

Continuous Analytics for a Near-Cloud Solution with Global-Active Device and Hitachi Thin Image: VSP E1090 to VSP 5200

Using Hitachi Cloud Connect for Equinix

Hitachi Vantara
July 2022

Table of Contents

Notices and Disclaimer	3
Executive Summary.....	4
About This Guide.....	5
Intended Audience	5
Document Revisions	5
References	5
Comments	5
Introduction.....	6
Solution Overview	6
Benefits	6
Key Components.....	7
Engineering Validation.....	8
Validation Method.....	8
High Level Diagram.....	9
Hardware and Software.....	9
Test Scenarios	10
Guidelines and Recommendations.....	13
Validation Results.....	14
Test 1: Prepare the Environment.....	14
Test 2: Create GAD Pairs with Cascaded HTI Snapshots.....	22
Test 3: Create HTI Pairs.....	36
Test 4: Access Snapshots from AWS EC2 while GAD is in Pair state	38
Test 5: Multigeneration HTI snapshot.....	42
Test 6: Planned Outage of the Primary Storage System.....	47
Test 7: Unplanned Path Failure of the Primary Storage System	57

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test 8: Unplanned Outage of On-Premises Servers with Data Return..... 63

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 screenshots 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

Global-active device (GAD) is a data mirroring technology that allows you to maintain synchronous copies of data at remote sites. Hitachi GAD provides a solution to avoid situations when a data center is affected by a disaster that stops operations for a long time, while providing continuous server I/O. Hitachi Thin Image (HTI) creates instantaneous point-in-time snapshots that are space efficient. This reference architecture documents how to set up a Hitachi GAD disaster recovery (DR) solution with cascading Hitachi Thin Image snapshots using Hitachi Ops Center Protector. By combining the two replication technologies, we have a solution that is always protected with site-to-site replication while providing the flexibility to quickly snap-off data copies for analytics, data mining, testing, or other ad hoc requirements in a cloud environment.

The environment used for the validation includes a Hitachi Virtual Storage Platform (VSP) E1090 as the primary storage system and a VSP 5200 as the secondary storage system. Both storage systems are in a near cloud location (Equinix – a colocation data center, located in California).

To access data on the secondary storage system at the DR site, we used cloud compute from Amazon Web Services (AWS) in the form of Elastic Compute Cloud (EC2) virtual machines. Because Hitachi Thin Image can create up to 1,024 snapshots per source volume, pairing it with the very well-known scale out capability of cloud computing makes for a great combination.

Because the cost of owning and operating a second data center for the purpose of disaster recovery is significant, leasing a small footprint in a colocation data center is a cost-attractive alternative. Equinix, Inc.™ is a leading provider of such services. Through our collaboration with Equinix, Hitachi Vantara created a new near-cloud hybrid solution: **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 an 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 allow 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>.

About This Guide

This reference architecture documents how to set up a disaster recovery solution by combining Hitachi GAD with cascading Hitachi Thin Image snapshots using Ops Center Administrator and Ops Center Protector. 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.

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	July 2022	Hitachi Vantara LLC	Initial Release

References

- Hitachi Global-Active Device User Guide:
https://knowledge.hitachivantara.com/Documents/Management_Software/SVOS/9.8.2/Global-Active_Device
- Hitachi Thin Image User Guide:
https://knowledge.hitachivantara.com/Documents/Management_Software/SVOS/9.8.2/Local_Replication/Thin_Image
- Hitachi Ops Center Protector User Guide:
[https://knowledge.hitachivantara.com/Documents/Management_Software/Ops_Center/10.8.x/Ops_Center_Protector_\(formerly_Data_Instance_Director\)](https://knowledge.hitachivantara.com/Documents/Management_Software/Ops_Center/10.8.x/Ops_Center_Protector_(formerly_Data_Instance_Director))
- Hitachi Ops Center Administrator User Guide:
https://knowledge.hitachivantara.com/Documents/Management_Software/Ops_Center/10.8.x/Administrator
- Hitachi Ops Center Administrator - High Availability User Guide:
https://knowledge.hitachivantara.com/Documents/Management_Software/Ops_Center/10.8.x/High_Availability
- Implementation Guide: Global-Active Device Cloud Quorum in AWS:
https://knowledge.hitachivantara.com/Documents/Management_Software/SVOS/Global_Active_Device_Cloud_Quorum/Global_Active_Device_Cloud_Quorum_Documentation_Library

Comments

Send us your 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!

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Introduction

The environment used for the validation included a Virtual Storage Platform (VSP) E1090 used as the primary storage system and a VSP 5200 used as the secondary storage system, located in an Equinix colocation data center in California. To access data on snapshot volumes and data at primary storage system during an outage, we used cloud compute from AWS in the form of EC2 virtual machines.

This information was accumulated while setting up our environment to validate use cases for the Hitachi Cloud Connect for Equinix solution.

- A near-cloud Equinix colocation data center (SV10) located in San Jose, California.
- A cloud hosted by AWS in Northern California.

The relationship between the Near Cloud and Cloud Service provider is shown in *Figure 1*:




Figure 1: High Level Diagram



Note: The information shared here is specific to our requirements and can be used as a guideline or a starting point. We recommend conducting a proof-of-concept in a non-production, isolated test environment matching your production environment before implementing this solution.

Solution Overview

Hitachi Global-Active Device (GAD) provides a solution to avoid situations when a data center is affected by a disaster that stops operations for a long period of time. Three storage systems are required in a GAD implementation. The secondary storage system is placed in a second data center that is far (up to 500 km) from the first data center, which contains the primary storage system. The third storage system is required for the quorum device. In this solution, AWS Cloud Quorum is used for the quorum device instead of a storage system. It is important to locate the two data centers far enough away from each other to reduce the chance that a single disaster brings down both data centers.

Hitachi Thin Image is a local replication solution for creating space efficient copies. Immediately after a copy (commonly referred to as snapshot) is created, it consists of only pointers to data in the source volume. As new data is written to the source volume or old data is overwritten, the data that existed at the time the snapshot was taken is moved to a snapshot pool. Only after this point do snapshots start consuming space.

By combining the two replication technologies, we have a solution that is always protected with site-to-site replication while providing the flexibility to quickly snap off data copies for analytics, data mining, testing, or other ad hoc requirements.

In addition, the VSP platform can provide storage to cloud-based virtual machines through iSCSI storage ports. Because Hitachi Thin Image can create up to 1,024 snapshots per source volume, pairing it with the very well-known scale out capability of cloud computing makes for a great combination.

Benefits

The following describes the benefits of a disaster recovery solution with Hitachi Global-Active Device with cascading Hitachi Thin Image snapshots

- The solution allows business to resume operations quickly when a disaster brings down the primary data center.
- Continuous server I/O when a failure prevents access to a data volume.
- Server failover and failback without any impact on the storage system.
- Because GAD is an active/active solution, the compute resources moved from the primary storage system to the secondary storage system, or vice-versa, without changing the underlying storage configuration.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- Hitachi Thin Image is space efficient. Therefore, less storage capacity is required.
- Leveraging cloud compute in the form of AWS EC2 virtual machines reduces infrastructure costs. At the same time, it provides benefits such as better scalability and availability.
- Replication is done at the block level by the storage systems and has the following advantages:
 - Dedicated software to move data between data centers is not required.
 - Server CPU processing cycles are not used to move data between data centers.
 - Various operating systems, clustering software programs, and applications are supported.

Key Components

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

- **Storage Systems:** Three storage systems are required. We used a VSP E1090 as the primary storage system and a VSP 5200 as the secondary storage system. A third storage system is required for the quorum device. In this solution, we used AWS Cloud Quorum for the quorum device instead of a storage system.
 - 10 GbE iSCSI storage system Channel Board: 10 GbE iSCSI ports were used to connect to the AWS cloud.
- **Global-Active Device:** GAD is embedded within the microcode, also known as firmware, of the storage system. It does not require any additional hardware. However, it must be activated using a license key on both storage systems.
- **Virtual Storage Machine (VSM):** A special resource group is created (normally) in the secondary storage system with the same model and serial number as the primary storage system. Storage virtualization enables servers to view GAD volumes located on different storage systems as though they were located on the same storage system.
- **Remote Connections:** Physical paths (Fibre Channel) for replicating data between the primary storage system and the secondary storage system.
- **Quorum Disk:** A volume from a third storage system that is used to determine GAD behavior when a storage system or path failure occurs.
- **Copy Pair:** A GAD pair consists of a volume in the primary storage system and a volume in the secondary storage system. Data is mirrored synchronously between the volumes in the copy pair.
 - **Primary Volume (P-VOL):** The source volume that is copied to another volume.
 - **Secondary Volume (S-VOL):** The volume that receives the replicated data from the source volume.
- **Hitachi Thin Image:** HTI is located in the microcode and must be activated using a license key.
- **Hitachi Ops Center Administrator:** Software for provisioning storage.
- **Hitachi Ops Center Protector:** Software for managing replication, data protection, recovery, and retention. It is deployed on a physical server or a virtual machine and can run on the same machine as Administrator.
 - **Intelligent Storage Manager (ISM):** Protector client with Command Control Interface software that acts as a proxy to Block storage devices (through an FC or IP command device).
- **Network Switches:** Cisco Nexus 9000 Series switches were used to connect the two data centers as well as to AWS Direct Connect. The following accessories are required for establishing a WAN between the two sites:
 - **10/25Gbase-LR-S Optics:** Long Range transceivers are required to connect long distances.
 - **Single-Mode Fiber Cables:** Required for long distance communications.
- **Equinix Fabric:** Connected equipment at the Equinix colocation data center that connects to the AWS cloud.
- **AWS Cloud:** Equipment at Equinix was connected to the AWS cloud using a 10 Gbps Direct Connect link. On AWS, a Virtual Private Cloud was created in the region us-west-1.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

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 and a standalone Microsoft SQL Server instance running on EC2 residing in AWS. Records were written to the SQL database (at the near cloud data center) before each failure test and then verified after the failure to ensure data consistency. Snapshots were verified in an EC2 instance in AWS.

Microsoft SQL Server cluster is configured on a Microsoft Windows Failover Cluster running on virtual machines (guests) on VMware ESXi Cluster.

Multipathing for storage volumes are managed at VMware level.

A single volume (multipath-ed) is mapped from the primary storage system and passed through to the virtual machines as Raw Device Mappings (RDM).

Note: Before creating GAD, each volume has a single path, and an alternate path is added to the volume at VMware level automatically when a GAD pair is established. No multipath configuration is done at Microsoft Windows Server 2019 guests.

The mapped (RDM) volume is used for Microsoft SQL Server (clustered) database files and logs.

A second volume (shared) was used for the Microsoft Windows Failover Cluster Quorum Witness and was not replicated.

For disaster recovery or analytics, an EC2 instance (cloud compute) is used from AWS to access the data replicated with GAD or snapshotted with HTI.

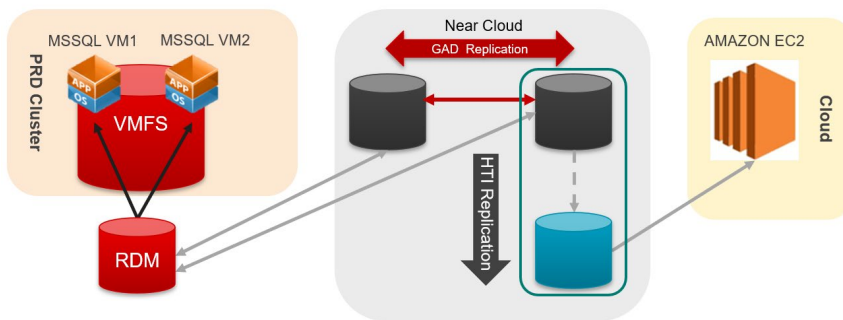


Figure 2: Illustrates the Use Case for Analytics

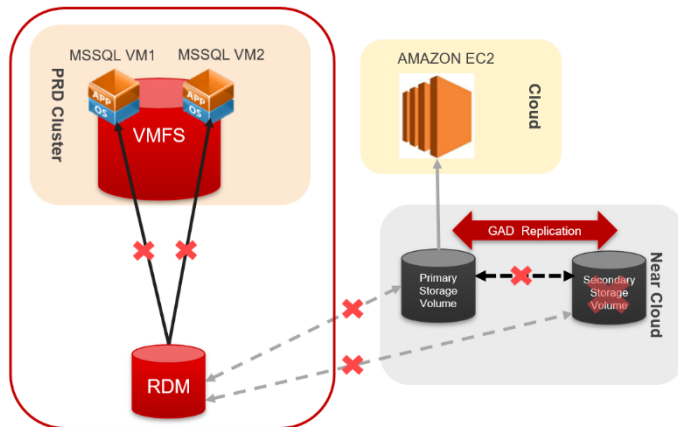


Figure 3: Illustrates Disaster Recovery for Multiple Failures

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

High Level Diagram

Figure 4 shows the test environment used to run the validation.

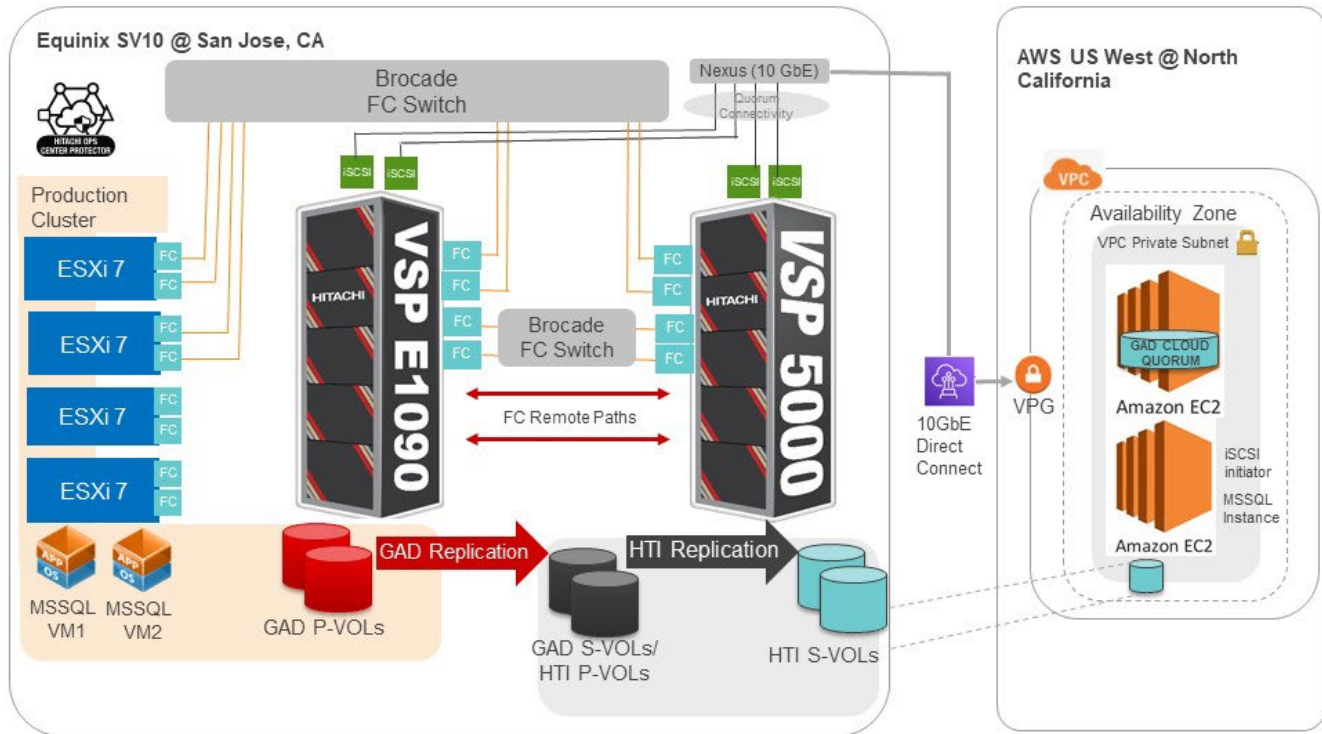


Figure 4: Test Environment

Hardware and Software

Table 1 provides the hardware specifications used in this validation.

Table 1: Hardware Components

Item	Description	Version	Function
Hitachi VSP E1090	1 TB cache (2) 20-core MPUs (3) RAID6 6D-2P parity groups (2) 32 Gbps FC ports (4) 10 GbE iSCSI ports	SVOS RF 9.8 93-06-21-80/00	Primary storage system
Hitachi VSP 5200	2 TB cache (2) 20-core MPUs (4) RAID6 6D-2P parity groups (2) 32 Gbps FC ports (4) 10 GbE iSCSI ports	SVOS RF 9.8 90-08-02-00/01-M079	Secondary 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 secondary vSphere cluster (DR)

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Item	Description	Version	Function
AWS EC2	(2) Intel Xeon E5-2676 @ 2.4 GHz vCPU 4 GB RAM	t2.medium	Compute in the cloud running SQL. Accessed data in VSP 5200.
Brocade G720	Gen 7 Fiber Channel switch	FOS 9.0.1a	Provided FC connectivity between VSP E790 and primary vSphere cluster
Cisco Nexus 93180YC-EX	(48) 1/10/25-Gbps fiber ports (6) 40/100-Gbps QSFP28 ports	NXOS 9.2(3)	Network switch at the primary data center
Cisco Nexus C93180YC-FX	(48) 1/10/25-Gbps fiber ports (6) 40/100-Gbps QSFP28 ports	NXOS 9.3(4)	Network switch at the secondary data center. Serviced Direct Connect to AWS.

Table 2 provides the software specifications used in this validation.

Table 2: Software Components

Item	Version	Function
VMware vSphere	7.0 U2 (17867351)	Hypervisor operating system
VMware vCenter Server Appliance	7.0 U3 (18700403)	Management interface for virtual environment
Hitachi Ops Center Protector	7.4.0.93070-R7.4	Management interface for Hitachi Universal Replicator, Hitachi Thin Image
Hitachi Ops Center Administrator	10.8.0-04.24071	Management interface for provisioning storage to servers
Microsoft Windows Server 2019 Datacenter	Microsoft Windows Server 2019 Datacenter	Guest operating system of SQL virtual machines and EC2 virtual machines
Microsoft SQL Server Enterprise	15.0.2000.5	Database application used to validate data consistency

Test Scenarios

Table 3 lists the test scenarios performed in the validation.

Table 3: Test Scenarios

Test	Description	Success Criteria
1	<p>Prepare VSP E1090 (GAD Primary Storage System):</p> <ol style="list-style-type: none"> Provision (2) 1 TB DP volumes to use as GAD P-VOLs. Provision (1) 200 GB DP volume to use as datastore. Provision (1) 50 GB DP volume to use as quorum for MSFC. <p>Prepare VSP 5200 (GAD Secondary Storage System & HTI):</p> <ol style="list-style-type: none"> Provision (1) 200 GB DP volume to use as datastore. <p>Prepare SQL application on the VSP E1090 storage system:</p> <ol style="list-style-type: none"> Deploy (2) Windows Server 2019 virtual machines on the 200 GB datastore. Map the (2) 1 TB DP volumes as raw device mappings on both Windows 2019 virtual machines. 	Environment is set up per specifications.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test	Description	Success Criteria
	<ol style="list-style-type: none"> 3. Install Microsoft Windows Failover Cluster. 4. Install SQL Server 2019 on the virtual machines. 5. Create a new database on the shared storage system. <p>Prepare SQL application on AWS:</p> <ol style="list-style-type: none"> 1. Deploy (1) Windows Server 2019 AWS EC2 instance. 2. Install SQL Server 2019 on EC2 instance. 3. Create a new database on EBS volume. <p>Prepare GAD Prerequisites:</p> <ol style="list-style-type: none"> 1. Configure FC-based Remote Paths between VSP E1090 and VSP 5200 using Ops Center Administrator. 2. Deploy GAD Cloud Quorum from Amazon Marketplace. 3. Discover GAD Cloud Quorum instance as external iSCSI target on VSP E1090 and VSP 5200. 4. Define GAD quorum on VSP E1090 and VSP 5200 storage systems. 	
2	<p>Create GAD pairs:</p> <ol style="list-style-type: none"> 1. Create GAD pairs using Ops Center Administrator and Protector. 	GAD pairs are created.
3	<p>Create HTI pairs:</p> <ol style="list-style-type: none"> 1. Create 1:1 HTI pairs using Ops Center Protector. 	HTI pairs are created.
4	<p>Access snapshot from AWS EC2 while GAD is in Pair state:</p> <ol style="list-style-type: none"> 1. Write new records to the SQL database. 2. While GAD pairs are in Pair state, suspend (1:1) HTI pairs. 3. Access HTI snapshots from EC2 instance using iSCSI. 4. Verify that the written records on the VM can be accessed from EC2. Write new records in the SQL database to demonstrate database is writable. 5. Detach database; delete snapshots. 	<ul style="list-style-type: none"> • Data created on GAD volumes are present on snapshots. • EC2 instance can write to snapshots without affecting GAD pairs.
5	<p>Multigeneration HTI snapshots:</p> <ol style="list-style-type: none"> 1. While GAD is in Pair state, suspend HTI pairs (1st generation). 2. Map 1st generation snapshots to EC2 instance. Bring copy of SQL database online. 3. Write new records in copy of SQL database. 4. Write new records in SQL database on GAD volumes. 5. Trigger 2nd generation snapshots. 6. Map 2nd generation snapshots to different EC2 instance. 	SQL database on 2 nd generation snapshots contains records written by the production SQL database but not records from the 1 st generation snapshots.
6	<p>Planned outage of the primary storage system (access snapshots from AWS EC2 while GAD is suspended):</p> <ol style="list-style-type: none"> 1. Suspend GAD pairs with "RS" parameter (P-VOLs are blocked and S-VOLs are SSWS). SQL application continues running through S-VOLs with no disruption in I/O. 2. Write new records to the SQL database. 3. Access HTI snapshots from EC2 instance: <ol style="list-style-type: none"> a. Split HTI snapshots. b. Map HTI snapshots to EC2 instance. 	Snapshots can still be taken while GAD roles are swapped.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test	Description	Success Criteria
	<ul style="list-style-type: none"> c. Verify that the written records on the VM can be accessed from AWS EC2. d. Detach database; delete snapshots. <ol style="list-style-type: none"> 4. Resync GAD pairs with “swaps” parameter. 5. Reverse GAD roles to the original direction. 	
7	<p>Unplanned path failure to the primary storage system (paths between ESXi cluster and primary storage system):</p> <ol style="list-style-type: none"> 1. Abruptly disable host connections between ESXi cluster and the primary storage system. Observation: <ul style="list-style-type: none"> • P-VOLs are not accessible from servers while S-VOLs continue to be accessible. • GAD volumes remain in Pair state. • Host I/O continues using alternate path from servers to GAD S-VOLs. SQL database is running from GAD S-VOLs. 2. Write new records to the SQL database. 3. Access HTI snapshots from EC2 instance: <ol style="list-style-type: none"> a. Split HTI snapshots. b. Map HTI snapshots to EC2 instance. c. Write new records to the copy of the SQL database. d. Detach database, delete snapshots. 4. Restore host connections. 	<ul style="list-style-type: none"> • SQL database continues to operate. Records can be written using paths to S-VOLs. • Snapshots can still be taken.
8	<p>Unplanned outage of on-premises servers with data return (simulate failing over to EC2 instance). Note: Downtime, deleting, and recreating GAD is involved.</p> <ol style="list-style-type: none"> 1. Abruptly power off both on-premises servers. Observation: <ul style="list-style-type: none"> • SQL virtual machines and database become offline. • GAD remains in Pair state. 2. Delete the GAD pairs. 3. Map ex-P-VOLs to EC2 instance. 4. Write records in SQL database running on EC2 instance. 5. Detach database. Un-map ex-P-VOLs from EC2 instance. 6. Re-create GAD pairs. 7. Restore on-premises servers. 8. Bring production SQL database online. 9. Verify that the data written in EC2 is present. 	<ul style="list-style-type: none"> • In the event of on-premises compute failure, operations can be restored using cloud compute. • Data written from EC2 instance to SQL server is verified when on-premises compute is restored.

