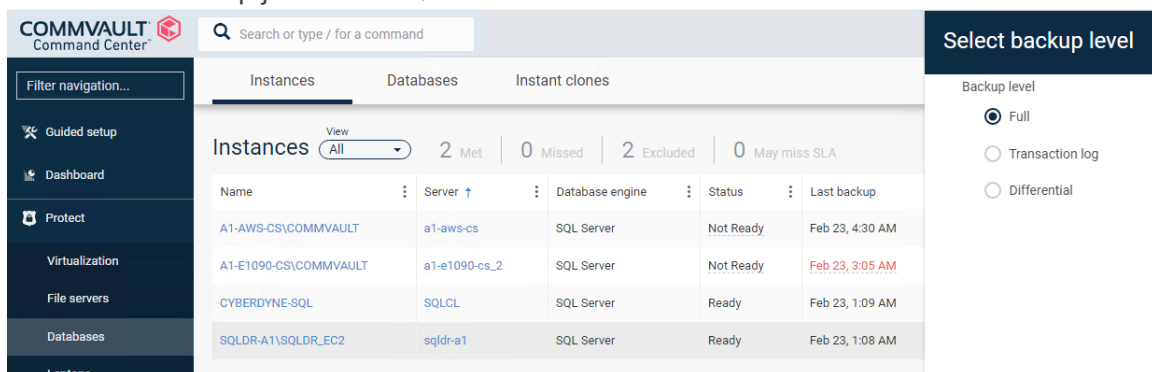
The following screenshot shows the first ten records contain different information now.

The screenshot shows the SQL Query Results window. The 'Results' tab is selected, displaying a table with 14 rows and 2 columns: 'Student Name' and 'Roll#'. The data is as follows:

	Student Name	Roll#
1	Christopher1976	SI15
2	Stephan1976	W4575
3	Brinson972	I
4	Carpenter1994	52
5	Doyle2	7T
6	Johnston1952	R
7	Booker6	7
8	Carvalho72	D80
9	Aguilera1993	X
10	Paris2022	E9D
11	Jerrod1982	L
12	Suarez82	2Q5
13	Abreu232	07C026
14	Shank212	2

Below the table, a green status bar indicates: 'Query executed successfully.'

- b. Initiate a full backup job of the SQL EC2 database instance.



- c. Review the Job summary to verify that the backup job completed successfully.

Active jobs / Job 21671 - [Backup]

Job summary

Type	Backup
Backup type	Full
Status	Completed
Source client computer	sqlldr-a1
Instance	SQLDR-A1\SQLDR_EC2
Subclient	default
Start time	Feb 23, 2022 4:45:50 AM
Job started by	admin (Started interactively)
Storage policy	Backup-Plan-AUXCopy
Encryption enabled	No
No of objects	4

Job details

Agent	SQL Server
Data transferred on network	N/A
Current throughput	N/A
Average throughput	N/A

5. The final step is to fail back to the primary site as follows:
- Power off the EC2 virtual machine hosting the standby database.
 - Power on the SQL virtual machines and active CommServe node at the primary site.
 - Perform an Out of Place Restore job with the source being the standby EC2 database instance and the destination being the production database instance.

Restore of PRD-A1-Test

In place

Out of place

Restore to disk

Restore database using a different name, file paths or instance.

Destination server

SQLCL

Destination instance

CYBERDYNE-SQL

Data files

Verify that the data file path and log file path exist on the destination server.

PRD-A1-Test

Destination database

PRD-A1-Test

-PRD-A1-Test

Data file path

S:\PRD-A1-DB

-PRD-A1-Test...

Log file path

T:\PRD-A1-log

- d. Review the Job summary to verify that the Restore job completed successfully.

Job summary		Job details	
Type	Restore	Total no. of files to be restored	1
Status	Completed	Media agent	a1-aws-ma,a1-aws-ma4
Source server	sqlldr-a1	No of files restored	1
Source instance	SQLDR-A1\SQLDR_EC2	Current throughput	N/A
Destination server	SQLCL	Average throughput	N/A
Destination instance	CYBERDYNE-SQL		
Start time	Feb 23, 2022 5:46:41 AM		
Job started by	admin (Started interactively)		
Total objects	1		
No. of successes	1		

6. On the production database instance, verify that the restored database is consistent with the deletion made during the outage.

Results		Messages	
	Student Name	Roll#	
1	Christopher1976	SI15	
2	Stephan1976	W4575	
3	Brinson972	I	
4	Carpenter1994	52	
5	Doyle2	7T	
6	Johnston1952	R	
7	Booker6	7	
8	Carvalho72	D80	
9	Aguilera1993	X	
10	Paris2022	E9D	
11	Jerrod1982	L	
12	Suarez82	2Q5	
13	Abreu232	07C026	
14	Shank212	2	
15	Chen70	7A	

✓ Query executed successfully.

7. To swap the CommServe nodes back to their original roles, click **Initiate Failover** again.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
hitachivantara.com | community.hitachivantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
hitachivantara.com/contact

HITACHI is a registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Vantara LLC. Microsoft, Azure and Windows are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks and company names are properties of their respective owners.

Author: S/Basu August 2022