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. The testing performed in this validation is crash consistent. Note: Protector supports application consistent backup for MS SQL with GAD, but it doesn't support iSCSI hosts. Because iSCSI is used to map snapshot volumes in an AWS EC2 instance, the application consistent backup is not being used in this solution. 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.
- Deploy Protector on a Windows Failover Cluster to protect it from an application failure or corruption. For customers without Windows Failover Cluster software, we recommended running Protector at the secondary data center. This allows you to perform failovers to the secondary storage system if the primary data center goes offline.
- Deploy at least one ISM instance at each site to proxy the storage systems local to that site. This allows Protector to continue managing the storage systems at the remaining data center in the case of a disaster.
- The version of Protector available during testing, version 7.4.0, cannot create GAD remote paths, Quorum disk Configuration, and VSM. Therefore, remote paths, Quorum, and VSM must be created using other tools, such as Ops Center Administrator and Storage Navigator.
- The version of Protector available during testing, version 7.4.0, cannot split the pair with P-VOL blocked (PSUS), and S-VOL writeable (SSWS). Therefore, we require `raidcom` command for creating this scenario.
 - The `raidcom` command is: `pairsplit -g d3_juno_gad -RS`
- Other Protector limitations in version 7.4.0 relevant to this solution:
 - Unable to map replicated volumes to iSCSI host groups. A workaround is to use Administrator to map volumes to iSCSI host groups.
 - Resync snapshot is not possible, which means that HTI pairs cannot be put in PAIR state.

Validation Results

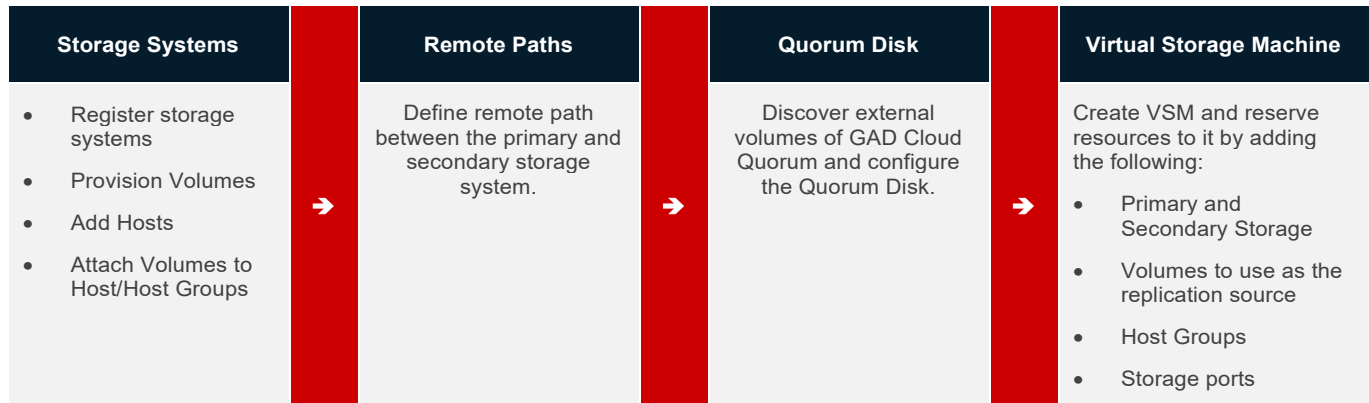
This section contains specific steps and screenshots for each test scenario.

Test 1: Prepare the Environment

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

Ops Center Administrator is used to prepare the environment and Ops Center Protector is used to configure GAD with the cascaded HTI.

The process performed in Ops Center Administrator is as follows:



Prerequisites

The following is a list of prerequisites for preparing the environment:

- Create a physical connection for the GAD environment.
- Configure two guest VMs on the ESXi Cluster using the recommended processor and memory for Windows 2019 with an SQL Server cluster. Resource allocation may vary depending on the database size.
- Install Windows Server 2019 on the guest VMs.
- Install and configure MSFC on the guest VMs.
- Install and configure the Microsoft SQL Server 2019 cluster. During installation, LUN (07:00) is used as a shared volume for the Microsoft SQL server cluster.
- Create two AWS EC2 Instances with Windows Server 2019 and Microsoft SQL Server 2019.

Procedure

To prepare the environment, complete the following steps:

1. LUN Provisioning: Provision one 1TB LUN (LDEV ID: 00:07:00) from the VSP E1090 storage system and mount it as an RDM disk to the Windows Guest (Node #1 and #2).

SELECT	ID	VIRTUAL ID	S/N	VIRTUAL S/N	LABEL	POOL ID ...	ATTRIBUTE	STATUS	CAPACITY SAVI...	TOTAL	USAGE	PROVISIONING...		
+														
<input type="radio"/>	1792	(00:07:...	1792	(00:07:00)	715006	733333	D3_GAD...	0	HA, Thin	Normal	No	1.00 TIB	0%	Attached




Note: You are not required to pre-create the GAD target volumes on the secondary storage system. Protector creates these volumes.

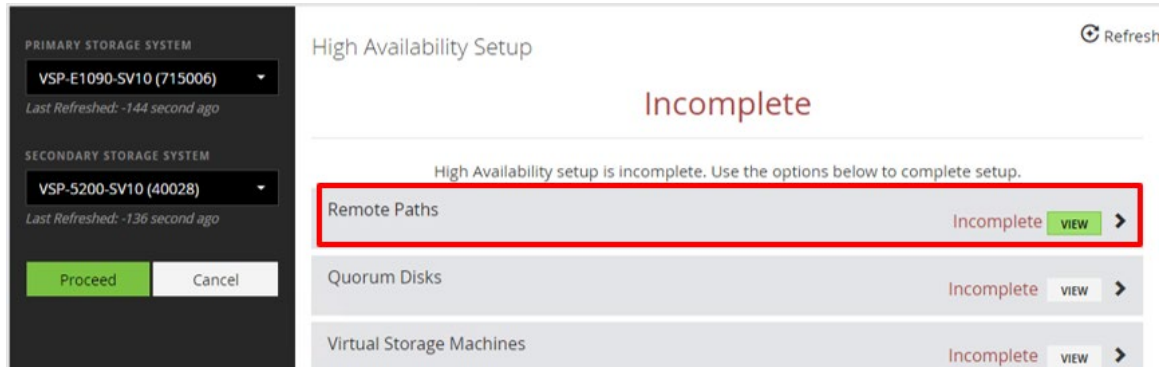
2. Remote Path Configuration: To configure a remote path, log in to Ops Center Administrator and complete the following steps:

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- a. From the **Dashboard**, click **Storage System**, and then click **High availability Setup**.

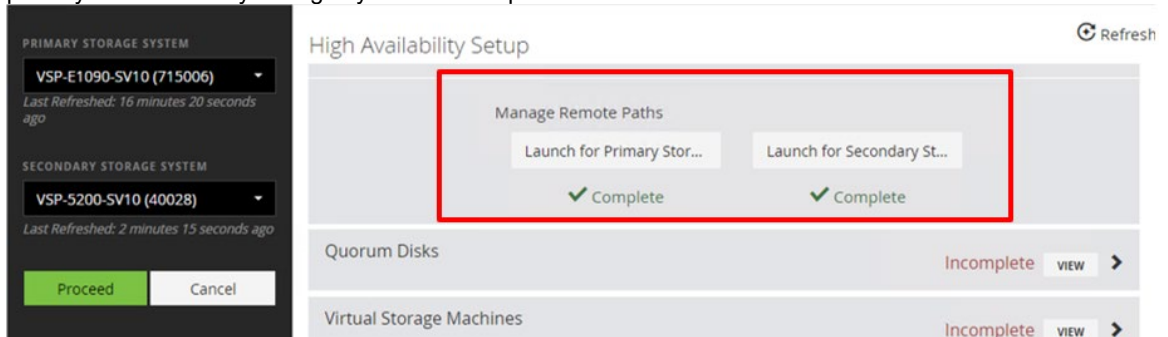


- b. In High Availability Setup, select **Primary** and **Secondary Storage system**.



- c. To launch Storage Navigator to configure remote path settings for the Primary Storage System, navigate to the **Remote Paths** tab and click **Primary Storage System**. Similarly, for configuring remote path settings for the Secondary Storage System, click **Secondary Storage System**.

For configuring remote path settings using Storage Navigator, see the [Hitachi Ops Center Administrator High Availability User Guide](#). Status is changed to **complete** after the remote path configuration between the primary and secondary storage systems is completed.



- d. Protector is used to view the GAD remote paths. The first of the following screenshots shows the paths from the primary storage system to the secondary storage system. The second screenshot shows the paths from the secondary storage system to the primary storage system. The same ports are used in both instances.

Remote Paths

Path Number	Master Port Number	Receiver Port Number	Port Type	Path Status	Master Control Unit	Receiver Control Unit	Control Unit Type
0	CL1-C	CL7-A	Fibre	✓ Normal	-	-	Receiver
1	CL2-C	CL8-B	Fibre	✓ Normal	-	-	Receiver

Remote Paths

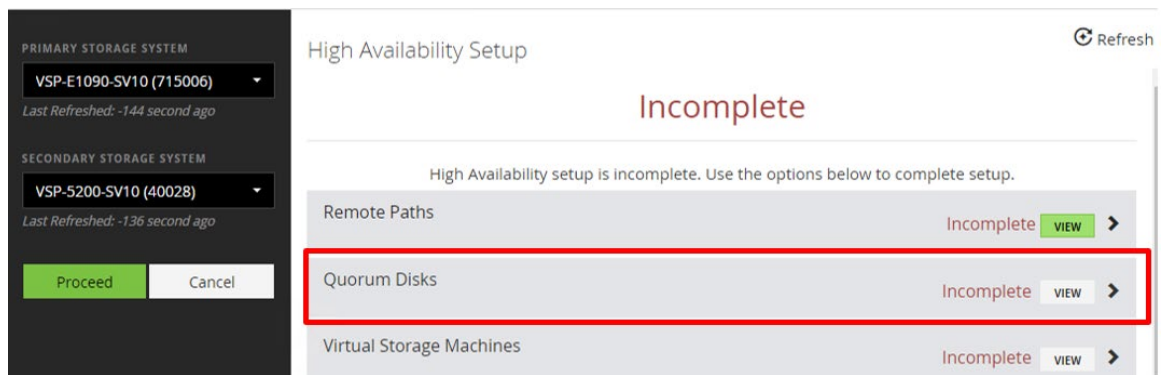
Path Number	Master Port Number	Receiver Port Number	Port Type	Path Status	Master Control Unit	Receiver Control Unit	Control Unit Type
0	CL7-A	CL1-C	Fibre	✓ Normal	-	-	Receiver
1	CL8-B	CL2-C	Fibre	✓ Normal	-	-	Receiver

3. **GAD Quorum Configuration:** In this solution, we used Global-Active Device Cloud Quorum in AWS for quorum volumes. To configure and discover the Quorum volume, complete the following steps:

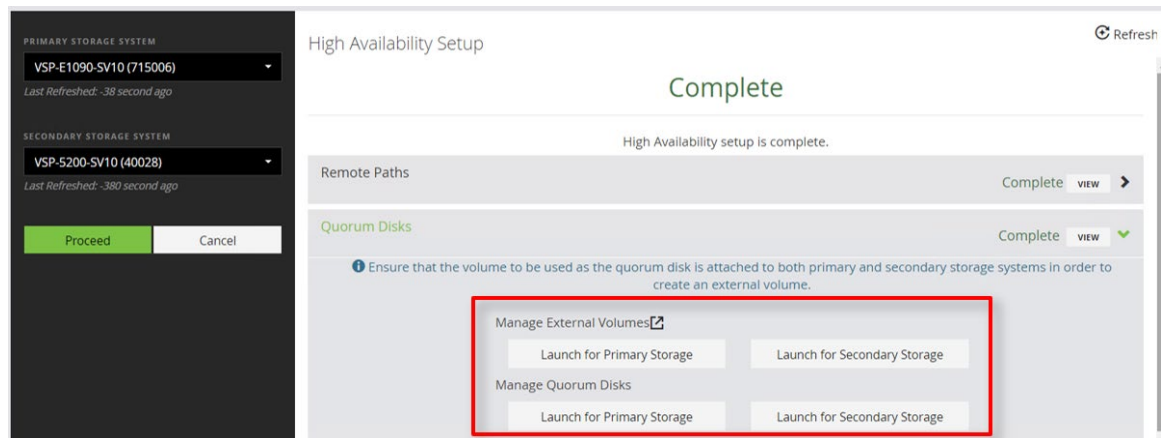
Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- a. GAD cloud quorum in AWS is an Amazon machine image provided by Hitachi Vantara through the AWS Marketplace. It simplifies and enhances GAD by replacing an on-premises quorum with an automatically configured, easy-to-use cloud quorum. In addition to being easier and faster to deploy, a cloud quorum also makes GAD more resilient against outages. To configure GAD Cloud Quorum volumes using an AWS virtual machine (EC2 instance), see [Global-Active Device Cloud Quorum in AWS](#).
- b. The iSCSI Cloud Quorum volume of AWS that was created previously can be discovered using Ops Center Administrator. To discover the volumes and turn them into GAD quorums, do the following:
 - i. Log in to Ops Center Administrator.
 - ii. From the **Dashboard**, click **Storage System**, and then click **High availability Setup**.
 - iii. In the High Availability Setup menu, select **Primary** and **Secondary Storage system**.
 - iv. To launch Storage Navigator to configure Quorum disk settings for the Primary Storage System Quorum Disks, click **Primary Storage System**.
 - v. Similarly, to configure Quorum disk settings for the Secondary Storage System, click **Secondary Storage System**.

For configuring Quorum Disk settings using Storage Navigator, see the [Hitachi Ops Center Administrator High Availability User Guide](#).



The status is changed to **Complete** after the Quorum Disk discovery for Primary and Secondary Storage System is completed.



- c. GAD Quorum Disks are connected to both storage systems. The first screenshot shows the Quorum disk status from the VSP E1090 Primary Storage System. The second screenshot shows the Quorum disk status from the VSP 5200 Secondary Storage System.

Quorum Disks

QUORUM DISKS THAT ARE CONNECTED TO BOTH SELECTED STORAGE SYSTEMS

QUORUM-ID	VOLUME ID	VOLUME LABEL	REMOTE STORAGE S...	STATUS	REMOTE STORAGE S...
0	8	-	715006	NORMAL	NORMAL

Quorum Disks

QUORUM DISKS THAT ARE CONNECTED TO BOTH SELECTED STORAGE SYSTEMS

QUORUM-ID	VOLUME ID	VOLUME LABEL	REMOTE STORAGE S...	STATUS	REMOTE STORAGE S...
0	2	-	40028	NORMAL	NORMAL



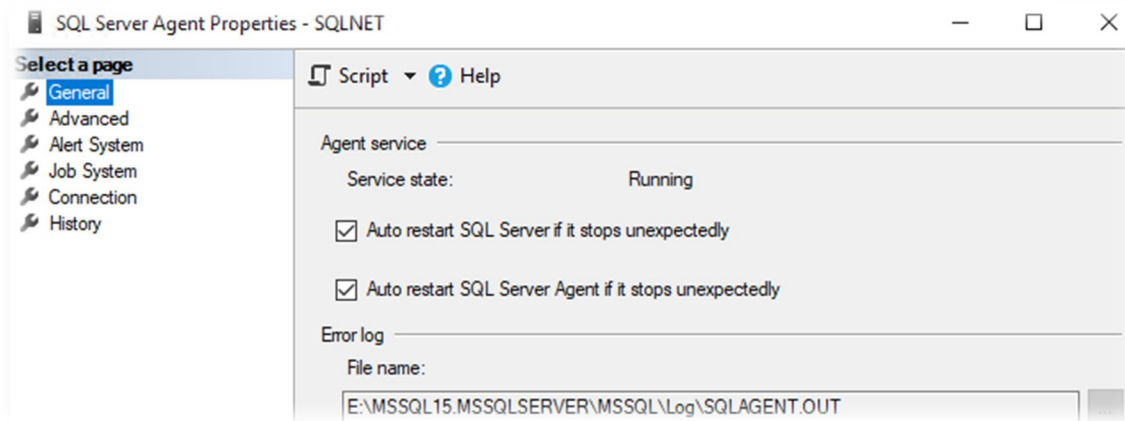
Note: You can configure GAD Quorum disks using the High availability Setup Option of OPS Center Administrator or Storage Navigator.

4. Window Failover Cluster: Install and configure MSFC, and then configure the Microsoft SQL Server 2019 cluster. The following screenshot shows the status of Nodes, Clustered Disks, and Cluster resources of a Microsoft SQL Server in a Windows Failover cluster:

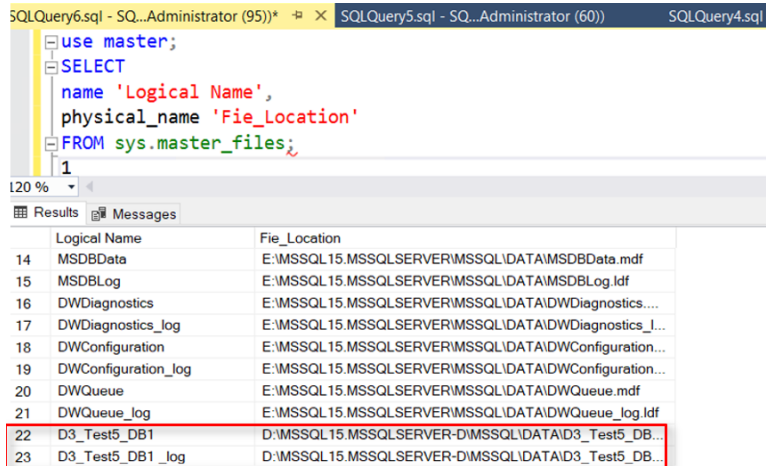
The screenshot displays the Windows Failover Cluster Manager interface. On the left, the tree view shows the cluster 'MSFCSQLGAD.pluto.com' with 'Roles', 'Nodes', 'Storage', 'Networks', and 'Cluster Events' visible. The main pane shows the 'Roles (1)' list with 'SQL Server (MSSQLSERVER)' listed as 'Running' on node '1090SQLNODE1'. Below this, the 'SQL Server (MSSQLSERVER)' role properties are shown, including a list of roles (Analysis Services, SQL Server Analysis Services CEIP, SQL Server CEIP, SQL Server launchpad, SQL Server Polybase Dms, SQL Server Polybase Engine) all in 'Online' status, and 'Storage' resources (Cluster Disk 1, Cluster Disk 2) also in 'Online' status. A 'Server Name' resource is listed as 'SQLNET' in 'Online' status. Under 'Other Resources', 'SQL Server' and 'SQL Server Agent' are also in 'Online' status. A 'SQL Server (MSSQLSERVER) Properties' dialog box is open, showing the 'Failover' tab. The 'Name' is 'SQL Server (MSSQLSERVER)'. The 'Preferred Owners' list contains '1090SQLNODE1' and '5200SQLNODE2'. The 'Priority' is set to 'Medium'. The 'Status' is 'Running' and the 'Node' is '1090SQLNODE1'. Buttons for 'OK', 'Cancel', and 'Apply' are visible at the bottom of the dialog.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

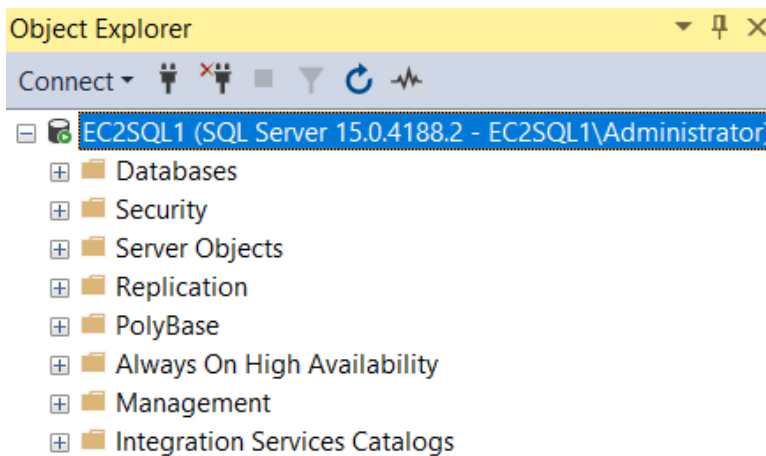
5. Microsoft SQL Server: The following figure shows the status of the SQL Server agent in running state:



6. An SQL query was used to locate the database files registered to the SQL instance. Entries 22 and 23, fenced in red, make up our database. The primary database file, D3_Test5_DB1.mdf, and the database log file, D3_Test5_DB1_log.ldf, are located on the cluster disk.



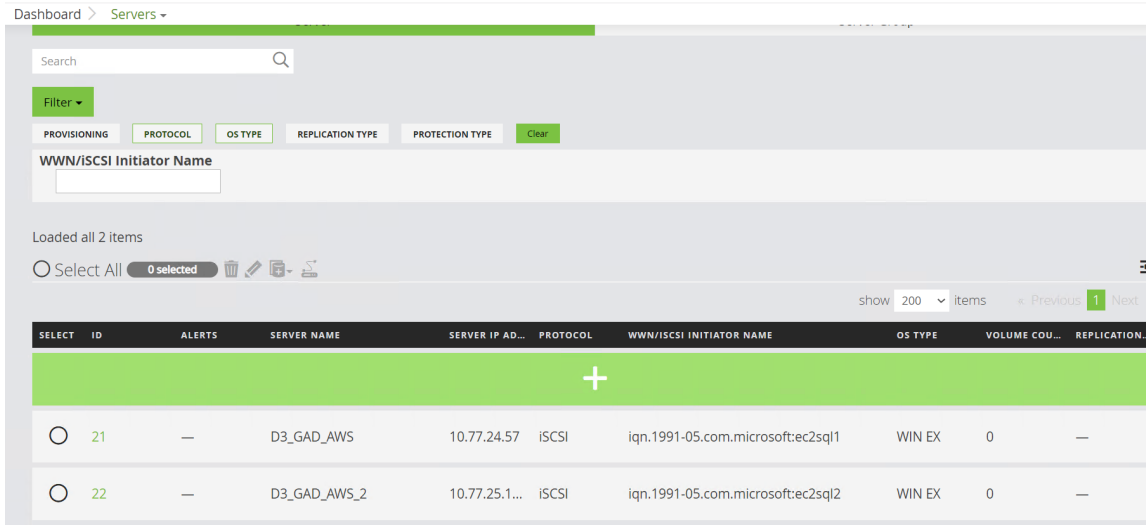
7. AWS EC2 Instance: An AWS EC2 virtual machine running Windows Server 2019 with a standalone SQL instance was used to access data at the Near Cloud site. The following screenshot shows the standalone SQL instance:



8. Add the AWS EC2 instance in Ops Center Administrator: To add the AWS EC2 instance as a server object in Ops Center Administrator, complete the following steps:


Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- a. Locate the iSCSI initiator name of the EC2 instance. You can retrieve it using the PowerShell command: Get-InitiatorPort.
- b. From the Ops Center Administrator Dashboard, click **Servers**.
- c. Click the + (add) symbol under iSCSI Servers.
- d. Enter the EC2 instance name and the iSCSI initiator name (iqn).
- e. Click **Submit**.



Repeat the procedure to add the second EC2 instance as a server object in Ops Center Administrator.

9. Register Ops Center Protector: To register Ops Center Protector from Ops center Administrator, navigate to **Settings** from the home page and click **Ops Center Protector Settings**.

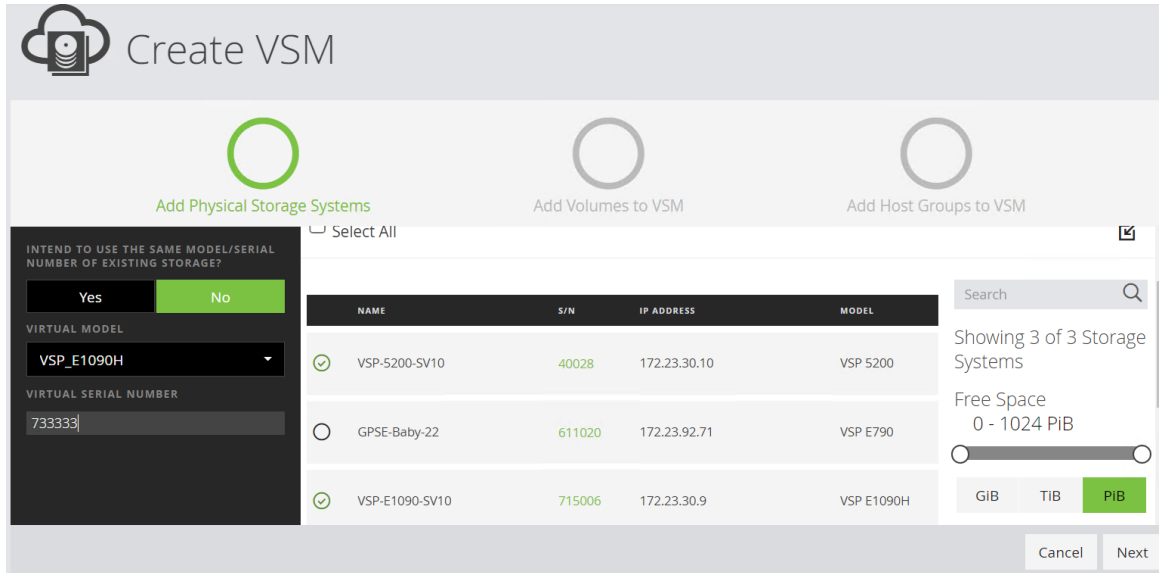


Provide the IP address of the Master Node, Account information, and then click **Submit** on the lower right-hand corner.

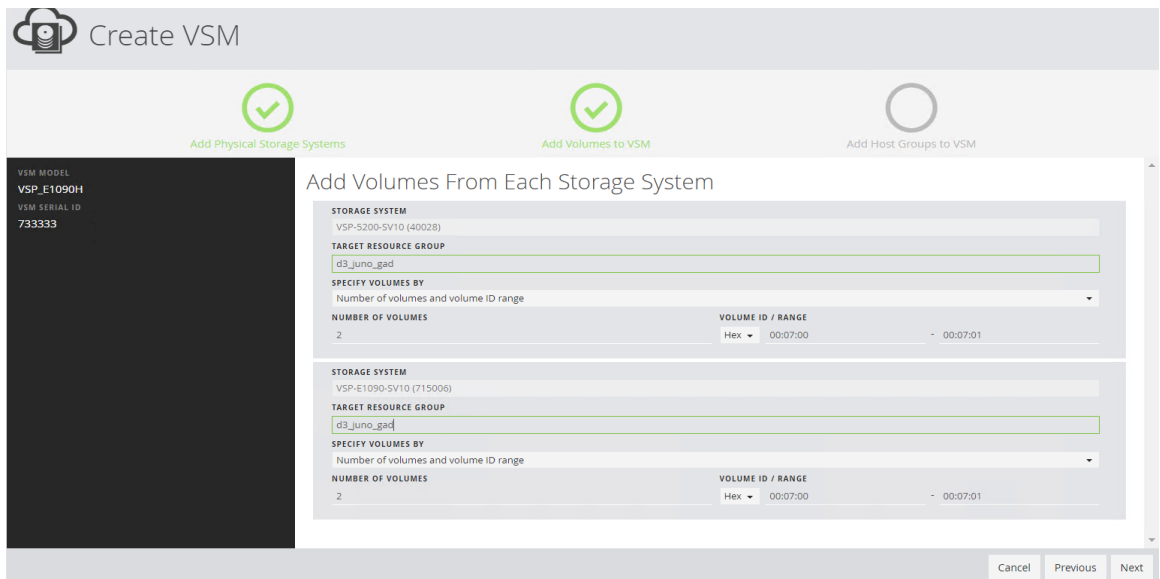


10. Create Virtual Storage Machine (VSM): To create a VSM, log in to Ops Center Administrator and complete the following steps:
 - a. From the dashboard, click **Virtual Storage Machines**. From the Virtual Storage Machines menu, click the + symbol to open the Create VSM menu.
 - b. In the Create VSM menu, select the **Primary Storage System**, **Secondary Storage System**, and **Virtual Storage model**, enter a **virtual serial number**, and click **Next**.

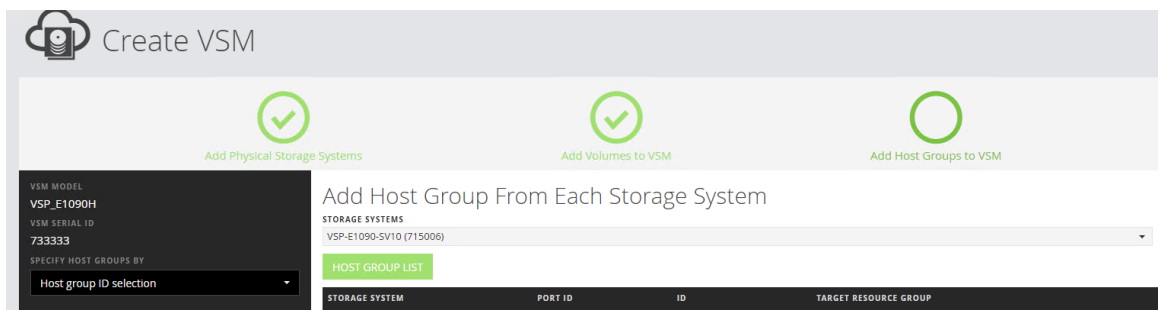
Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image



c. Add LUNs to the VSM from each storage system with the same LDEV ID and click **Next**.



d. Add a Host Group to the VSM and click **Submit**.



The following screenshots show the status of VSM after configuration:

High Availability Setup

PRIMARY STORAGE SYSTEM
VSP-E1090-SV10 (715006)
Last Refreshed: 7 minutes 21 seconds ago

SECONDARY STORAGE SYSTEM
VSP-5200-SV10 (40028)
Last Refreshed: 348 second ago

Proceed
Cancel

High Availability Setup 🔄 Refresh

Complete

High Availability setup is complete.

Remote Paths	Complete view >
Quorum Disks	Complete view >
Virtual Storage Machines	Complete view >

Virtual Storage Machines

VSMS THAT INCLUDE THE SELECTED STORAGE SYSTEMS

VSM ID	VIRTUAL STORAGE SYSTEM ID	VIRTUAL STORAGE MACHINE MODEL
733333-VSPE1090H	733333	VSP E1090H

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test 2: Create GAD Pairs with Cascaded HTI Snapshots

This test case walks you through the process of creating Hitachi Global-Active Device pairs with cascaded Hitachi Thin Image pairs in Protector. The workflow process is as follows:



1. To register the primary and secondary storage systems in Hitachi Ops Center Protector as a Hitachi Block Device, complete the following steps:

- a. Click **Nodes** and click the **+** (plus) symbol.
- b. Select **Storage** and then select **Hitachi Block Device**. Click **Next**.
- c. Enter a node name, add tags (optional), and check the **I confirm...** checkbox. Click **Next**.

d. Click **Next**.

e. From the dropdown list, select the matching proxy node and click **Next**.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- f. Enter a directory for the proxy node to store the metadata and click **Next**.

The screenshot shows the 'Create Node - Hitachi Block Device' configuration page. Under the 'Common Metadata Directory' section, the text input field contains '/root/ISMMetadata' and a 'Browse' button is to its right. A yellow note box below states: 'Note: The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.'

- g. If you are using FC (and already mapped to the proxy host), select the command device. If you are using IP, enter the **IP address** of the SVP/controller. Click **Next**.

The screenshot shows the 'Specify Device' configuration page. The 'Specify by IP or Hostname with a port' radio button is selected. The text input fields contain '172.23.30.11' and '31001'. A note below states: 'To be used as a fall-back option, or if fiber command devices cannot be configured.'

- h. Enter the storage system credentials and click **Next**.

The screenshot shows the 'Specify credentials for device' configuration page. The 'Storage Device Serial Number' field contains '715006'. The 'Username' field contains 'maintenance'. The 'Password' field is masked with dots. A note below states: 'The device account requires the following roles: Storage Administrator (Provisioning, Local Copy, Remote Copy), Security Administrator (View Only, View and Modify) and Support Personnel.'

- i. In the Specify Configuration for Global Replication Reports screen, click **Next**.

- j. Select **All** for LDEV Range and click **Next**.

The screenshot shows the 'Specify LDEV Provisioning Range' configuration page. The 'All' radio button under 'LDEV Range' is selected. The 'Start' and 'End' text input fields both contain '0x00'.

- k. Continue to click **Next** until you reach the summary screen.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- I. Review the information in the summary screen and click **Finish**.

Create Node - Hitachi Block Device

Summary of 'BLOCKSTR_E1090_30_9'

Proxy Node
sv10probenost

Storage Device Serial Number
715006

Username
maintenance

LDEV Provisioning Range
All Available

Configured Command Devices

Type	LDEV ID	IP Address	Port
IP	-	172.23.30.11	31001

Repeat the procedure to register the Hitachi VSP 5200 storage system as a Hitachi Block Device.

Create Node - Hitachi Block Device

Summary of 'BLOCKSTR_VSP5200_30_10'

Proxy Node
sv10probenost

Storage Device Serial Number
540028

Username
opscenter

LDEV Provisioning Range
All Available

Configured Command Devices

Type	LDEV ID	IP Address	Port
IP	-	172.23.30.10	31001

2. Create a Hitachi Block Host node to designate the GAD source volumes.

- a. Click **Nodes** and click the **+** (plus) symbol.
- b. Select **Host** and then select **Hitachi Block Host**. Click **Next**.
- c. Enter a name and tags (optional) and click **Next**.

Create Node - Hitachi Block Host

Specify Node name

Node Name

Must be between 2 and 64 characters, contain only letters, numbers, underscores, hyphens and full stops.

Tags

Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To define a key/value tag, separate the key and value with a colon.

- d. In the Allocate Node to Access Control Resource Group screen, click **Next** without changing the default group.
- e. Select the source storage system and click **Next**.

Create Node - Hitachi Block Host

Select Hitachi Block Device

Hitachi Block Device

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- f. Filter for specific volumes by volume ID or host group. For example, we entered the IDs of a volume that we want to replicate with GAD. Click **Next**.

Create Node - Hitachi Block Host

Specify Logical Devices

Enter Logical Devices using any of the following formats:

- *LDEV_ID* - for a single logical device, e.g., 100, 0x10
- *LDEV_ID-LDEV_ID* - for a logical device range, e.g., 200-299, 0x01-0x0F
- *Host Group ID* - for all logical devices within the host group, e.g., CL1-A-0, CL10-A-0, CL10-A-0xA

Included Logical Devices

0x700

- g. Review the information on the summary screen and click **Finish**.

Create Node - Hitachi Block Host

Summary of 'MCU_E1090'

Hitachi Block Device
BLOCKSTR_E1090_30_9

Logical Devices
0x700

Repeat the procedure to register the VSP 5200 secondary storage system as a Hitachi Block Host.

3. Create a Policy in Hitachi Ops Center Protector.

- a. Click **Policies** and click the **+** (plus) symbol.
- b. Enter a name, description (optional), and tags (optional). Click **Next**.

Create Policy

Specify name and description

Name
GAD_HTI_BLOCK_D3

Description

Tags
Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To define a key/value tag, separate the key and value with a colon.

Add

- c. In the Add One or More Classifications screen, click the **+** (plus) symbol.
- d. Select **Physical > Path > Next**.
- e. Select **Use Hitachi Block Host selections** and click **Apply**.

Create Policy

Specify Hitachi Block Storage classification attributes

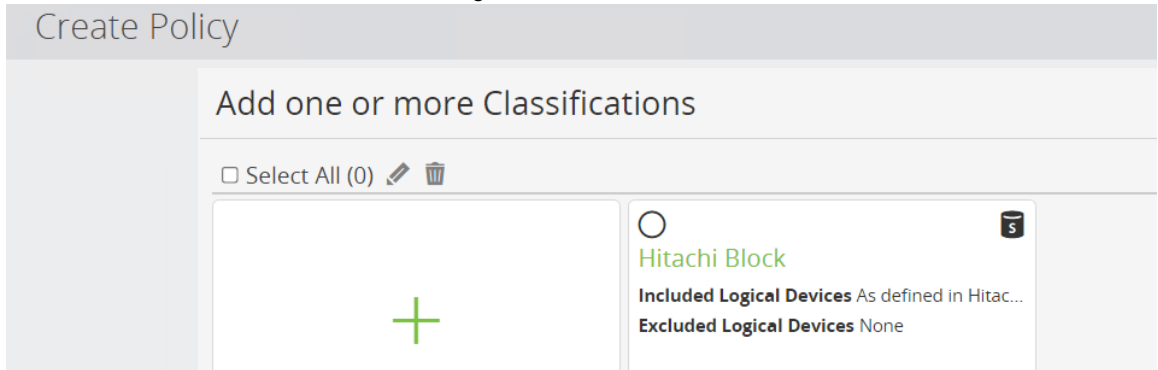
Use Hitachi Block Host selections
 Specify additional selections

Enter values using any of the following formats:

- *Serial/LDEV_ID* - for a single logical device, e.g., 12345/100, 12345/0x10
- *Serial/LDEV_ID-LDEV_ID* - for a logical device range, e.g., 12345/200-299, 12345/0x01-0x0F
- *Serial/Host Group ID* - for all logical devices within the host group, e.g., 12345/CL1-A-0, 12345/CL10-A-0, CL10-A-0xA

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

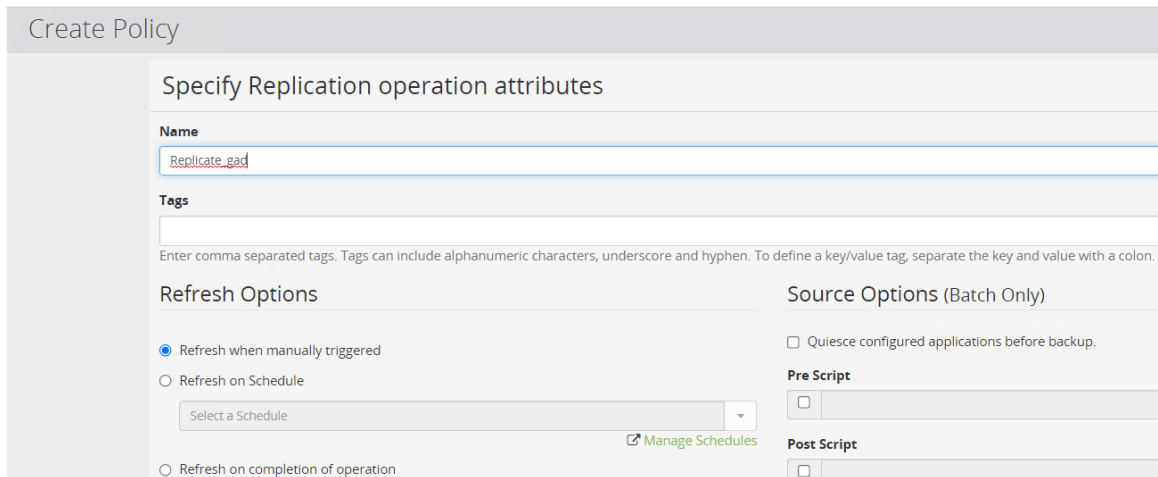
- f. In the next window, as shown in the following screenshot, click **Next**.



- g. In the Add One or More Operations screen, click the + (plus) symbol.
h. Select **Replicate** and click **Next**.

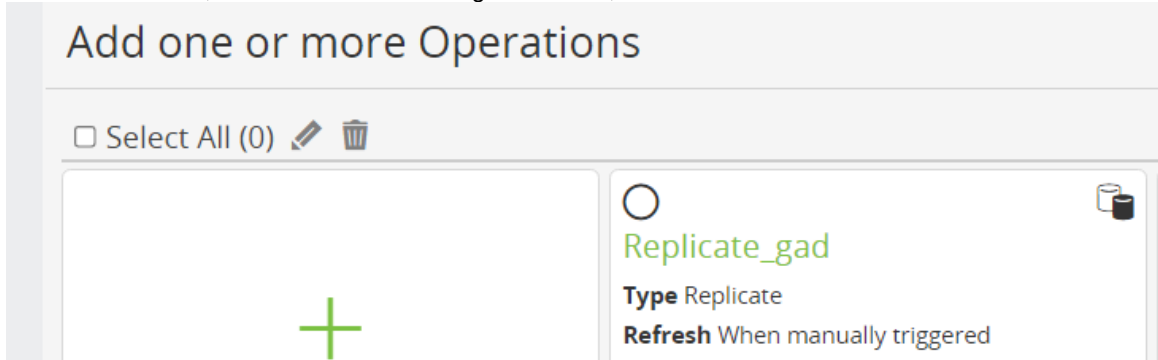


- i. Deselect **Quiesce...** and then click **Apply**.

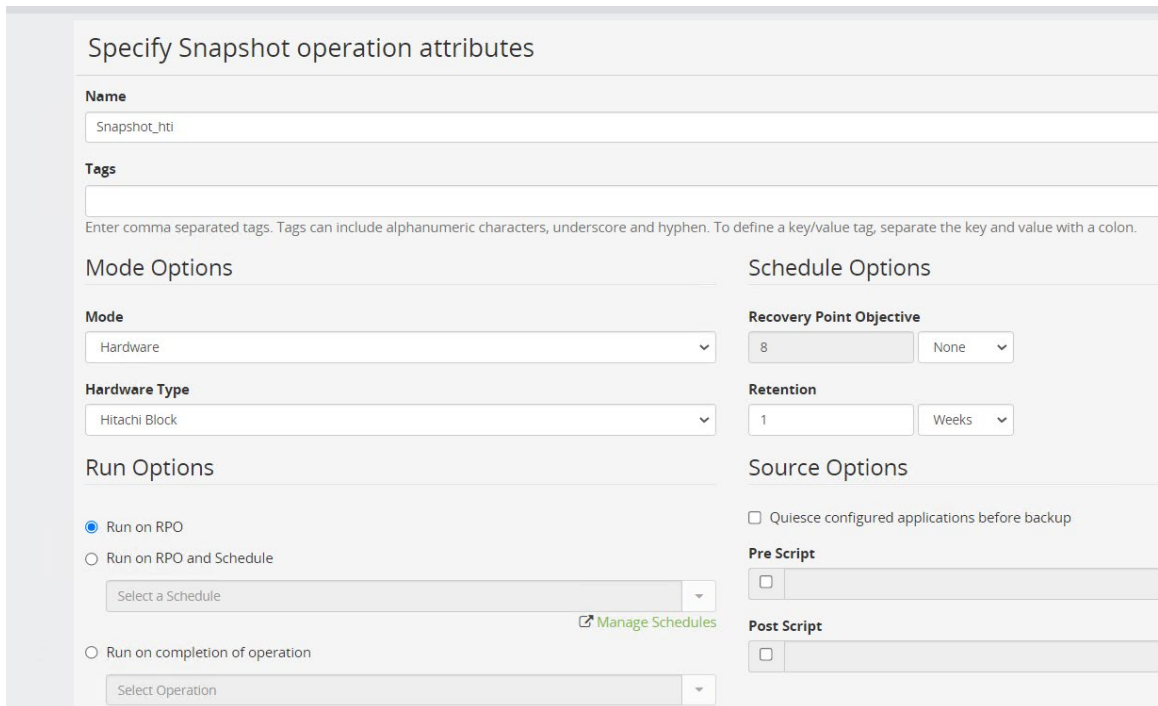


Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- j. In the next window, as shown in the following screenshot, click **Next**.



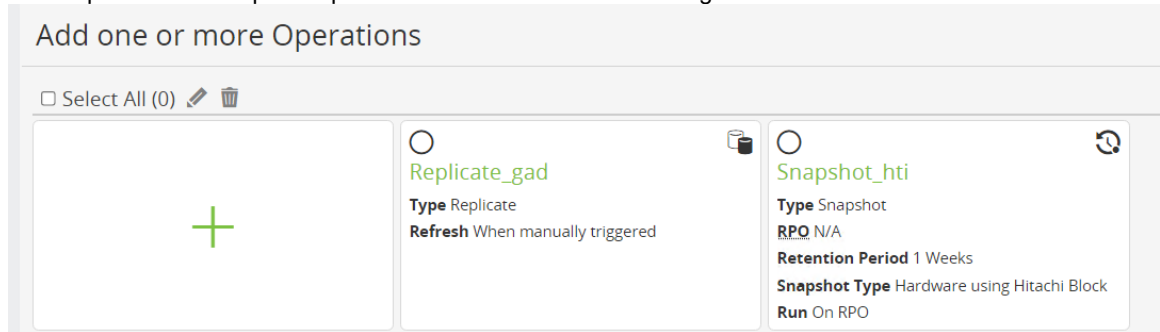
- k. In the Add One or More Operations screen, click the + (plus) symbol again to add a snapshot operation.
- l. Select **Snapshot** and click **Next**.
- m. In the **Specify snapshot operation attributes** screen, set **Recovery Point Objective=None** and **Retention=1 week** (In this scenario, we have set the retention period to 1 week, but it may vary depending on the use case). Deselect **Quiesce...** and then click **Apply**.



Note: **Recovery Point Objective=None** indicates that the snapshot will not generate automatically. It must be triggered manually.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

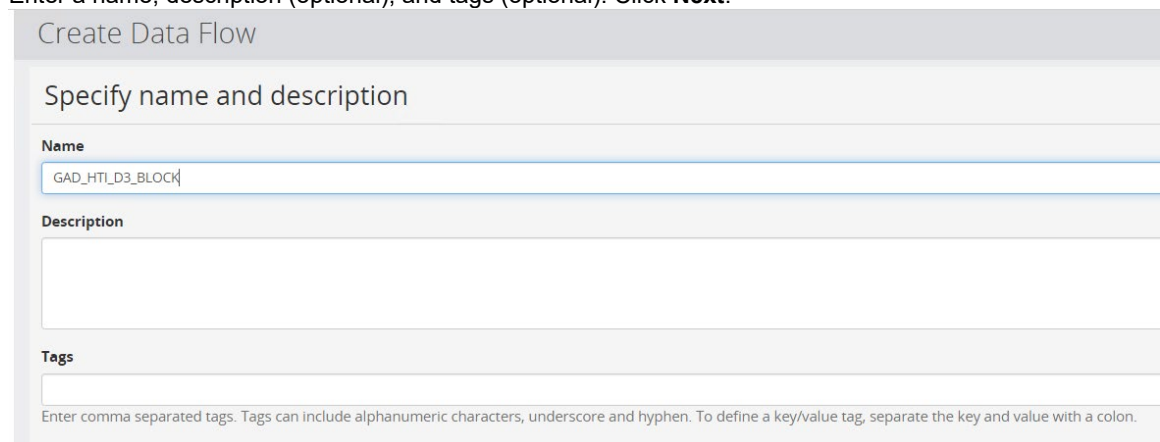
- n. The Replicate and Snapshot operations are shown in the following screenshot. Click **Next**.




- o. Click **Finish**.

- 4. The final step is to create a data flow and activate it to initiate GAD replication. We will manually trigger the snapshot after replication is completed.

- a. Click **Data Flows** and click the + (plus) symbol.
- b. Enter a name, description (optional), and tags (optional). Click **Next**.

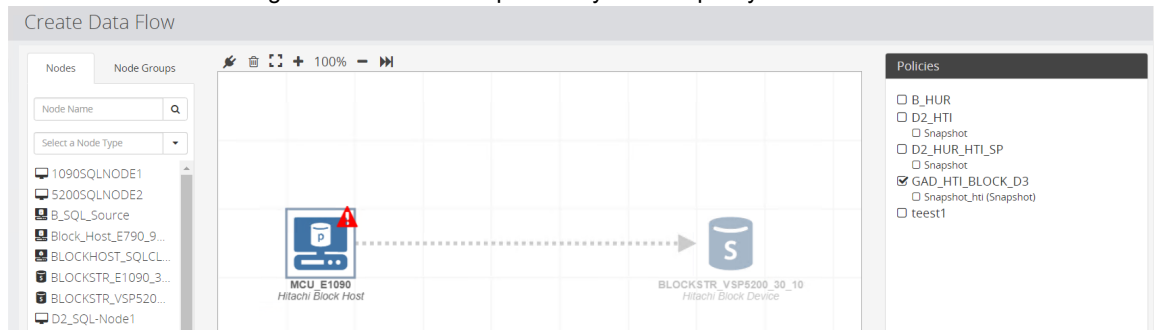


- c. Drag the source **Block Host node** to the whiteboard.
- d. Drag the target **Block Device node** to the whiteboard, over the source Block Host node. This creates an arrow between the two nodes.
- e. Select the **Arrow** and change **Transfer Type** to **Continuous**.

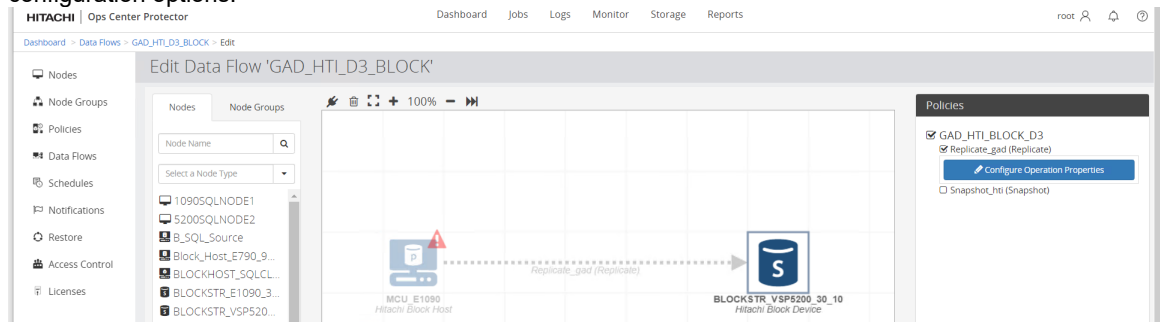


Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- f. Select the source node again and activate the previously created policy from the Policies list.



- g. Select the **Target Node** and then on the right, click **Replicate** first. This brings up a prompt for Replicate configuration options.



- h. Select **Configure new replication** and click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10'

Select Creation Mode

- Configure new replication
- Adopt an existing replication

- i. Select **Active-Active Remote Clone (Global-Active Device)** and click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10'

Select Replication Type

- In-System Clone (*ShadowImage*)
- Snapshot (*Thin Image*)
- Asynchronous Remote Clone (*Universal Replicator*)
- Asynchronous Remote Failover (*Universal Replicator*)
Can only be applied to replication operations over the failover mover.
- Synchronous Remote Clone (*TrueCopy*)
- Active-Active Remote Clone (*Global-Active Device*)

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- j. Select the **Pool** on the secondary storage system where the target volumes will be created, select the **Target Quorum**, and click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10' ×

Configure Active-Active Remote Clone (Global-Active Device)

Configure Replication Settings

Please Note: Changing the replication configuration will cause the replication to be torn down and recreated whether the replication was initialized by Protector or previously adopted.

Pool

dr_pool - 57.52 TB of 82.50 TB free ▼

Target Quorum

Quorum0 ▼

Mirror Unit

Allocate Automatically ▼

Copy Pace

Medium (8) ▼

Use Consistency Group

- k. Select the **Remote Replication Path Group** and click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10'

Select Remote Path Group

- Automatically Selected
- User Selected

Select Source Node

Select a Source Remote Path Group

Only paths that are suitable for this replication type will be shown.

- l. Select the **Resource Group** and click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10' >

Configure Active-Active Remote Clone (Global-Active Device)

Configure Resource Group

- Automatically Selected
Resource Group used by existing SVOLs. Otherwise Resource Group 0.
- User Selected

d3_juno_gad ▼

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- m. Specify the host groups on the secondary storage system where the GAD secondary volumes will be mapped. Click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10' ×

Configure Active-Active Remote Clone (*Global-Active Device*)

Secondary Volume Host Groups

Block replication technologies require all P-VOLs and S-VOLs to have at least one existing LUN path. Options for configuring such paths for the S-VOLs are presented below.

Use Automatically Provisioned Host Group
A LUN path will be created in a placeholder host group for each provisioned S-VOL. If not selected, at least one host group must be specified below.

Enforce LUN ID Matching (fail if primary LUN IDs are not available in the destination host groups)

Optionally specify one or more host groups on the destination storage system. If specified and possible, Protector will create a LUN path from each S-VOL in each of these host groups.

▼ ×
 D3_GAD_5A (CL5-A-2)

▼ ×
 D3_GAD_6B (CL6-B-1)

▼ ×
 Select a Host Group



Note: The version of Protector available during testing (7.4.0) does not map volumes to an iSCSI host group. Volumes must be mapped using Ops Center Administrator.

- n. Specify the naming option for the target volumes. The default option **Match Origin** names the volumes with the same name from the source volumes. Click **Next**.

Replicate_gad configuration on 'BLOCKSTR_VSP5200_30_10' ×

Configure Active-Active Remote Clone (*Global-Active Device*)

Specify Naming Options

Secondary Logical Device Name

Match Origin

Custom

Protector_LDEV_%SECONDARY_LDEV_ID%_%CREATION_DATE%_%CREATION_TIME%

Logical device names are limited to 32 characters, after variable resolution.

Display variables which can be used for the secondary LDEVs' name [▶](#)

- o. Review the information on the summary screen and click **Finish**.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- p. Click **Snapshot** to bring up Snapshot configuration options, select the **snapshot pool**, click **Advanced Configuration**, and then click **Next**.

Snapshot_hti configuration on 'BLOCKSTR_VSP5200_30_10' ×

Configure Snapshot Settings

Storage Node

BLOCKSTR_VSP5200_30_10 ▼

Snapshot Pool

dr_pool (Dynamic Pool) - 57.52 TB of 82.50 TB free ▼

Advanced Configuration

- q. Select the required **Resource Group** and click **Next**.

Snapshot_hti configuration on 'BLOCKSTR_VSP5200_30_10' ×

Configure Resource Group

Automatically Selected
Resource Group used by existing SVOLs, or the Resource Group used by the PVOLs, if all PVOLs are in the same group. Otherwise Resource Group 0.

User Selected

d3_juno_gad ▼

- r. Select all three options: **Consistency group**, **Fully provisioned**, and **Cascade mode**. Select the appropriate cascade pool and click **Next**.

Snapshot_hti configuration on 'BLOCKSTR_VSP5200_30_10' ×

Configure Provisioning Options

Consistency group

Fully provisioned
Creates secondary devices for this snapshot. If not enabled floating devices will be used.

Cascade mode
Allows the creation of snapshots / clones of this snapshot.

Cascade Pool

dr_pool - 57.52 TB of 82.50 TB free ▼

When using fully provisioning devices with cascade mode, a cascade pool may be required.



Note: **Fully provisioned** and **Cascaded mode** allow us to quickly access the HTI secondary volumes as required.

- s. In the Specify Naming Options screen, you can choose the default option **Match Origin** or the naming option for the individual HTI secondary volumes, as well as the naming option for the whole snapshot group. Click **Next**.

Snapshot_hti configuration on 'BLOCKSTR_VSP5200_30_10'

Specify Naming Options

Secondary Logical Device Name

Match Origin

Custom

Protector_LDEV_%SECONDARY_LDEV_ID%_%CREATION_DATE%%CREATION_TIME%

Logical device names are limited to 32 characters, after variable resolution.

Display variables which can be used for the secondary LDEVs' name >

Snapshot Group Name

Automatically Generated

Custom

- t. For Protection Type, select **None**. Click **Next**.

Snapshot_hti configuration on 'BLOCKSTR_VSP5200_30_10'

Configure DRU Options

DRU protection is only available when using fully provisioned devices.

Protection Type

None

Duration of Settings Lock (Days)

0

⚠ DRU settings **cannot be removed** until this duration has expired.

- u. Review the information on the summary screen and click **Finish**.

Snapshot_hti configuration on 'BLOCKSTR_VSP5200_30_10'

Configuration Summary

Type

Differential snapshot (using Thin Image)

Snapshot Pool

dr_pool on BLOCKSTR_VSP5200_30_10

Resource Group

d3_juno_gad on BLOCKSTR_VSP5200_30_10

Provisioning Options

Using consistency group
Using fully provisioned devices
Using cascade mode

Cascade Pool

dr_pool on BLOCKSTR_VSP5200_30_10

Volume Naming

Automatically Generated

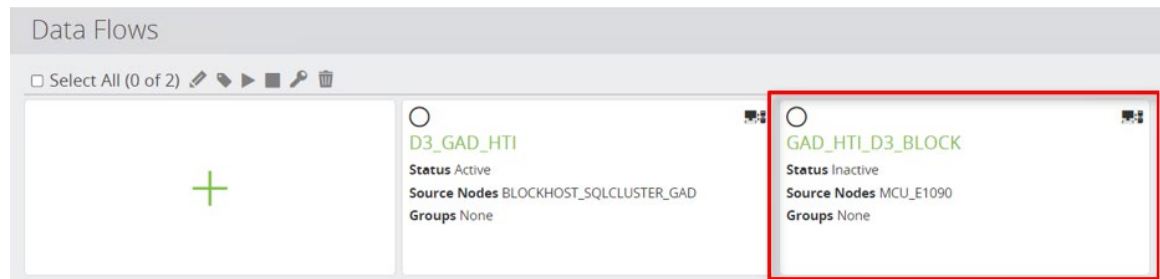
Snapshot Group Naming

Automatically Generated

DRU Protection

Disabled

The newly created data flow is shown as follows:



Multipath status on the VMware ESXi host where Windows VMs are located before the GAD pair is created is shown as follows:

```
[root@SISDS220-21:~] esxcfg-mpath -b
naa.60060e80238235005070823500000700 : HITACHI Fibre Channel Disk
(naa.60060e80238235005070823500000700)
  vmhba2:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:00:10:9b:5b:70:e0 WWPN:
10:00:00:10:9b:5b:70:e0 Target: WWNN: 50:06:0e:80:23:3a:9e:20 WWPN:
50:06:0e:80:23:3a:9e:20
[root@SISDS220-24:~] esxcfg-mpath -b
naa.60060e80238235005070823500000700 : HITACHI Fibre Channel Disk
(naa.60060e80238235005070823500000700)
  vmhba2:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:00:10:9b:5b:6f:39 WWPN:
10:00:00:10:9b:5b:6f:39 Target: WWNN: 50:06:0e:80:23:3a:9e:30 WWPN:
50:06:0e:80:23:3a:9e:30
```

- v. To initiate replication, select the newly created data flow and click **Play**.

- w. Review the information in the prompt and then click **Activate**.

Activate Data Flow(s) ✕

Compiling rules for data flow 'GAD_HTI_D3_BLOCK'
 Rules summary for data flow 'GAD_HTI_D3_BLOCK'
 'MCU_E1090' implements:
 'GAD_HTI_BLOCK_D3'
 'MCU_E1090' sends:
 'GAD_HTI_BLOCK_D3' from the local filesystem to 'BLOCKSTR_VSP5200_30_10'
 'BLOCKSTR_VSP5200_30_10' receives:
 'GAD_HTI_BLOCK_D3' from 'MCU_E1090' for replication
 Compilation of data flow 'GAD_HTI_D3_BLOCK' completed successfully

Compilation succeeded. Ready to activate rules.

Cancel
Activate

- x. To track the progress, click **Monitor**.

GAD Pair status from Ops Center Protector is shown as follows:

HITACHI | Ops Center Protector Dashboard Jobs Logs Monitor **Storage** Reports

Dashboard > Storage > BLOCKSTR_VSP5200_30_10 > Replications and Clones > 05/02/2022 10:08:42 > Pairs

- Nodes
- Node Groups
- Policies
- Data Flows

'BLOCKSTR_VSP5200_30_10' Replication '05/02/2022 10:08:42' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	%	I/O M Mode	ID	Storage	Status	Attribute	%	I/O M Mode	Mirror Unit	Fence Type	Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	L/M	→ 0x0700	540028	PAIR	S-VOL	100%	L/M	0	GAD	NEVER	.

Multipath status on the ESXi host where Windows VMs are located after GAD pair is created is shown as follows:

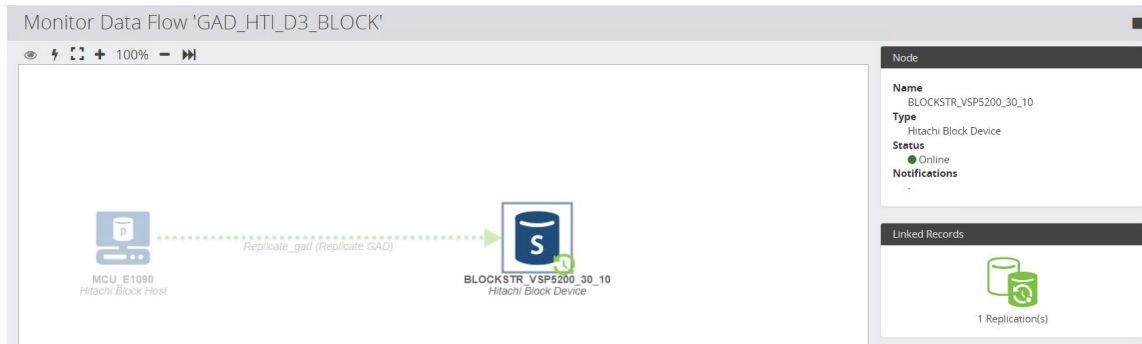
```
[root@SISDS220-21:~] esxcfg-mpath -b |more
naa.60060e80238235005070823500000700 : HITACHI Fibre Channel Disk
(naa.60060e80238235005070823500000700)
  vmhba3:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:00:10:9b:5b:70:e1 WWPN:
10:00:00:10:9b:5b:70:e1 Target: WWNN: 50:06:0e:80:08:9c:5c:40 WWPN:
50:06:0e:80:08:9c:5c:40
  vmhba2:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:00:10:9b:5b:70:e0 WWPN:
10:00:00:10:9b:5b:70:e0 Target: WWNN: 50:06:0e:80:23:3a:9e:20 WWPN:
50:06:0e:80:23:3a:9e:20
[root@SISDS220-24:~] esxcfg-mpath -b
naa.60060e80238235005070823500000700 : HITACHI Fibre Channel Disk
(naa.60060e80238235005070823500000700)
  vmhba3:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:00:10:9b:5b:6f:3a WWPN:
10:00:00:10:9b:5b:6f:3a Target: WWNN: 50:06:0e:80:08:9c:5c:51 WWPN:
50:06:0e:80:08:9c:5c:51
  vmhba2:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:00:10:9b:5b:6f:39 WWPN:
10:00:00:10:9b:5b:6f:39 Target: WWNN: 50:06:0e:80:23:3a:9e:30 WWPN:
50:06:0e:80:23:3a:9e:30
```

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

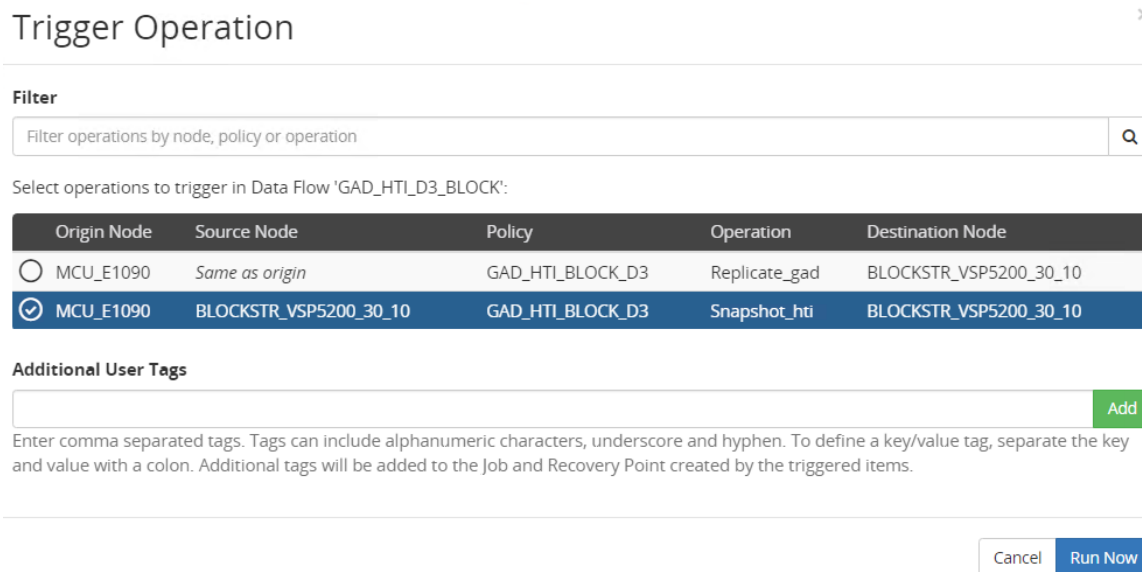
Test 3: Create HTI Pairs

To create Hitachi Thin Image cascade snapshots, complete the following steps:

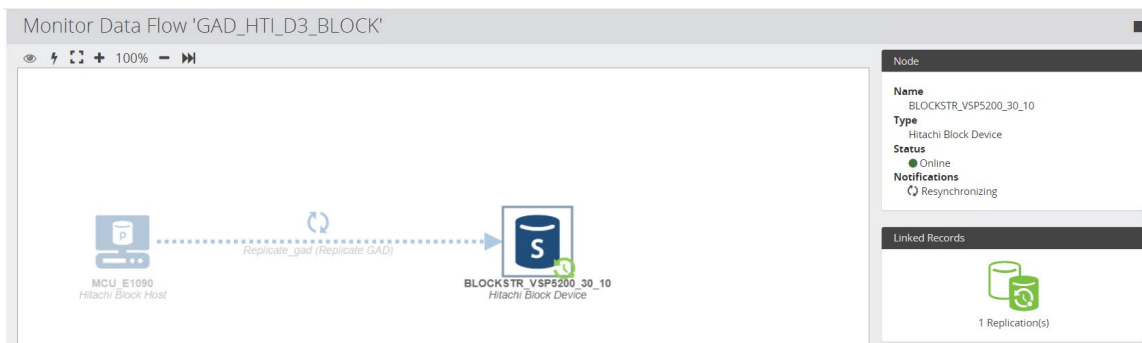
1. Click **Monitor** at the top and click the previously activated data flow. The data flow appears as follows:



2. Select the secondary storage system and click the **Trigger (lightning)** symbol. Click **Run Now**.



Snapshot creation is in progress and is shown as follows:



3. To view the snapshot details after the task has completed, select **Storage > Secondary Storage System > Snapshots**.
 - a. Click the previously created snapshot object.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- b. To display the HTI snapshot details, click **Pairs**.

Dashboard > Storage > BLOCKSTR_VSP5200_30_10 > Snapshots > 05/02/2022 10:19:04 > Pairs

'BLOCKSTR_VSP5200_30_10' Snapshot '05/02/2022 10:19:04' Pairs

Original Primaries							Original Secondaries							Properties			
ID	Storage	Status	Attribute	% M	I/O Mode		ID	Storage	Status	Attribute	% M	I/O Mode	Mirror Unit	Type	Fence Level	Quorum	
0x0700	540028	-	P-VOL	-	-/-		→ 0x0012	540028	-	S-VOL	-	-/-	3	TI	-	-	

- c. Because we chose **Fully provisioned**, the HTI secondary volumes are automatically mapped to random ports. For example, volume 00:12 was mounted to port CL3-A.

HITACHI | Ops Center Administrator Dashboard Jobs Monitoring

Dashboard > Storage Systems > 40028 > Volumes > 18

Storage volume 18 (00:00:12)

Data Protection Alerts

0

VOLUME LABEL
D3_GAD_PVOL

VIRTUAL VOLUME ID
18 (00:00:12)

VIRTUAL STORAGE SYSTEM ID
733333

VIRTUAL MODEL
VSP E1090H

POOL ID
0

POOL NAME
dr_pool

TYPE
Thin

POOL TIER
Platinum

TIERING POLICY
0: All

ATTRIBUTE
Thin

STATUS
NORMAL

PROTECTED

CAPACITY SAVING
No

COMPRESSION ACCELERATION
-

TOTAL
1.00 TiB

BLOCK SIZE
2147483648 Blocks

THIN FREE
1.00 TiB

THIN USED
0 GiB

USAGE
0%

PROVISIONING STATUS
UNMANAGED

ATTACHED SERVERS
-

PORT ID
CL3-A

HOST GROUP NAME/ISCSI TARGET ALIAS
HDIDProvisionedHostGro
up_RG2

LUN
2047 (07FF)

NUMBER OF LUN PATHS
0

ASSIGNED TO MIGRATION
No

ALUA ENABLED
No

T10PI
Disabled

External

TECHNOLOGY	# OF BACKUPS	COMMENT
Snap on Snap	N/A	N/A
SCHEDULE	CONSISTENCY	DATA FLOW
N/A	N/A	

1792 (00:07:00) D3_GAD_PVOL PSUS

COPY GROUP HDID4UA39BPMD...

DATE + TIME 1651488095000

CONSISTENCY ID N/A

18 (00:00:12) D3_GAD_PVOL SSUS

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test 4: Access Snapshots from AWS EC2 while GAD is in Pair state

This test case demonstrates how HTI snapshots can be used to instantly create copies of production data without suspending the GAD replication between sites. The snapshots are then mapped to an AWS EC2 virtual machine so the data can be read.

1. Create a database and populate new records to the SQL database.

Results	Messages	staf_id	first_name	last_name	email	phone	store_id	manager_id
1		196986915	Nickie	Waite	Gutierrez73@example.com	(350) 234-6559	559600483	1020214938
2		411735280	Brad	Craven	Bray162@example.com	(513) 643-0254	1983882757	1228800146
3		548150613	Letha	Wahl	Bolt@nowhere.com	(459) 613-9467	922941124	89729574
4		762898979	Boyd	NULL	NULL	(157) 934-3116	1171267201	363268285
5		868816774	Harlan	Ludwig	qpzre411@example.com	(681) 747-0352	1346404200	1335703973
6		868816775	Alonzo	Skaggs	FelipaAbernathy@nowhere.com	(465) 454-8347	748410784	145232362
7		1083565139	Kraig	Boucher	Morehead@example.com	(730) 955-6764	1410715718	NULL
8		1083565140	Bennie	Flowers	xmhiz1900@nowhere.com	(943) 664-5705	NULL	614412025
9		1518311404	Francis	Conway	Monroe.Finehart@nowhere.com	(282) 510-9326	256308514	34645117
10		1908000233	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	(776) 960-1978	1916920311	1942501253

2. Trigger snapshots using the instructions provided in section [Test 3](#).
3. Map the HTI snapshots to an EC2 virtual machine.
 - a. In Ops Center Administrator, select the **Secondary Storage System**.
 - b. Click **Volumes**.
 - c. Select the snapshot volumes (00:12 in our case) and click **Attach Volumes**.

HITACHI Ops Center Administrator Dashboard Jobs Monitoring

Dashboard > Storage Systems > 40028 > Volumes

Volume ID	Snapshot ID	Storage System	Capacity	Usage	Replication Type	Status
5 (00:00:05)	5 (00:00:05)	40028	40028	DDB_2TB_DR...	0	Thin Normal No 2.00 TIB 0% Attached
6 (00:00:06)	6 (00:00:06)	40028	40028	JNL_0_D2_HU...	0	Thin Normal No 400.00 GiB 1% Unattached
18 (00:00:12)	18 (00:00:12)	40028	733333	D3_GAD_PVOL	0	Thin Normal No 1.00 TIB 0% Unmanaged

- d. In the Attach Volumes window, select the required server and click **Next**.

Attach Volumes

○
 Select Servers

○
 Attach Settings

○
 Create Paths

Select Servers ✖

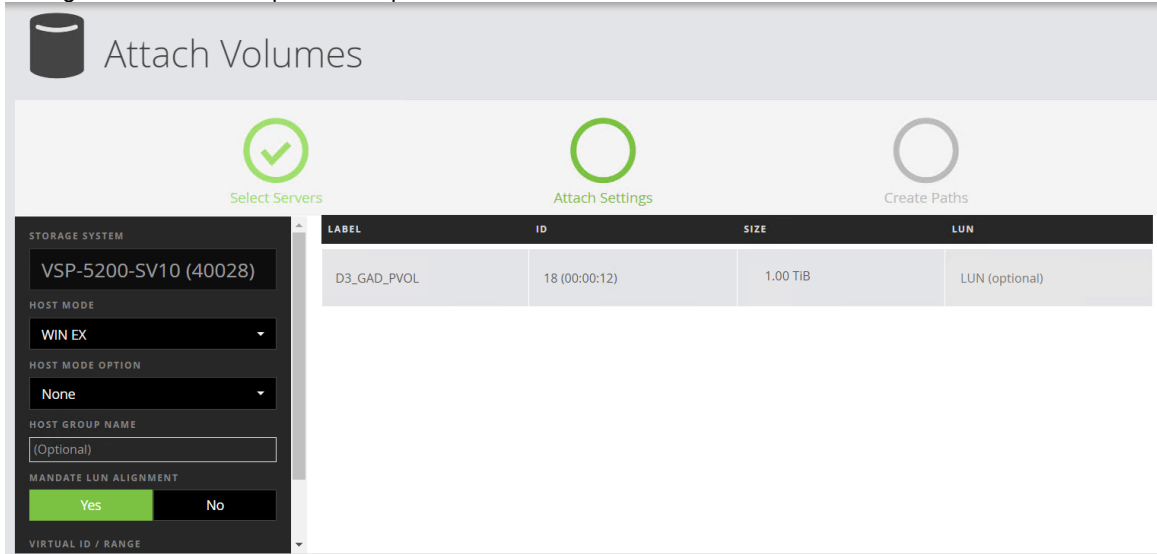
Select All

SERVER ID...	SERVER LABEL	SERVER IP ADDRESS	PROTOCO...	OS TYPE	VOLUME COU...	REPLICATION TYPE
21	D3_GAD_AWS	10.77.24.57	iSCSI	WIN EX	0	-

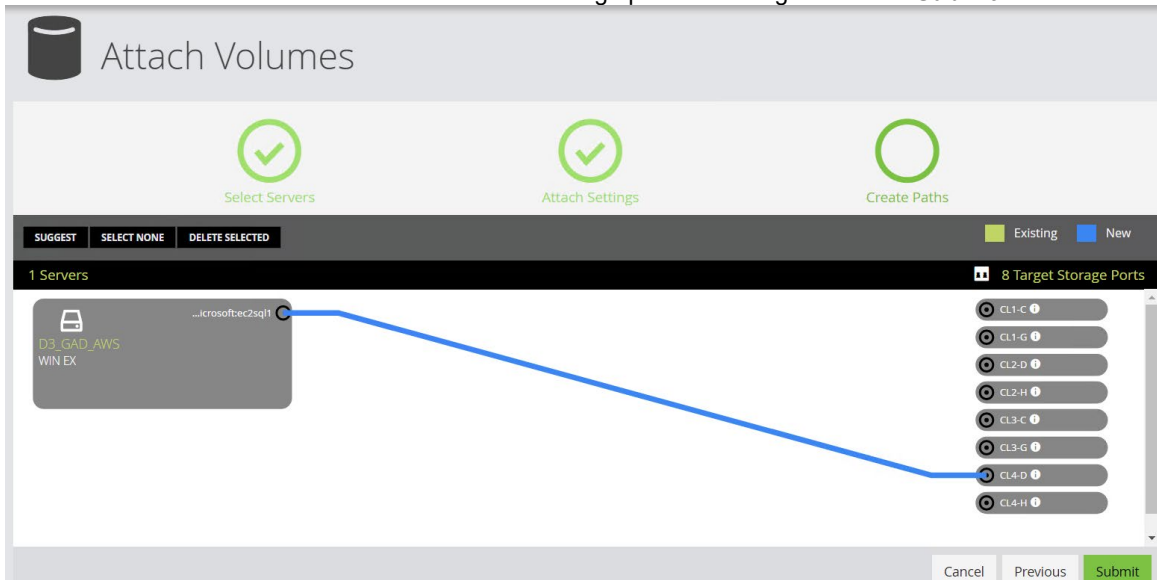
Showing 1 of 14 Servers
 Provisioning
 Provisioned Not Provisioned
 Protocol
 Fibre iSCSI

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- e. Change the Host Mode options if required and click **Next**.



- f. Draw a line between the server on the left and the storage ports on the right and click **Submit**.



This creates a host group with the EC2 virtual machine initiator name on storage port CL4-D and maps the volume to this host group.

○	18 (00:00:12)	18 (00:00:12)	40028	733333	D3_GAD_PV...	0	Thin	Normal	No	1.00 TIB	0%	Attached
---	---------------	---------------	-------	--------	--------------	---	------	--------	----	----------	----	----------

- 4. Discover the iSCSI target and volumes on the EC2 virtual machine.
 - a. Identify the IP address of the storage ports mapped to the EC2 virtual machine.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

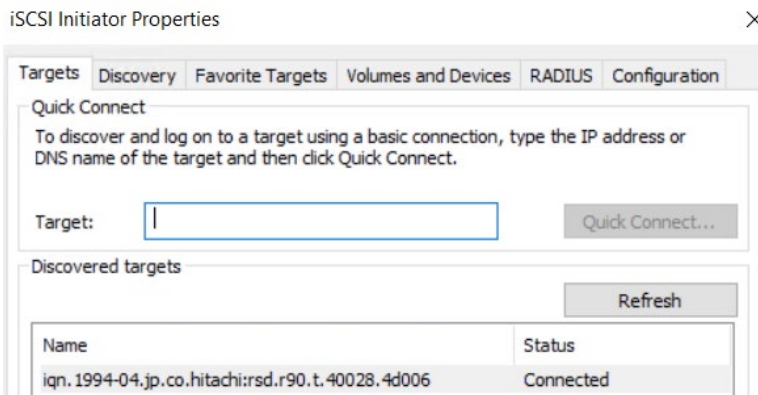
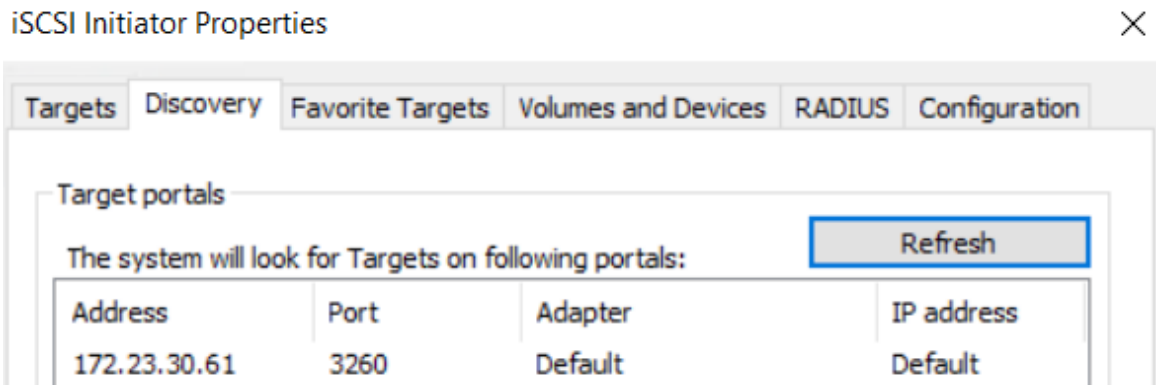
Dashboard > Storage Systems > 40028 > Ports

Fibre


Select All

ALERTS	ID	ISCSI NAME	IPV4	IPV6(GLOBAL) IPV6(LINK LOCAL)	SECURITY	VSM PORT	ATTRIBUTE
	CL1-C	sd.r90.i.089c5c.1c	172.23.30.58	:: fe80::	Enabled	No	Target
	CL4-D	sd.r90.i.089c5c.4d	172.23.30.61	:: fe80::	Enabled	Yes	Target

- b. On the EC2 virtual machine, select **Administrative Tools** > **iSCSI Initiator**. Enter the IP address of the storage port in the **Target** field and click **Quick Connect**.



5. Select **Computer Management** > **Disk Management**. The volume should be listed. Right-click on the volumes and click **Online**.



6. Verify that the previously written records on the VM can be accessed from EC2. The new records in the SQL database demonstrate that the database is writable.

- a. In Microsoft **SQL Server Management Studio** (SSMS), import the database by attaching the database files.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

b. Verify that the previously written records are present.

Results	Messages						
staf_id	first_na...	last_name	email	phone	store_id	manager_id	
1	196986915	Nickie	Waite	Gutierrez73@example.com	(350) 234-65...	559600483	10202149...
2	411735280	Brad	Craven	Bray162@example.com	(513) 643-02...	19838827...	12288001...
3	548150613	Letha	Wahl	Bolt@nowhere.com	(459) 613-94...	922941124	89729574
4	762898979	Boyd	NULL	NULL	(157) 934-31...	11712672...	363268285
5	868816774	Harlan	Ludwig	qpzre411@example.com	(681) 747-03...	13464042...	13357039...
6	868816775	Alonzo	Skaggs	FelipaAbernathy@nowhere.com	(465) 454-83...	748410784	145232362
7	10835651...	Kraig	Boucher	Morehead@example.com	(730) 955-67...	14107157...	NULL
8	10835651...	Bennie	Flowers	xmhz1900@nowhere.com	(943) 664-57...	NULL	614412025
9	15183114...	Francis	Conway	Monroe.Rinehart@nowhere.c...	(282) 510-93...	256308514	34645117
10	19080002...	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	(776) 960-19...	19169203...	19425012...

Query executed successfully.

Data is written to the database. Record updated successfully.

Results	Messages						
staf_id	first_na...	last_name	email	phone	store_id	manager_id	
1	196986915	Nickie	Waite	Gutierrez73@example.com	(350) 234-65...	559600483	10202149...
2	411735280	Brad	Craven	Bray162@example.com	(513) 643-02...	19838827...	12288001...
3	548150613	Letha	Wahl	Bolt@nowhere.com	(459) 613-94...	922941124	89729574
4	555555559	ab	kumar	abc@example.com	(556) 777-14...	NULL	NULL
5	762898979	Boyd	NULL	NULL	(157) 934-31...	11712672...	363268285
6	868816774	Harlan	Ludwig	qpzre411@example.com	(681) 747-03...	13464042...	13357039...
7	868816775	Alonzo	Skaggs	FelipaAbernathy@nowhere.com	(465) 454-83...	748410784	145232362
8	10835651...	Kraig	Boucher	Morehead@example.com	(730) 955-67...	14107157...	NULL
9	10835651...	Bennie	Flowers	xmhz1900@nowhere.com	(943) 664-57...	NULL	614412025
10	15183114...	Francis	Conway	Monroe.Rinehart@nowhere.c...	(282) 510-93...	256308514	34645117
11	19080002...	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	(776) 960-19...	19169203...	19425012...

Query executed successfully.

Note: Updates are happening on the HTI snapshots that are completely independent to the production database. Data written to the snapshots is discarded when the snapshots are deleted.

7. Delete the HTI snapshots.

- a. In **SQL Server Management Studio (SSMS)**, detach the databases.
- b. Open **Computer Management > Disk Management**. Right-click on the drives and click **Offline**.
- c. In Administrator, un-map the volumes from the EC2 virtual machine by selecting the volumes and clicking **Detach**.
- d. In Protector, select the snapshot and click **Delete**.

Confirm Record Deletion ✕

The following LDEVs will be deleted:

ID	Current Name	Historic Name
0x0012	D3_GAD_PVOL	D3_GAD_PVOL

Deleting the selected record(s) from 'BLOCKSTR_VSP5200_30_10' (540028), will delete all associated LDEVs.

Confirmation Word

DELETE

If you are certain you want to perform this operation enter 'DELETE' in the field above.

Cancel
Delete

Test 5: Multigeneration HTI snapshot

This test case demonstrates how to create multiple generations of HTI snapshots from the source. Each snapshot generation is a point-in-time copy of the source data. Each generation is independent of the other generations (although all generations share some amount of data with the source volumes).

1. Create HTI pairs (First Generation). When GAD is in Pair state, generate new records on the GAD volumes by creating a database (D3_Test5_DB1) and a table.

	Roll_no	First_name	Last_name	Class	Total_marks
1	548150612	Letha	Wahl	8	499
2	548150613	Nickie	Waite	7	824
3	762898978	Brad	Craven	6	363
4	762898979	Boyd	NULL	7	458
5	868816773	Francis	Conway	9	791
6	1083565138	Kraig	Boucher	5	964
7	1303563038	Alonzo	Skaggs	7	342
8	1518311403	Bennie	Flowers	6	791
9	1908000232	Enrique	Sizemore	5	110
10	2147483647	Harlan	Ludwig	9	156


2. Mount the first generation HTI snapshots to the AWS EC2 instance.
 - a. Trigger snapshots of GAD Volumes as described in section [Test 3](#).
 - b. Map the snapshot volumes to the AWS EC2 instance as described in section [Test 4.3](#).
 - c. In EC2, bring the drives online and attach the database (D3_Test5_DB1) as described in section [Test 4.5](#).
 - d. Verify that the previously written records are present.

	Roll_no	First_name	Last_name	Class	Total_marks
1	548150612	Letha	Wahl	8	499
2	548150613	Nickie	Waite	7	824
3	762898978	Brad	Craven	6	363
4	762898979	Boyd	NULL	7	458
5	868816773	Francis	Conway	9	791
6	10835651...	Kraig	Boucher	5	964
7	13035630...	Alonzo	Skaggs	7	342
8	15183114...	Bennie	Flowers	6	791
9	19080002...	Enrique	Sizemore	5	110
10	21474836...	Harlan	Ludwig	9	156

Query executed successfully.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

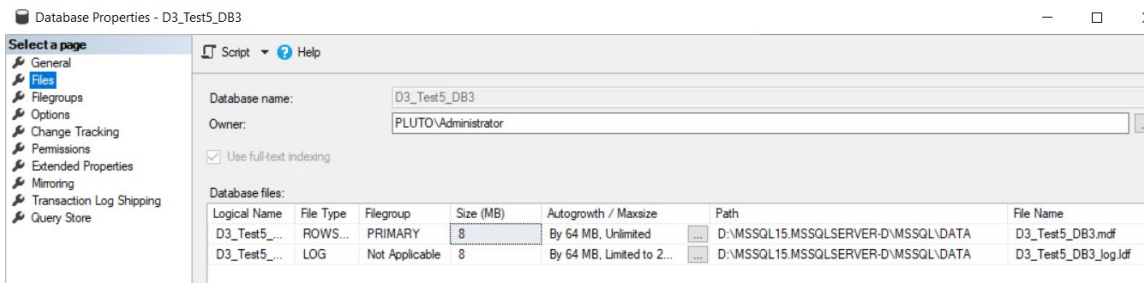
3. Create a new database (D3_Test5_DB2) on the same HTI snapshot drives.



Note: Because the new database (D3_Test15_DB2) was created on the first-generation snapshots, it is not present in subsequent snapshot generations. Snapshot generations are independent, which is part of the design.

4. Create HTI pairs (second generation).

Generate new records on the GAD volume by creating a database named D3_Test5_DB3 and a table. The database is used to validate the second generation of HTI snapshots.



EMP_ID	FIRST_NAME	LAST_NAME	DESIG	DEPT	SAL	
1	08276	Brad	Craven	DY	CG	314700
2	15878	Boyd	Skaggs	R	YH	47072
3	19222	Kraig	Boucher	AAJ	UC	267612
4	23003	Enrique	Sizemore	NWWTKQA	W	33274
5	40093	Alonzo	Peacock	OST	MCEG	131011
6	44963	Nickie	Waite	ERT	I	332456
7	47864	Francis	Conway	R	AB	453349
8	68767	Bennie	Flowers	CQ	ZXXZFB	360675
9	80460	Letha	Wahl	A	H	225097
10	86849	Harlan	Ludwig	CSK	RIL	324325

✓ Query executed successfully.

5. Create the second generation of HTI snapshots as described in section **Test 3**. After the operation is completed, two snapshot relationships are visible.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

HITACHI | Ops Center Protector Dashboard Jobs Logs Monitor **Storage** Reports

Dashboard > Storage > BLOCKSTR_VSP5200_30_10 > Snapshots

'BLOCKSTR_VSP5200_30_10' Snapshots

Select All (0 of 2)
 ↺ ↻ ↻ ↻ ↻ ↻

<div style="text-align: center; border-bottom: 1px solid #ccc; border-radius: 5px; width: 20px; margin: 0 auto;">○</div> <p style="text-align: center; color: green; font-weight: bold;">05/03/2022 10:56:13</p> <p>Type Static Thin Snapshot Data Origin MCU_E1090 Application Hitachi Block Host Policy GAD_HTI_BLOCK_D3 (Snapshot_hti) Expiry Date 05/03/2023 10:56:13</p>	<div style="text-align: center; border-bottom: 1px solid #ccc; border-radius: 5px; width: 20px; margin: 0 auto;">○</div> <p style="text-align: center; color: green; font-weight: bold;">05/03/2022 05:55:36</p> <p>Type Static Thin Snapshot Data Origin MCU_E1090 Application Hitachi Block Host Policy GAD_HTI_BLOCK_D3 (Snapshot_hti) Expiry Date 05/03/2023 05:55:36</p>
--	--

In the second-generation snapshot, the HTI primary volumes remain the same, but the HTI secondary volumes are different (from the first snapshot generation HTI secondary volumes).

'BLOCKSTR_VSP5200_30_10' Snapshot '05/03/2022 10:56:1...

Original Primaries							Original Secondaries							Properties		
ID	Storage	Status	Attribute	%	M	I/O Mode	ID	Storage	Status	Attribute	%	M	I/O Mode	Mirror Unit	Fence Type	Level
0x0700	540028	-	P-VOL	-	-	-/-	→ 0x0014	540028	-	S-VOL	-	-	-/-	4	TI	-

6. Add Undefined Resources to the VSM.

A free host group is required for mapping second generation snapshots to an EC2 instance, and the Host group ID must be added to the VSM before using it to map volumes.

- a. Open the Ops Center Administrator **Dashboard** and select **Virtual Storage Machine > Required VSM > Add Undefined Host Groups to VSM**.
- b. Select the **Storage System** and click **Host Group List**.

Add Undefined Host Groups to VSM

VSM MODEL
VSPE1090H

VSM SERIAL ID
733333

SPECIFY HOST GROUPS BY

Host group ID selection

Add Host Group From Each Storage System

STORAGE SYSTEMS

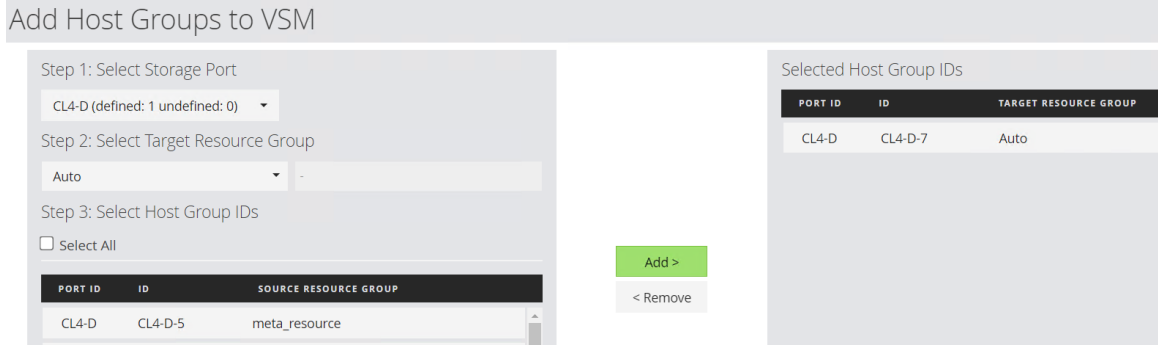
VSP-5200-SV10 (40028)

HOST GROUP LIST

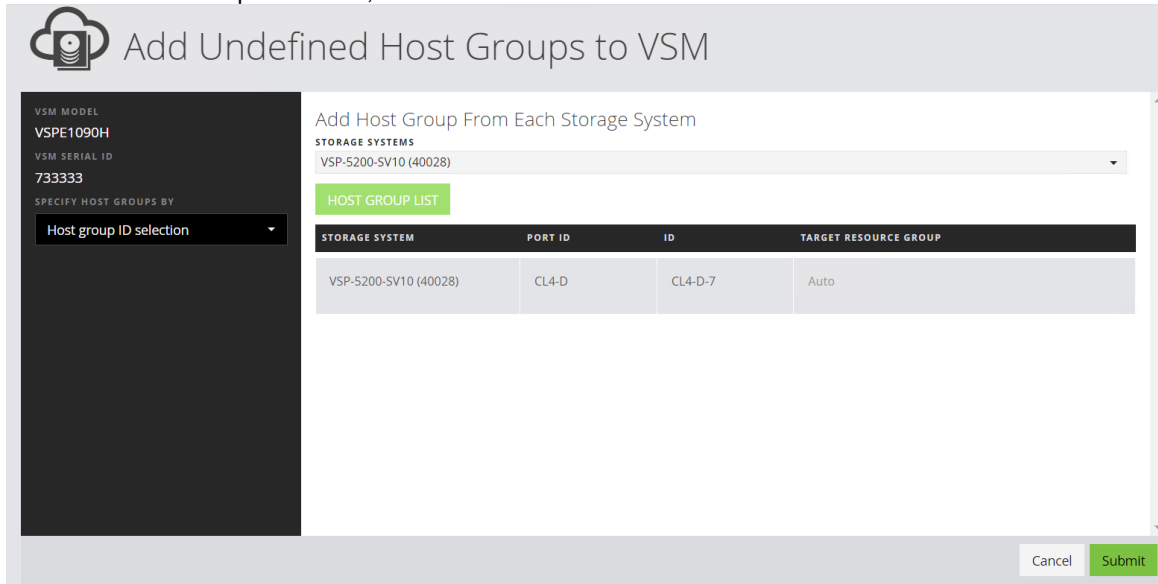
STORAGE SYSTEM	PORT ID	ID	TARGET RESOURCE GROUP

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

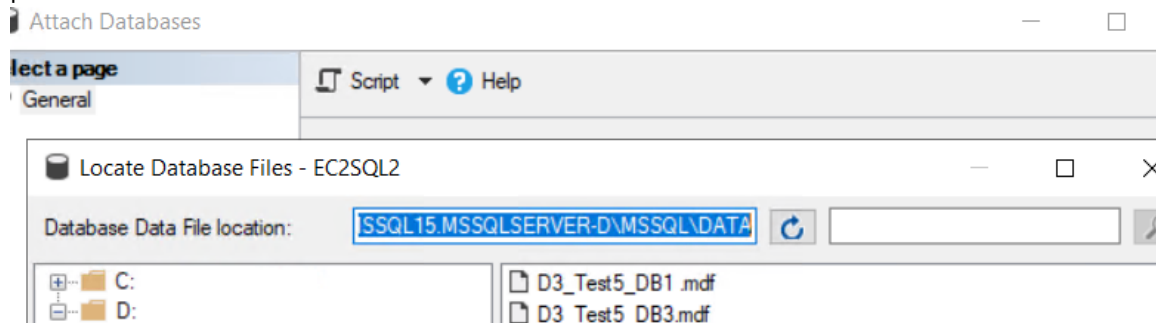
- c. Select and add the required Host Group ID and click **Update**.



- d. To add the Host Group ID to VSM, click **Submit**.



- 7. Add the second EC2 instance in Ops Center Administrator. To add the second AWS EC2 instance as a server object in Ops Center Administrator, see section [Test 1.9](#).
- 8. Map the second generation snapshots to a second EC2 virtual machine, as described in section [Test 4.3](#).
- 9. In EC2, bring the drives online. In Microsoft SQL Server Management Studio (SSMS), import databases D3_Test5_DB1 and D3_Test5_DB3. Notice that the database, D3_Test5_DB2, that was created on the first-generation snapshots, is not present.



Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

10. Verify that the previously written records in D3_test5_DB3 are present.

Results		Messages				
	EMP_ID	FIRST_NAME	LAST_NAME	DESIG	DEPT	SAL
1	08276	Brad	Craven	DY	CG	314700
2	15878	Boyd	Skaggs	R	YH	47072
3	19222	Kraig	Boucher	AAJ	UC	267612
4	23003	Enrique	Sizemore	NWWTK...	W	33274
5	40093	Alonzo	Peacock	OST	MCEG	131011
6	44963	Nickie	Waite	ERT	I	332456
7	47864	Francis	Conway	R	AB	453349
8	68767	Bennie	Flowers	CQ	ZXXZFB	360675
9	80460	Letha	Wahl	A	H	225097
10	86849	Harlan	Ludwig	CSK	RIL	324325

✓ Query executed successfully.

11. Verify that the previously written records in D3_test5_DB1 are present.

Results		Messages			
	Roll_no	First_name	Last_name	Class	Total_marks
1	548150612	Letha	Wahl	8	499
2	548150613	Nickie	Waite	7	824
3	762898978	Brad	Craven	6	363
4	762898979	Boyd	NULL	7	458
5	868816773	Francis	Conway	9	791
6	1083565138	Kraig	Boucher	5	964
7	1303563038	Alonzo	Skaggs	7	342
8	1518311403	Bennie	Flowers	6	791
9	1908000232	Enrique	Sizemore	5	110
10	2147483647	Harlan	Ludwig	9	156

✓ Query executed successfully.

12. Delete the HTI snapshots. When you are ready to delete the HTI snapshots, follow the instructions in section [Test 4.7](#).

Test 6: Planned Outage of the Primary Storage System

This test case demonstrates how GAD can be suspended in a planned outage; for example, you can perform maintenance in the primary data center and business can continue operations on the secondary storage system. It also shows how HTI snapshots can still be created and accessed during an outage. After the outage, data created on the secondary storage system is replicated back to the primary storage system.

1. Initiate the planned outage.

Hitachi Command Control Interface (CCI) is used to conduct the planned outage. To configure CCI, see the [Hitachi Command Control Interface User and Reference Guide](#). Create an HORCM file for the primary and secondary storage system. The following files are the instance definition files of CCI:

- horcm10191.conf: Defines the GAD primary storage volume for GAD pair operation.
- horcm5201.conf: Defines the GAD secondary storage volume for GAD pair operation.

```
[root@svl0probehost etc]# cat horcm10191.conf
HORCM_MON
#ip address      service      poll(10ms)  timeout(10ms)
127.0.0.1       horcm10191  1000       3000

HORCM_CMD
\\.\IPCMD-172.23.30.11-31001

HORCM_LDEV
# DeviceGroup   DeviceName   Serial#   CU:LDEV(LDEV#)   MU#
d3_juno_gad    dev1         715006   07:00            0

HORCM_INST
#dev_group   ip_address   service
d3_juno_gad  127.0.0.1   horcm5201

[root@svl0probehost etc]#
[root@svl0probehost etc]# cat horcm5201.conf
HORCM_MON
#ip address      service      poll(10ms)  timeout(10ms)
127.0.0.1       horcm5201   1000       3000

HORCM_CMD
\\.\IPCMD-172.23.30.10-31001

HORCM_LDEV
# DeviceGroup   DeviceName   Serial#   CU:LDEV(LDEV#)   MU#
d3_juno_gad    dev1         540028   07:00            0

HORCM_INST
#dev_group   ip_address   service
d3_juno_gad  127.0.0.1   horcm10191
[root@svl0probehost etc]#
```

Start CCI instance (horcm10191.conf) on the primary storage system by running the following commands:

```
[root@svl0probehost etc]# horcmstart.sh 10191
starting HORCM inst 10191
HORCM inst 10191 starts successfully.
[root@svl0probehost etc]#
[root@svl0probehost etc]# raidcom -login -I10191
User for Serial#[715006] : maintenance
Password :
[root@svl0probehost etc]#
```

Start CCI instance (horcm5201.conf) on the secondary storage system by running the following commands:

```
[root@svl0probehost etc]# horcmstart.sh 5201
starting HORCM inst 5201
HORCM inst 5201 starts successfully.
[root@svl0probehost etc]#
[root@svl0probehost etc]# raidcom -login -I5201
User for Serial#[540028] : maintenance
```

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

```
Password :
[root@svl0probehost etc]#
```

Pair status before the split:

```
[root@svl0probehost etc]# pairdisplay -g d3_juno_gad -IH10191 -fxec
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,  %,P-LDEV# M CTG JID AP EM
E-Seq# E-LDEV# R/W QM DM P PR
d3_juno_gad dev1(L) (CL3-A-1, 0, 0)715006 700.P-VOL PAIR NEVER , 100 700 - 2 0 2
- - - L/M AA S N D
d3_juno_gad dev1(R) (CL5-A-2, 0, 0)540028 700.S-VOL PAIR NEVER , 100 700 - 2 0 2
- - - L/M AA S N D
[root@svl0probehost etc]#
```

To split pair with PVOL blocked (PSUS), and SVOL writeable (SSWS), run the pairsplit command with RS option.

```
[root@svl0probehost etc]# pairsplit -g d3_juno_gad -RS -IH5201
[root@svl0probehost etc]# pairdisplay -g d3_juno_gad -IH10191 -fxec
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,  %,P-LDEV# M CTG JID AP EM
E-Seq# E-LDEV# R/W QM DM P PR
d3_juno_gad dev1(L) (CL3-A-1, 0, 0)715006 700.P-VOL PSUS NEVER , 100 700 - 2 0 2
- - - B/B - S N D
d3_juno_gad dev1(R) (CL5-A-2, 0, 0)540028 700.S-VOL SSWS NEVER , 100 700 - 2 0 2
- - - L/L - S N D
[root@svl0probehost etc]#
```



Note: Because Ops Center protector does not have the option to split the pair with PVOL blocked (PSUS), and SVOL writeable (SSWS), you must use raidcom to perform this activity.

Pair status from Ops Center Protector:

CKSTR_VSP5200_30_10 > Replications and Clones > 05/03/2022 10:55:58 > Pairs

'BLOCKSTR_VSP5200_30_10' Replication '05/03/2022 10:55:58' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	%	I/O M Mode	ID	Storage	Status	Attribute	%	I/O M Mode	Mirror Unit	Fence Type	Level	Quorum
✓ 0x0700	715006	PSUS	P-VOL	100%	- B/B	→ 0x0700	540028	SSWS	S-VOL	100%	- L/L	0	GAD	NEVER	-

Path Status in VMware ESXi Server:

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore
<input type="checkbox"/>	Local AVAGO Disk (naa.6001636002a953a0240c539bd2a299c7)	0	disk	1.75 TB	local
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E10
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Con

1 EXPORT 1 - 20 of 26 items

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:C0:T0:L0	Active (I/O)	50:06:0e:80:08:9c:5c:40 50:06:0e:80:08:9c:5...	vmhba3:C0:T0:L0	
<input type="radio"/>	vmhba2:C0:T0:L0	Dead	50:06:0e:80:23:3a:9e:20 50:06:0e:80:23:3a:9...	vmhba2:C0:T0:L0	

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datstore
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E10
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Con
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e000000010)	2	disk	2.00 TB	Not Con

1 EXPORT 1 - 20 of 26 items |< < 1 / 2 > >|

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:CO:T0:L0	Active (I/O)	50:06:0e:80:08:9c:5c:51 5...	vmhba3:CO:T0:L0	
<input type="radio"/>	vmhba2:CO:T0:L0	Dead	50:06:0e:80:23:3a:9e:30 5...	vmhba2:CO:T0:L0	

2. Create a database when GAD is in pair split condition and write new records to the SQL database.

In Split condition, database D3_Test6_DB1 can be created when the primary storage system is blocked.

Database Properties - D3_Test6_DB1

Select a page: General, Files, Filegroups, Options, Change Tracking, Permissions, Extended Properties, Mirroring, Transaction Log Shipping, Query Store

Database name: D3_Test6_DB1
Owner: PLUTO\Administrator
 Use full-text indexing

Logical Name	File Type	Filegroup	Size (MB)	Autogrowth / Maxsize	Path	File Name
D3_Test6_...	ROWS...	PRIMARY	8	By 64 MB, Unlimited	D:\MSSQL15.MSSQLSERVER-D\MSSQL\DATA	D3_Test6_DB1.mdf
D3_Test6_...	LOG	Not Applicable	8	By 64 MB, Limited to 2...	D:\MSSQL15.MSSQLSERVER-D\MSSQL\DATA	D3_Test6_DB1_log.ldf

Results Messages

	BOOK_ID	BOOK_NAME	AUTHOR_NAME	ISSUE_DATE	RECEIVE_DATE	STUDENT_NAME	CLASS
1	1001	Letha	Abraham	2021-01-17	2022-01-04	Conrad	6
2	1002	Stephenson	Latonya	2021-04-27	2022-04-15	Damaris	6
3	1003	Sharen	Matthew	2021-03-28	2022-01-08	Drew	8
4	1004	Dillon	Abram	2021-02-10	2022-01-19	Johnston	6
5	1005	Bernardo	Thigpen	2021-01-11	2022-03-19	Jeremy	8
6	1006	Emmelinda	Drury	2021-03-01	2022-03-21	Alease	9
7	1007	Spring	Mcnally	2021-03-13	2022-03-25	Richmond	9
8	1008	Daniell	Omar	2021-03-05	2022-02-05	Frederick	8
9	1009	Jeremy	Vickery	2021-01-11	2022-01-12	Hyman	8
10	1010	Ismael	Broome	2021-01-08	2022-04-11	Trejo	6

Query executed successfully.

3. Create a Block host of the GAD secondary storage volume from the VSP 5200 storage system.

To create a Block host, see section **Test 2.2**.

4. Create a Policy with the Hitachi Block classification and Snapshot operation.

- a. Click **Policies** and click the + (plus) symbol.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- b. Enter a name, description (optional), and tags (optional). Click **Next**.

Create Policy

Specify name and description

Name
D3_SVOL_HT

Description

Tags

Add

- c. In the Add One or More Classifications screen, click the + (plus) symbol.
- d. Select **Physical > Path > Next**.
- e. Select **Use Hitachi Block Host selections** and click **Apply**.

Create Policy

Specify Hitachi Block Storage classification attributes

Use Hitachi Block Host selections

Specify additional selections

- f. The previous selection displays as shown in the following screenshot. Click **Next**.

Create Policy

Add one or more Classifications

Select All (0) [edit] [trash]

Hitachi Block [3]

Included Logical Devices As defined in Hitachi Block ...

Excluded Logical Devices None

- g. In the Add One or More Operations screen, click the + (plus) symbol.

Create Policy

Add one or more Operations

Select All (0) [edit] [trash]

No policy operations. Click to add an operation.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- h. Select **Snapshot** and click **Next**.

Create Policy

Select Operation

- Access
- Backup
- Mount
- Replicate
- Snapshot**
- Tier

- i. In the Specify snapshot operation attributes screen, set **Recovery Point Objective=None** and **Retention=1 year**. Deselect **Quiesce...** and then click **Apply**.

Create Policy

Specify Snapshot operation attributes

Name: Snapshot

Tags: Add

Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To define a key/value tag, separate the key and value with a colon.

Mode Options

Mode: Hardware

Hardware Type: Hitachi Block

Run Options

Run on RPO

Run on RPO and Schedule

Run on completion of operation

Schedule Options

Recovery Point Objective: 8 None

Retention: 1 Years

Source Options

Quiesce configured applications before backup

Pre Script:

Post Script:

Buttons: Cancel Discard Previous Apply



Note: Because Recovery Point Objective=None indicates snapshot will not generate automatically, it must be triggered manually.

- j. Snapshot operations are displayed as shown in the following screenshot. Click **Finish**.

Create Policy

Add one or more Operations

Select All (0)

Snapshot

Type Snapshot

RPO N/A

Retention Period 1 Years

Snapshot Type Hardware using Hitachi Block

Run On RPO

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

5. Create **Data Flows** with the new policy.
Create a dataflow with the configuration shown in the following screenshots. See section [Test 2.4](#).

Snapshot configuration on 'RCU_VSP5200'

Configuration Summary

Type
Differential snapshot (using Thin Image)

Snapshot Pool
dr_pool on BLOCKSTR_VSP5200_30_10

Resource Group
d3_juno_gad on BLOCKSTR_VSP5200_30_10

Provisioning Options
Using consistency group
Using fully provisioned devices
Using cascade mode

Cascade Pool
dr_pool on BLOCKSTR_VSP5200_30_10

Volume Naming
Automatically Generated

Snapshot Group Naming
Automatically Generated

DRU Protection
Disabled

Cancel Previous **Finish**

The screenshot shows a management interface with a grid of resources. One resource is visible: 'RCU_VSP5200 Hitachi Block Host'. To the right, a 'Policies' sidebar lists:

- B_HUR
- D2_HTI
 - Snapshot
- D2_HUR_HTI_SP
 - Snapshot
- D3_SVOL_HTI
 - Snapshot

 A configuration window for 'D3_SVOL_HTI' is open, displaying:

- Type**: Differential snapshot (using Thin Image)
- Snapshot Pool**: dr_pool on BLOCKSTR_VSP5200_30_10
- Resource Group**: d3_juno_gad on BLOCKSTR_VSP5200_30_10
- Provisioning Options**: Using consistency group, Using fully provisioned devices, Using cascade mode
- Cascade Pool**: dr_pool on BLOCKSTR_VSP5200_30_10
- Volume Naming**: Automatically Generated
- Snapshot Group Naming**: Automatically Generated
- DRU Protection**: Disabled

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

6. **Activate** the dataflow.

Activate Data Flow(s) ✕

Compiling rules for data flow 'D3_HTI_SVOL_Snapshot'
 Rules summary for data flow 'D3_HTI_SVOL_Snapshot'
 'RCU_VSP5200' implements:
 'D3_SVOL_HTI'
 Compilation of data flow 'D3_HTI_SVOL_Snapshot' completed successfully

Compilation succeeded. Ready to activate rules.

Cancel
Activate

7. Create a snapshot.

- a. Navigate to **Monitor** and select the **data flow**.
- b. Select the **block host** and click **Trigger**.

The screenshot shows a window titled 'Monitor Data Flow 'D3_HTI_SVOL_Snapshot''. On the left, there is a node icon for 'RCU_VSP5200 Hitachi Block Host'. On the right, a 'Node' panel displays the following details:

- Name:** RCU_VSP5200
- Type:** Hitachi Block Host
- Status:** Online (indicated by a green dot)
- Notifications:** -

- c. From the Trigger Operation menu, click **Run Now**.

Trigger Operation ✕

Filter

Filter operations by node, policy or operation 🔍

Select operations to trigger for source node 'RCU_VSP5200':

Origin Node	Source Node	Policy	Operation	Destination Node
<input checked="" type="checkbox"/> RCU_VSP5200	Same as origin	D3_SVOL_HTI	Snapshot	RCU_VSP5200

Additional User Tags

Add

Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To define a key/value tag, separate the key and value with a colon. Additional tags will be added to the Job and Recovery Point created by the triggered items.

Cancel
Run Now

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- d. After the snapshots are created, note the HTI snapshot details.

'BLOCKSTR_VSP5200_30_10' Snapshot '05/05/2022 07:12:18' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	% M	I/O Mode	ID	Storage	Status	Attribute	% M	I/O Mode	Mirror Unit	Type	Fence Level	Quorum
0x0700	540028	-	P-VOL	-	-/-	→ 0x0015	540028	-	S-VOL	-	-/-	3	TI	-	-

8. Map the snapshot volume to an EC2 virtual machine as described in section [Test 4.3](#).

HITACHI | Ops Center Administrator

Dashboard > Storage Systems > 40028 > Volumes

Volume ID	Size	Used	File System	Format	Mount Point	Mount Status	Thin	Normal	No	Size	Usage	Attachment
14 (00:00:0E)	40028	40028	-	0	Thin	Normal	No	1.00 GiB	0%	Unattached		
21 (00:00:15)	40028	733333	D3_GAD_PVOL	0	Thin	Normal	No	1.00 TiB	0%	Attached		

9. In the EC2 virtual machine, bring the drive online. In Microsoft SQL Server Management Studio (SSMS), import the database D3_Test6_DB1 by attaching the database files.

Database Properties - D3_Test6_DB1

Select a page: General, Files, Filegroups, Options, Change Tracking, Permissions, Extended Properties, Mirroring, Transaction Log Shipping, Query Store

Script Help

Database name: D3_Test6_DB1

Owner: []

Use full-text indexing

Logical Name	File Type	Filegroup	Size (MB)	Autogrowth / Maxsize	Path	File Name
D3_Test6_D...	ROWS ...	PRIMARY	8	By 64 MB, Unlimited	D:\MSSQL15.MSSQLSERVER-DIMSSQLDATA	D3_Test6_DB1.m...
D3_Test6_D...	LOG	Not Applicable	8	By 64 MB, Limited to ...	D:\MSSQL15.MSSQLSERVER-DIMSSQLDATA	D3_Test6_DB1_lo

10. Verify that the previously written records in D3_test6_DB1 are present.

Results Messages

	BOOK_ID	BOOK_NA...	AUTHOR_NAME	ISSUE_DATE	RECEIVE_DATE	STUDENT_NAME	CLASS
1	1001	Letha	Abraham	2021-01-17	2022-01-04	Conrad	6
2	1002	Stephenson	Latonya	2021-04-27	2022-04-15	Damaris	6
3	1003	Sharen	Matthew	2021-03-28	2022-01-08	Drew	8
4	1004	Dillon	Abram	2021-02-10	2022-01-19	Johnston	6
5	1005	Bernardo	Thigpen	2021-01-11	2022-03-19	Jeremy	8
6	1006	Ermelinda	Drury	2021-03-01	2022-03-21	Alease	9
7	1007	Spring	Mcnally	2021-03-13	2022-03-25	Richmond	9
8	1008	Daniell	Omar	2021-03-05	2022-02-05	Frederick	8
9	1009	Jeremy	Vickery	2021-01-11	2022-01-12	Hyman	8
10	1010	Ismael	Broome	2021-01-08	2022-04-11	Trejo	6

Query executed successfully.

11. Detach the database and take the disk from disk management of the Windows OS offline.
12. Remove the snapshot volume from the Host Group. When you are ready to delete the HTI snapshots, follow the instructions in section [Test 4.7](#).
13. Resync GAD pairs with the swaps option as follows:

```
[root@svl0probest etc]# pairresync -g d3_juno_gad -swaps -IH5201
[root@svl0probest etc]#
[root@svl0probest etc]# pairdisplay -g d3_juno_gad -IH10191 -fxec
Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence, % ,P-LDEV# M CTG JID AP EM
E-Seq# E-LDEV# R/W QM DM P PR
d3_juno_gad dev1(L) (CL3-A-1, 0, 0)715006 700.S-VOL PAIR NEVER, 100 700 - 2 0 2
- - - L/M AA S N D
d3_juno_gad dev1(R) (CL5-A-2, 0, 0)540028 700.P-VOL PAIR NEVER, 100 700 - 2 0 2
- - - L/M AA S N D
[root@svl0probest etc]#
```

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

14. Reverse GAD roles to the original direction by running the following commands:

```
[root@sv10probehost etc]# pairresync -g d3_juno_gad -swaps -IH10191
[root@sv10probehost etc]#
[root@sv10probehost etc]# pairdisplay -g d3_juno_gad -IH10191 -fxec
Group   PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,  %,P-LDEV# M CTG JID AP EM
E-Seq# E-LDEV# R/W QM DM P PR
d3_juno_gad dev1(L) (CL3-A-1, 0, 0)715006 700.P-VOL PAIR NEVER , 100 700 - 2 0 2
- - - L/M AA S N D
d3_juno_gad dev1(R) (CL5-A-2, 0, 0)540028 700.S-VOL PAIR NEVER , 100 700 - 2 0 2
- - - L/M AA S N D
[root@sv10probehost etc]#
```

Path status from the server after GAD pair resync:

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E10
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Con
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000010)	2	disk	2.00 TB	Not Con

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:CO:TO:LO	Active (I/O)	50:06:0e:80:08:9c:5c:51 5...	vmhba3:CO:TO:LO	
<input type="radio"/>	vmhba2:CO:TO:LO	Active (I/O)	50:06:0e:80:23:3a:9e:30 5...	vmhba2:CO:TO:LO	

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore
<input type="checkbox"/>	Local AVAGO Disk (naa.6001636002a953a0240c539bd2a299c7)	0	disk	1.75 TB	locc
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E10
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Con

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:CO:TO:LO	Active (I/O)	50:06:0e:80:08:9c:5c:40 5...	vmhba3:CO:TO:LO	
<input type="radio"/>	vmhba2:CO:TO:LO	Active (I/O)	50:06:0e:80:23:3a:9e:20 5...	vmhba2:CO:TO:LO	

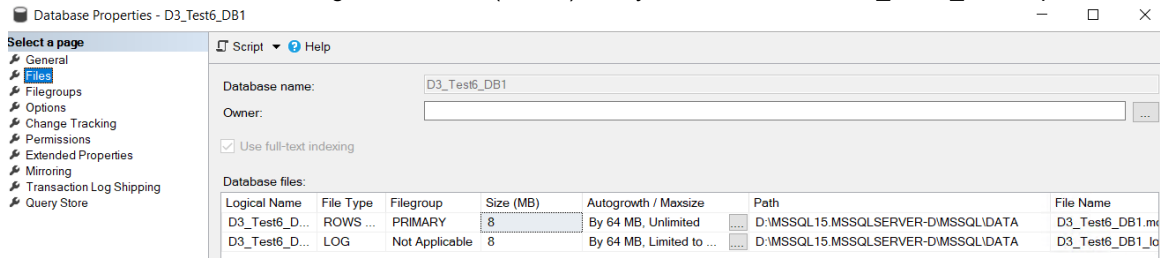
GAD Pair status:

'BLOCKSTR_VSP5200_30_10' Replication '05/03/2022 10:55:58' Pairs

Original Primaries						Original Secondaries						Properties		
ID	Storage	Status	Attribute	%	I/O Mode	ID	Storage	Status	Attribute	%	I/O Mode	Mirror Unit	Fence Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	L/M	→ 0x0700	540028	PAIR	S-VOL	100%	L/M	0	GAD NEVER	-

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

15. In Microsoft SQL Server Management Studio (SSMS), verify that the database D3_Test6_DB1 is present.



16. Verify that the previously written records in D3_test6_DB1 are present.

Results Messages

	BOOK_ID	BOOK_NA...	AUTHOR_NAME	ISSUE_DATE	RECEIVE_DATE	STUDENT_NAME	CLASS
1	1001	Letha	Abraham	2021-01-17	2022-01-04	Conrad	6
2	1002	Stephenson	Latonya	2021-04-27	2022-04-15	Damaris	6
3	1003	Sharen	Matthew	2021-03-28	2022-01-08	Drew	8
4	1004	Dillon	Abram	2021-02-10	2022-01-19	Johnston	6
5	1005	Bernardo	Thigpen	2021-01-11	2022-03-19	Jeremy	8
6	1006	Ermelinda	Drury	2021-03-01	2022-03-21	Alease	9
7	1007	Spring	Mcnally	2021-03-13	2022-03-25	Richmond	9
8	1008	Daniell	Omar	2021-03-05	2022-02-05	Frederick	8
9	1009	Jeremy	Vickery	2021-01-11	2022-01-12	Hyman	8
10	1010	Ismael	Broome	2021-01-08	2022-04-11	Trejo	6

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test 7: Unplanned Path Failure of the Primary Storage System

This test case demonstrates how business can resume operations after a sudden outage at the primary data center. It also shows how HTI snapshots can still be created during the outage and can be accessed from the EC2 instance. In this scenario, the path failure occurred to the primary storage system from both nodes of the MSFC cluster.

1. Status before the outage.

- GAD Pair Status before the outage:

'BLOCKSTR_VSP5200_30_10' Replication '05/03/2022 10:55:58' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	%	I/O M Mode	ID	Storage	Status	Attribute	%	I/O M Mode	Mirror Unit	Fence Type	Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	- L/M	→ 0x0700	540028	PAIR	S-VOL	100%	- L/M	0	GAD	NEVER	-

- Path Status in ESXi Servers:

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datstore
<input type="checkbox"/>	Local AVAGO Disk (naa.6001636002a953a0240c539bd2a299c7)	0	disk	1.75 TB	local
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E10
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Con

1 - 20 of 26 items

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:C0:T0:L0	Active (I/O)	50:06:0e:80:08:9c:5c:40 5...	vmhba3:C0:T0:L0	
<input type="radio"/>	vmhba2:C0:T0:L0	Active (I/O)	50:06:0e:80:23:3a:9e:20 5...	vmhba2:C0:T0:L0	

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datstore
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E10
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Con
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000010)	2	disk	2.00 TB	Not Con

1 - 20 of 26 items

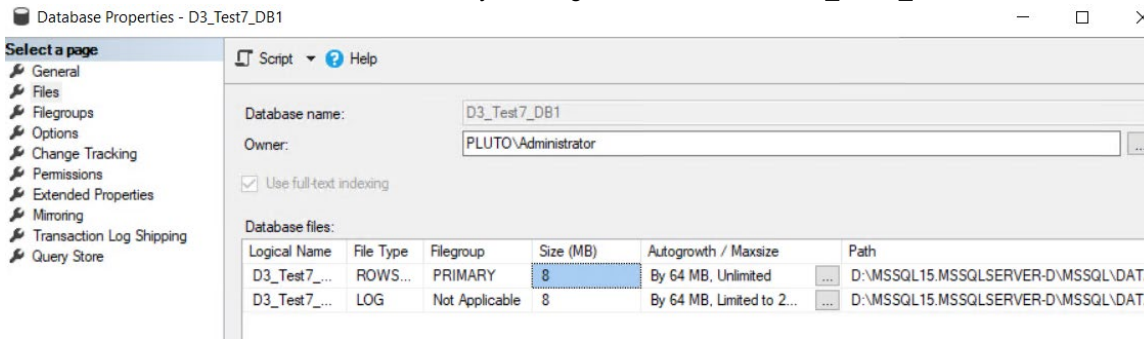
Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:C0:T0:L0	Active (I/O)	50:06:0e:80:08:9c:5c:51 5...	vmhba3:C0:T0:L0	
<input type="radio"/>	vmhba2:C0:T0:L0	Active (I/O)	50:06:0e:80:23:3a:9e:30 5...	vmhba2:C0:T0:L0	

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

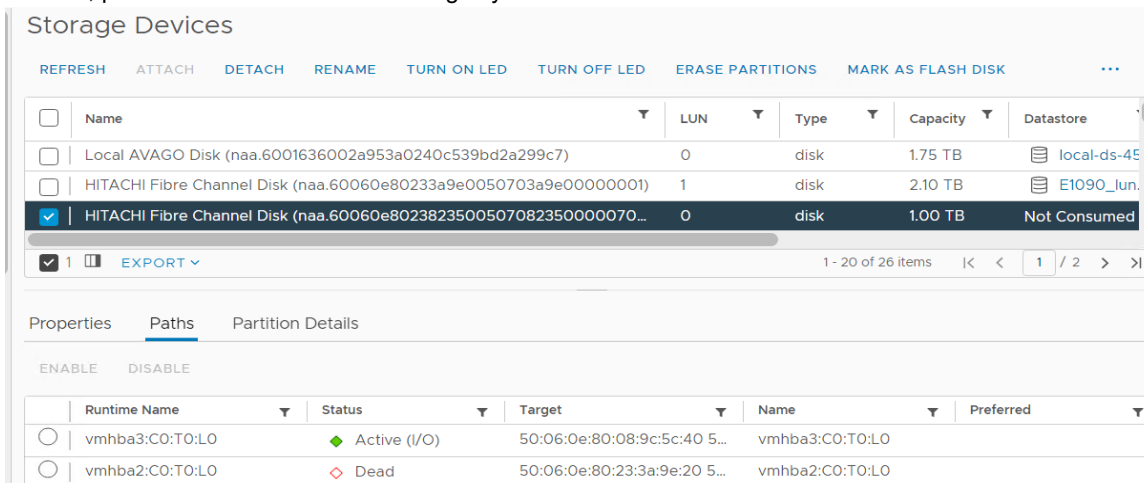
2. Generate new records on the GAD volume by creating a database named D3_Test7_DB1 and a table named Employee.



EMP_CODE	DESIG	FIRST_NAME	LAST_NAME	BASIC	HRA	TOTAL_SAL
2428	XUSZ	Letha	Wahl	262665.6673	67926.56	760424.9626
7940	SX	Kraig	Boucher	212243.8249	86344.994	746919.5346
6574	T	Harlan	Ludwig	333312.5014	81108.9261	656757.0082
5563	Z	Brad	Craven	375154.8406	54632.2129	777731.0444
8587	RR	Enrique	Sizemore	406366.5585	57982.2825	686743.7916
3562	Y	Bennie	Flowers	429458.0114	NULL	723985.527
1160	SZ	Francis	Conway	401142.8497	81826.1271	652786.2531
2155	SSURR	Nickie	Waite	317093.9404	52985.6098	602185.0652
4244	YYU	Boyd	NULL	492931.0614	70726.7528	571781.2206
7961	S	Alonzo	Skaggs	279295.541	68115.19	626341.7442

Query executed successfully.

3. Initiate the outage by disabling FC ports between the VSP E1090 storage system and the Production cluster. In ESXi Servers, paths from the VSP E1090 storage system are in Dead state.



Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

GAD Pair Status after the outage:

'BLOCKSTR_VSP5200_30_10' Replication '05/03/2022 10:55:58' Pairs

Original Primaries							Original Secondaries							Properties			
ID	Storage	Status	Attribute	%	M	I/O Mode	ID	Storage	Status	Attribute	%	M	I/O Mode	Mirror Unit	Type	Fence Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	-	L/M	→ 0x0700	540028	PAIR	S-VOL	100%	-	L/M	0	GAD	NEVER	-

4. Verify that the previously written records are present.

Results Messages

	EMP_CODE	DESIG	FIRST_NAME	LAST_NAME	BASIC	HRA	TOTAL_SAL
1	2428	XUSZ	Letha	Wahl	262665.6673	67926.56	760424.9626
2	7940	SX	Kraig	Boucher	212243.8249	86344.994	746919.5346
3	6574	T	Harlan	Ludwig	333312.5014	81108.9261	656757.0082
4	5563	Z	Brad	Craven	375154.8406	54632.2129	777731.0444
5	8587	RR	Enrique	Sizemore	406366.5585	57982.2825	686743.7916
6	3562	Y	Bennie	Flowers	429458.0114	NULL	723985.527
7	1160	SZ	Francis	Conway	401142.8497	81826.1271	652786.2531
8	2155	SSURR	Nickie	Waite	317093.9404	52985.6098	602185.0652
9	4244	YYU	Boyd	NULL	492931.0614	70726.7528	571781.2206
10	7961	S	Alonzo	Skaggs	279295.541	68115.19	626341.7442

Query executed successfully.

5. Generate new records on the GAD volume by creating a database named D3_Test7_DB2 during outage.

Database Properties - D3_Test7_DB2

select a page

Script Help

Database name: D3_Test7_DB2
Owner: PLUTO\Administrator
Use full-text indexing:

Logical Name	File Type	Filegroup	Size (MB)	Autogrowth / Maxsize	Path	File Name
D3_Test7_...	ROWS...	PRIMARY	8	By 64 MB, Unlimited	D:\MSSQL15.MSSQLSERVER-D\MSSQL\DATA	D3_Test7_DB2.mdf
D3_Test7_...	LOG	Not Applicable	8	By 64 MB, Limited to 2...	D:\MSSQL15.MSSQLSERVER-D\MSSQL\DATA	D3_Test7_DB2_log.ldf

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

EMP_ID	DESIG	FIRST_NAME	LAST_NAME	ADDRESS	AGE	
1	1000	DEAE	Letha	Wahl	461 NW Riddle Hill Lane, Concord, NH, 04357	43
2	1001	CD	Kraig	Boucher	11 Red Lake Hwy, Harrisburg, Pennsylvania, 30254	31
3	1002	D	Harlan	Ludwig	1753 North Church Pkwy, Boise, Idaho, 95705	21
4	1003	E	Brad	Craven	1695 Hidden Deepwood Cir, Superior Bldg, Indiana...	36
5	1004	DC	Enrique	Sizemore	970 Riddle Hill Pkwy, Comcast Building, Albany, Ne...	34
6	1005	C	Bennie	Flowers	1955 Front Road, Madison, WI, 37006	24
7	1006	AD	Francis	Conway	2844 North Woodfort St, Raleigh, NC, 26635	40
8	1007	DBEDE	Nickie	Walte	2989 East Pine Tree Road, Topeka, KS, 16758	52
9	1008	BAD	Boyd	Skaggs	62 Ashwood Pkwy, 3rd Floor, Madison, Wisconsin, ...	NULL
10	1009	A	Alonzo	Peacock	40 Bayview Lane, Topeka, Kansas, 52885	36

Query executed successfully.

6. Create HTI snapshots as described in section [Test 3](#). Pair status of the snapshot is as follows:

'BLOCKSTR_VSP5200_30_10' Snapshot '05/05/2022 12:33:53' Pairs

Original Primaries							Original Secondaries							Properties			
ID	Storage	Status	Attribute	%	M	I/O Mode	ID	Storage	Status	Attribute	%	M	I/O Mode	Mirror Unit	Type	Fence Level	Quorum
0x0700	540028	-	P-VOL	-	-	-/-	→ 0x0016	540028	-	S-VOL	-	-	-/-	3	TI	-	-

7. Map the snapshot volumes to the EC2 instance as described in [Test 4.3](#).

○	22 (00:00:16)	22 (00:00:16)	40028	733333	D3_GAD_PVOL	0	Thin	Normal	No	1.00 TiB	0%	Attached
---	---------------	---------------	-------	--------	-------------	---	------	--------	----	----------	----	----------

8. In the EC2 instance, bring the drives online. In this scenario, we have mounted the drive as drive D.

Disk 0

Basic
65.00 GB
Online

(C:)
65.00 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

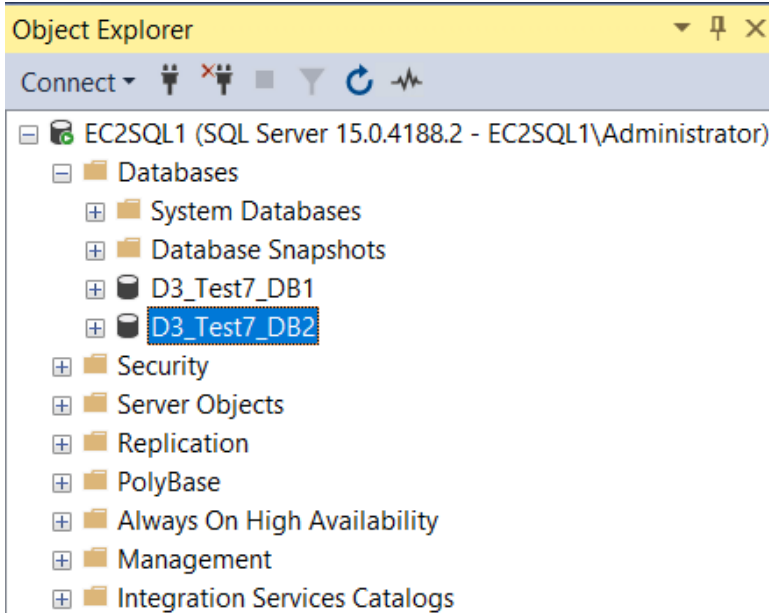
Disk 1

Basic
1024.00 GB
Online

D1 (D:)
1024.00 GB NTFS
Healthy (Primary Partition)

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- In Microsoft SSMS, import the database (D3_Test7_DB1 and D3_Test7_DB2) by attaching the database files.



- Verify that the previously written records are present.

	EMP_CODE	DESIG	FIRST_NAME	LAST_NAME	BASIC	HRA	TOTAL_SAL
1	2428	XUSZ	Letha	Wahl	262665.66...	67926.56	760424.96...
2	7940	SX	Kraig	Boucher	212243.82...	86344.994	746919.53...
3	6574	T	Harlan	Ludwig	333312.50...	81108.9261	656757.00...
4	5563	Z	Brad	Craven	375154.84...	54632.2129	777731.04...
5	8587	RR	Enrique	Sizemore	406366.55...	57982.2825	686743.79...
6	3562	Y	Bennie	Flowers	429458.01...	NULL	723985.527
7	1160	SZ	Francis	Conway	401142.84...	81826.1271	652786.25...
8	2155	SSURR	Nickie	Waite	317093.94...	52985.6098	602185.06...
9	4244	YYU	Boyd	NULL	492931.06...	70726.7528	571781.22...
10	7961	S	Alonzo	Skaggs	279295.541	68115.19	626341.74...

Query executed successfully.

	EMP_ID	DESIG	FIRST_NAME	LAST_NAME	ADDRESS	AGE
1	1000	DEAE	Letha	Wahl	461 NW Riddle Hill Lane, Concord, NH, 04357	43
2	1001	CD	Kraig	Boucher	11 Red Lake Hwy, Harrisburg, Pennsylvania, 30254	31
3	1002	D	Harlan	Ludwig	1753 North Church Pkwy, Boise, Idaho, 95705	21
4	1003	E	Brad	Craven	1695 Hidden Deepwood Cir, Superior Bldg, Indiana...	36
5	1004	DC	Enrique	Sizemore	970 Riddle Hill Pkwy, Comcast Building, Albany, Ne...	34
6	1005	C	Bennie	Flowers	1955 Front Road, Madison, WI, 37006	24
7	1006	AD	Francis	Conway	2844 North Woodfort St, Raleigh, NC, 26635	40
8	1007	DBEDE	Nickie	Waite	2989 East Pine Tree Road, Topeka, KS, 16758	52
9	1008	BAD	Boyd	Skaggs	62 Ashwood Pkwy, 3rd Floor, Madison, Wisconsin, ...	NU...
10	1009	A	Alonzo	Peacock	40 Bayview Lane, Topeka, Kansas, 52885	36

Query executed successfully.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

11. Ensure that the new record can be written in the database (D3_Test7_DB2). In the following screenshot, record number 11 is written in the database.

Results		Messages					
	EMP_ID	DESIG	FIRST_NAME	LAST_NAME	ADDRESS	AGE	
1	1000	DEAE	Letha	Wahl	461 NW Riddle Hill Lane, Concord, NH, 04357	43	
2	1001	CD	Kraig	Boucher	11 Red Lake Hwy, Harrisburg, Pennsylvania, 30254	31	
3	1002	D	Harlan	Ludwig	1753 North Church Pkwy, Boise, Idaho, 95705	21	
4	1003	E	Brad	Craven	1695 Hidden Deepwood Cir, Superior Bldg, Indiana...	36	
5	1004	DC	Enrique	Sizemore	970 Riddle Hill Pkwy, Comcast Building, Albany, Ne...	34	
6	1005	C	Bennie	Flowers	1955 Front Road, Madison, WI, 37006	24	
7	1006	AD	Francis	Conway	2844 North Woodfort St, Raleigh, NC, 26635	40	
8	1007	DBEDE	Nickie	Waite	2989 East Pine Tree Road, Topeka, KS, 16758	52	
9	1008	BAD	Boyd	Skaggs	62 Ashwood Pkwy, 3rd Floor, Madison, Wisconsin, ...	NU...	
10	1009	A	Alonzo	Peacock	40 Bayview Lane, Topeka, Kansas, 52885	36	
11	1010	AE	Getrk	Qhayi	40.gtys	40	

✔ Query executed successfully.

12. Detach the database and take the disk from disk management of the Windows OS offline. Remove the snapshot volume from the Host or host group.
13. Delete the D3_test7_DB2 snapshot.
14. Recover the failed paths between the ESXi cluster and the primary storage system.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

Test 8: Unplanned Outage of On-Premises Servers with Data Return

This test case demonstrates how business can resume operations using AWS EC2 cloud compute in the event of an on-premises compute failure. Data written from the EC2 instance to the SQL server is verified when the on-premises compute is restored.

1. Status before the outage:

- GAD Pair status:

'BLOCKSTR_VSP5200_30_10' Replication '05/05/2022 12:33:45' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	%	I/O M Mode	ID	Storage	Status	Attribute	%	I/O M Mode	Mirror Unit	Fence Type	Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	- L/M	→ 0x0700	540028	PAIR	S-VOL	100%	- L/M	0	GAD	NEVER	-

- Path status in ESXi servers:

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore
<input type="checkbox"/>	Local AVAGO Disk (naa.6001636002a953a0240c539bd2a299c7)	0	disk	1.75 TB	loca
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E105
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Cons

1 - 20 of 26 items | 1 / 2

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:C0:T0:L0	Active (I/...	50:06:0e:80:08:9c:5c:40 50:06:0e:80:08:9c:5c:40	vmhba3:C0:T0:L0	
<input type="radio"/>	vmhba2:C0:T0:L0	Active (I/...	50:06:0e:80:23:3a:9e:20 50:06:0e:80:23:3a:9e:20	vmhba2:C0:T0:L0	

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS FLASH DISK ...

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000001)	1	disk	2.10 TB	E105
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80238235005070823500000700)	0	disk	1.00 TB	Not Cons
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80233a9e0050703a9e00000010)	2	disk	2.00 TB	Not Cons

1 - 20 of 26 items | 1 / 2

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba3:C0:T0:L0	Active (I/...	50:06:0e:80:08:9c:5c:51 50:06:0e:80:08:9c:5c:51	vmhba3:C0:T0:...	
<input type="radio"/>	vmhba2:C0:T0:L0	Active (I/...	50:06:0e:80:23:3a:9e:30 50:06:0e:80:23:3a:9e:30	vmhba2:C0:T0:...	

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- SQL Database status:

Roles (1)

Name	Status	Type	Owner Node	Priority	Information
SQL Server (MSSQLSERVER)	Running	Other	1090SQLNODE1	Medium	

SQL Server (MSSQLSERVER) Properties

General | Fallover

Name: SQL Server (MSSQLSERVER)

Preferred Owners: 1090SQLNODE1, 5200SQLNODE2

Priority: Medium

Status: Running

Node: 1090SQLNODE1

SQL Server Agent Properties - SQLNET

Service state: Running

Auto restart SQL Server if it stops unexpectedly

Auto restart SQL Server Agent if it stops unexpectedly

Error log File name: E:\MSSQL15.MSSQLSERVER\MSSQL\Log\SQLAGENT.OUT

2. Generate new records on the GAD volume by creating a database named D3_Test8_DB1.

Database Properties - D3_Test8_DB1

Database name: D3_Test8_DB1

Owner: PLUTO\Administrator

Use full-text indexing

Logical Name	File Type	Filegroup	Size (MB)	Autogrowth / Maxsize	Path	File N
D3_Test8_...	ROWS...	PRIMARY	8	By 64 MB, Unlimited	D:\MSSQL15.MSSQLSERVER-D\MSSQL\DATA	D3_1
D3_Test8_...	LOG	Not Applicable	8	By 64 MB, Limited to 2...	D:\MSSQL15.MSSQLSERVER-D\MSSQL\DATA	D3_1

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

	PURCHASE_ID	VENDOR_ID	VENDOR_NAME	ITEM_NAME	QTY
1	101	5001	ADX	BOOKS	5
2	102	5002	XXD	PEN	20
3	1103	3003	GTX	PENCIL	300

Query executed successfully.

3. Initiate the outage by abruptly powering off servers of the Production cluster.

- Status of Windows VM:

D3-SQL-NODE1-OS-E1090-DG | [Play] [Stop] [Refresh] [Share] [Refresh] | : ACTIONS

Summary | Monitor | Configure | Permissions | Datastores | Networks | Snapshots | Updates

Powered Off

Guest OS: Microsoft Windows Server 2019 (64-bit)

Compatibility: ESXi 7.0 U2 and later (VM version 19)

VMware Tools: Not running, not installed

[MORE INFO](#)

DNS Name:

IP Addresses:

Host: 172.23.30.45

LAUNCH WEB CONSOLE

D3-SQL-NODE1-OS-VSP5200-DG | [Play] [Stop] [Refresh] [Share] [Refresh] | : ACTIONS

Summary | Monitor | Configure | Permissions | Datastores | Networks | Snapshots | Updates

Powered Off

Guest OS: Microsoft Windows Server 2019 (64-bit)

Compatibility: ESXi 7.0 U2 and later (VM version 19)

VMware Tools: Not running, not installed

[MORE INFO](#)

DNS Name:

IP Addresses:

Host: 172.23.30.48

LAUNCH WEB CONSOLE

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- Pairs status after the outage:

'BLOCKSTR_VSP5200_30_10' Replication '05/06/2022 10:28:29' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	%	I/O M Mode	ID	Storage	Status	Attribute	%	I/O M Mode	Mirror Unit	Fence Type	Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	- L/M	→ 0x0700	540028	PAIR	S-VOL	100%	- L/M	0	GAD	NEVER	-

4. Delete the GAD Pair.

To fail over the workload of an on-premises GAD volume to an EC2 instance, you must delete the GAD pair. After the pair is deleted, the volume from the primary storage system is mounted to the AWS EC2 instance.

- Navigate to the Protector **Dashboard**, select **Monitor**, select the active **Data flow** you want to deactivate, and click **Deactivate**.

Deactivate Data Flow(s) ✕

Deactivating the selected data flow(s) will immediately stop all *scheduled* activity.

Any hardware replication pairs will be marked as "eligible for teardown". To manage the replication pairs, navigate to the Storage screen for the destination array and filter for replications eligible for teardown, then either:

- Use the teardown action to remove the replication.
- Use the dissociate action to keep the replication but remove it from the inventory.

Cancel Deactivate

- Navigate to the Protector **Dashboard**, select **Storage**, and then select the **Secondary Storage System**. Click **Replication**, select the pair you want to tear down, and then click **Teardown**.

Teardown Hitachi Block Replication ✕

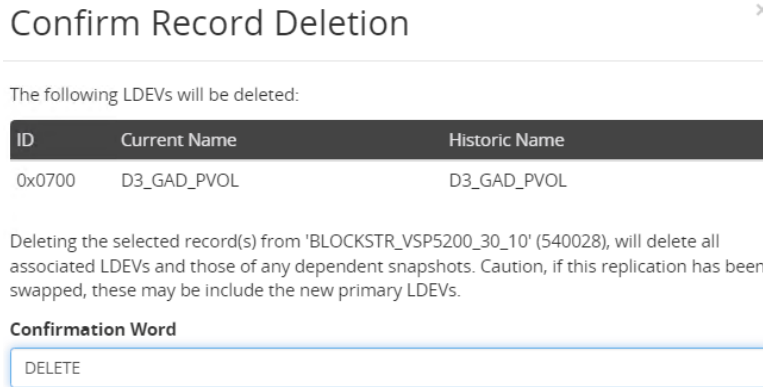
Confirmation Word

TEARDOWN

Tearing down a replication will **remove the volume pairings on the array**. If you are certain you want to perform this operation enter 'TEARDOWN' in the field above.

Cancel Teardown


- Navigate to the Protector **Dashboard**, select **Storage**, and then select the **Secondary Storage System**. Click **Replication**, select the pair you want to delete, and then click **Delete**.



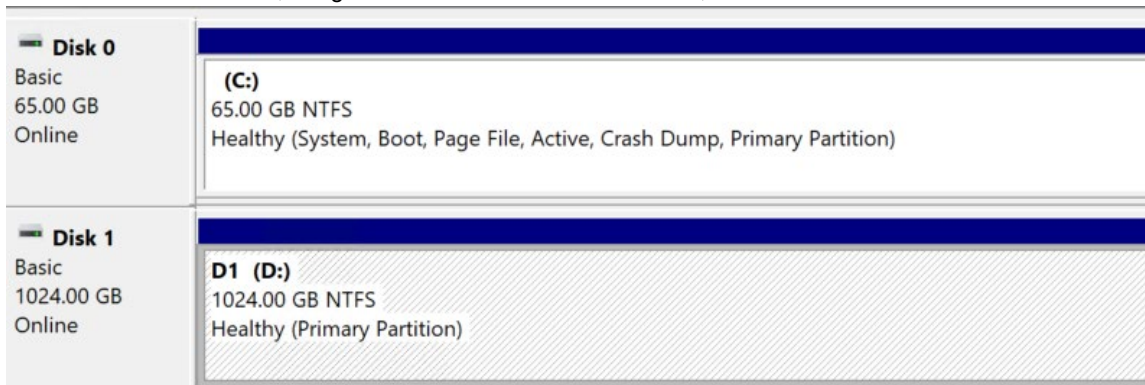
5. Mount volume of the primary storage system to the EC2 instance.

The iSCSI port of the primary storage system and AWS Direct Connect between Near Cloud and AWS are used to mount the volume in the AWS EC2 instance. The storage port (CL1-B) is used in this configuration and one Host Group from port CL1-B is required to mount the volume.

- a. Add the Host Group ID in the VSM in Ops Center Administrator before mounting the volume. To add the Host Group ID to the VSM, see section [Test 5.6](#).
- b. Attach the Primary GAD volume to the iSCSI port (CL1-B) and map the volume to the EC2 instance as described in section [Test 4.3](#).

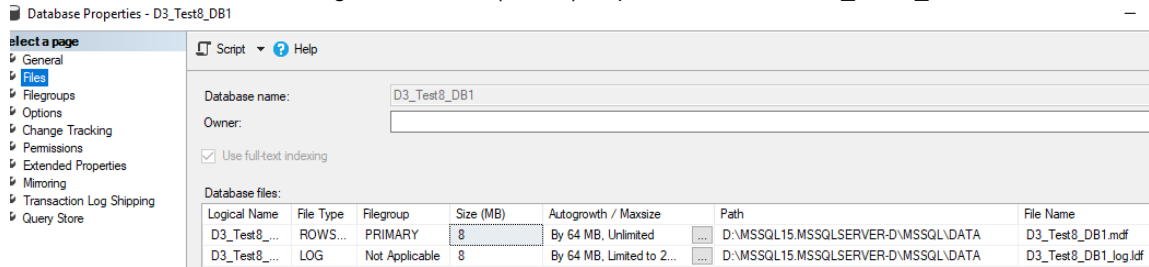


6. In the AWS EC2 instance, bring the drives online. In this scenario, the drive is mounted as drive D.



Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

7. In Microsoft SQL Server Management Studio (SSMS), import the database D3_Test8_DB1.



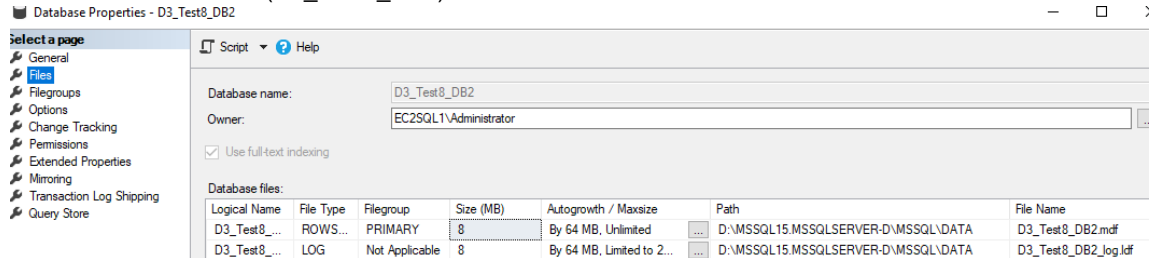
8. Verify that the previously written records are present.

	PURCHASE_ID	VENDOR_ID	VENDOR_NAME	ITEM_NAME	QTY
1	101	5001	ADX	BOOKS	5
2	102	5002	XXD	PEN	20
3	1103	3003	GTX	PENCIL	300

9. Ensure that the new record can be written. In the following screenshot, a record with the PURCHASE_ID 103 is added.

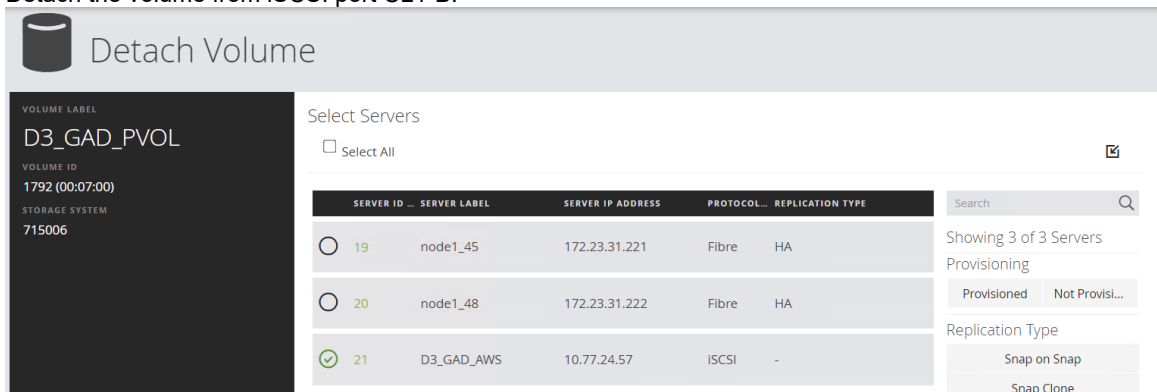
	PURCHASE_ID	VENDOR_ID	VENDOR_NAME	ITEM_NAME	QTY
1	101	5001	ADX	BOOKS	5
2	102	5002	XXD	PEN	20
3	103	4456	MBO	COPY BO...	400
4	1103	3003	GTX	PENCIL	300

10. Create a new database (D3_Test8_DB2) in the EC2 instance.



11. Recover from the outage.

- a. Detach the database and take the disk from the disk management of the Windows OS offline.
- b. Detach the volume from iSCSI port CL1-B.



Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

- c. Activate the GAD dataflow as described in section [Test 2.4.v](#).

'BLOCKSTR_VSP5200_30_10' Replication '05/06/2022 13:04:35' Pairs

Original Primaries						Original Secondaries						Properties			
ID	Storage	Status	Attribute	%	I/O M Mode	ID	Storage	Status	Attribute	%	I/O M Mode	Mirror Unit	Fence Type	Level	Quorum
✓ 0x0700	715006	PAIR	P-VOL	100%	- L/M	→ 0x0700	540028	PAIR	S-VOL	100%	- L/M	0	GAD	NEVER	-

- 12. To bring the primary data center back online, power on the servers of the Production Cluster.

After the primary data center is back online, use the Windows Failover Cluster Manager to verify that the nodes are online. Navigate to role and ensure that the status is online.

SQL Cluster status:

The screenshot displays the Failover Cluster Manager interface for the cluster 'MSFCSQLGAD.pluto.com'. The 'Roles (1)' section shows the 'SQL Server (MSSQLSERVER)' role is in a 'Running' state on node '1090SQLNODE1'. The 'SQL Server (MSSQLSERVER) Properties' dialog box is open, showing the 'Failover' tab. Under 'Preferred Owners', both '1090SQLNODE1' and '5200SQLNODE2' are checked. The 'Priority' is set to 'Medium', and the 'Status' is 'Running' on node '1090SQLNODE1'. Below the main role list, various services like 'Analysis Services', 'SQL Server Analysis Services CEIP', 'SQL Server CEIP (MSSQLSERVER)', 'SQL Server launchpad (MSSQLSERVER)', 'SQL Server Polybase Dms (MSSQLSERVER)', and 'SQL Server Polybase Engine (MSSQLSERVER)' are all shown as 'Online'. Storage elements 'Cluster Disk 1' and 'Cluster Disk 2' are also 'Online', and the 'Server Name' is 'SQLNET'.

- a. In Microsoft SQL Server Management Studio (SSMS), verify that database D3_Test8_DB1 (created while running on GAD) and D3_Test8_DB2 (created while running on the AWS EC2 instance) are present.

Continuous Analytics with Near-Cloud Solution with Global-Active Device and Hitachi Thin Image

```

use master;
SELECT
  name 'Logical Name',
  physical_name 'File_Location'
FROM sys.master_files;
1

```

	Logical Name	File_Location
22	D3_Test5_DB1	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test5_DB...
23	D3_Test5_DB1_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test5_DB...
24	D3_Test5_DB3	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test5_DB...
25	D3_Test5_DB3_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test5_DB...
26	juno_gad5	E:\MSSQL15.MSSQLSERVER\MSSQL\DATA\juno_gad5.mdf
27	juno_gad5_log	E:\MSSQL15.MSSQLSERVER\MSSQL\DATA\juno_gad5_log.ldf
28	Juno-X	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\Juno-X.mdf
29	Juno-X_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\Juno-X_log.ldf
30	D3_Test6_DB1	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test6_DB...
31	D3_Test6_DB1_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test6_DB...
32	D3_Test7_DB1	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test7_DB...
33	D3_Test7_DB1_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test7_DB...
34	D3_Test7_DB2	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test7_DB...
35	D3_Test7_DB2_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test7_DB...
36	D3_Test8_DB1	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test8_DB...
37	D3_Test8_DB1_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test8_DB...
38	D3_Test8_DB2	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test8_DB...
39	D3_Test8_DB2_log	D:\MSSQL15.MSSQLSERVER\DIMSSQL\DATA\D3_Test8_DB...

Query executed successfully.

- b. Verify whether the new records written to D3_Test8_DB1 on the AWS EC2 instance are present.

	PURCHASE_ID	VENDOR_ID	VENDOR_NAME	ITEM_NAME	QTY
1	101	5001	ADX	BOOKS	5
2	102	5002	XXD	PEN	20
3	103	4456	MBO	COPY BOOK	400
4	1103	3003	GTX	PENCIL	300