HITACHI
Inspire the Next

Cloud-based Analytics at Equinix with Hitachi Universal Replicator and Hitachi Thin Image: VSP E790 and VSP 5200

Using Hitachi Cloud Connect for Equinix

Reference Architecture

March 2022

Hitachi Vantara

Contents

Notices and Disclaimer

Executive Summary

Hitachi Universal Replicator (HUR) is a proven, hardware-based replication solution for copying data to a remote location. Hitachi Thin Image (HTI) creates instantaneous point-in-time snapshots that are space efficient. This reference architecture documents how to set up a disaster recovery (DR) solution with Hitachi Universal Replicator 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 on-demand analytics, data mining, DR testing, development testing, and similar use cases.

The environment used for the validation included a Virtual Storage Platform (VSP) E790 used as the primary storage system located in a traditional, “on-premises” data center in Colorado and a VSP 5200 used as the secondary storage system located in an Equinix colocation data center in California.

The cost of owning and operating a second data center for the purpose of disaster recovery is not insignificant, so leasing a small footprint in a colocation data center is a cost-attractive alternative. Equinix, Inc. is a leading provider of these services. Through our collaboration with Equinix, Hitachi Vantara has 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 one agreement and invoice, greatly simplifying and accelerating their time to market. By using Equinix IBX data centers and Equinix Fabric™ to interconnect sources of data to applications, VSP storage systems enable organizations to locate their data next to clouds while still maintaining control by enabling applications such as data protection and back-up for hybrid- and multi-cloud data availability.

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

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.

About This Guide

Introduction

This reference architecture documents how to set up a disaster recovery solution with Hitachi Universal Replicator with cascading Hitachi Thin Image snapshots using 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.

The environment used for the validation included a VSP E790 used as the primary storage system and located in a traditional, “on-premises” data center in Colorado and a VSP 5200 used as the secondary storage system located in an Equinix colocation data center in California. To access data on the secondary storage system at the DR site, we used cloud compute from AWS in the form of EC2 virtual machines.

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

Intended Audience

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

Document Revisions

Revision Number	Date	Author	Details
1.0	March 2022	Hitachi Vantara LLC	Initial Release

References

- *Hitachi Universal Replicator User Guide*
- *Hitachi Thin Image User Guide*
- *Hitachi Ops Center Protector User Guide*
- *Hitachi Ops Center Administrator User Guide*

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!

Solution Overview

Hitachi Universal Replicator provides a solution to avoid situations when a data center is affected by a disaster that stops operations for a long period of time. Two storage systems are required in an HUR implementation. Typically, the secondary storage system is located in a second data center that is far from the first data center that contains the primary storage system. It is important to locate the two data centers far 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 into 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 on-demand analytics, data mining, DR testing, development testing, and similar use cases.

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 Universal Replicator with cascading Hitachi Thin Image snapshots:

- The solution allows business to resume operations quickly when a disaster brings down the primary data center.
- Interment wide area network (WAN) failures between data centers do not stop operations. Hitachi Universal Replicator can store un-replicated data in local journals until the WAN connection is recovered.
- Hitachi Thin Image is space efficient, so less storage capacity is needed.
- Replication is done at the block level by the storage systems, which has the following advantages:
  - Dedicated software to move data between data centers is not required.
  - CPU processing cycles are not used to move data between data centers.
  - Various operating systems, clustering software programs, and applications are supported.
- 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.

Key Components

The major components of the solution are described below. Technical details are provided in the Hardware and Software section.

- Storage Systems: Two storage systems are required in this solution. This validation used a VSP E790 as the primary storage system and a VSP 5200 as the secondary storage system.
  - iSCSI Cards: 10 GbE iSCSI ports were used to connect to the AWS cloud.
- Hitachi Universal Replicator: HUR exists in the microcode, also known as firmware, of the storage system. It does not require any additional hardware. However, it must be activated by a license key on both storage systems.
- Hitachi Thin Image: HTI exists in the microcode and must be activated by 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 virtual machine and can run on the same machine as Ops Center 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 needed 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.

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 at the primary site and a standalone SQL instance running on AWS EC2 was used at the DR site. Records were written to the SQL database before each failure test and then verified after the failure to ensure data consistency. Snapshots were verified in a similar manner.

The primary SQL cluster resided on a Microsoft Windows Failover Cluster running on virtual machines. Two volumes were mapped from the VSP E790 storage system and passed through to the virtual machines as Raw Device Mappings (RDM). These volumes hosted the database files and logs. The two volumes were replicated with HUR to the VSP 5200 storage system at the DR site.

A third volume was configured as a VMFS datastore. The virtual machines with the SQL cluster instance resided on this datastore. A fourth volume was used for the Microsoft Windows Failover Cluster Quorum Witness. These two volumes were not replicated.

At the DR site, we used cloud compute from AWS in the form of EC2 virtual machines to access the data replicated with HUR or snapshotted with HTI.

Figure 1 illustrates the SQL setup.

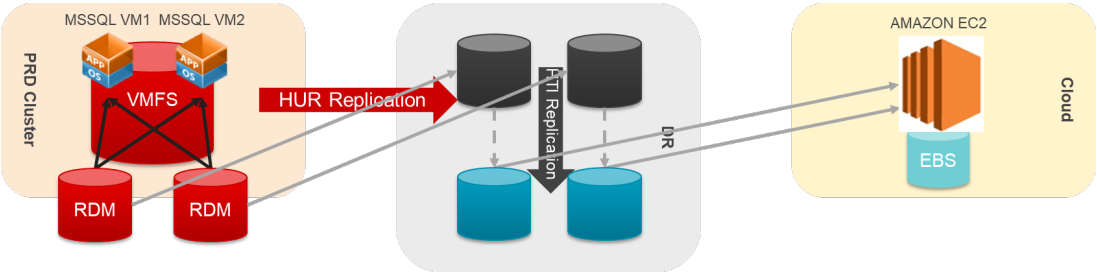

Figure 1. *SQL Setup*

High Level Diagram

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

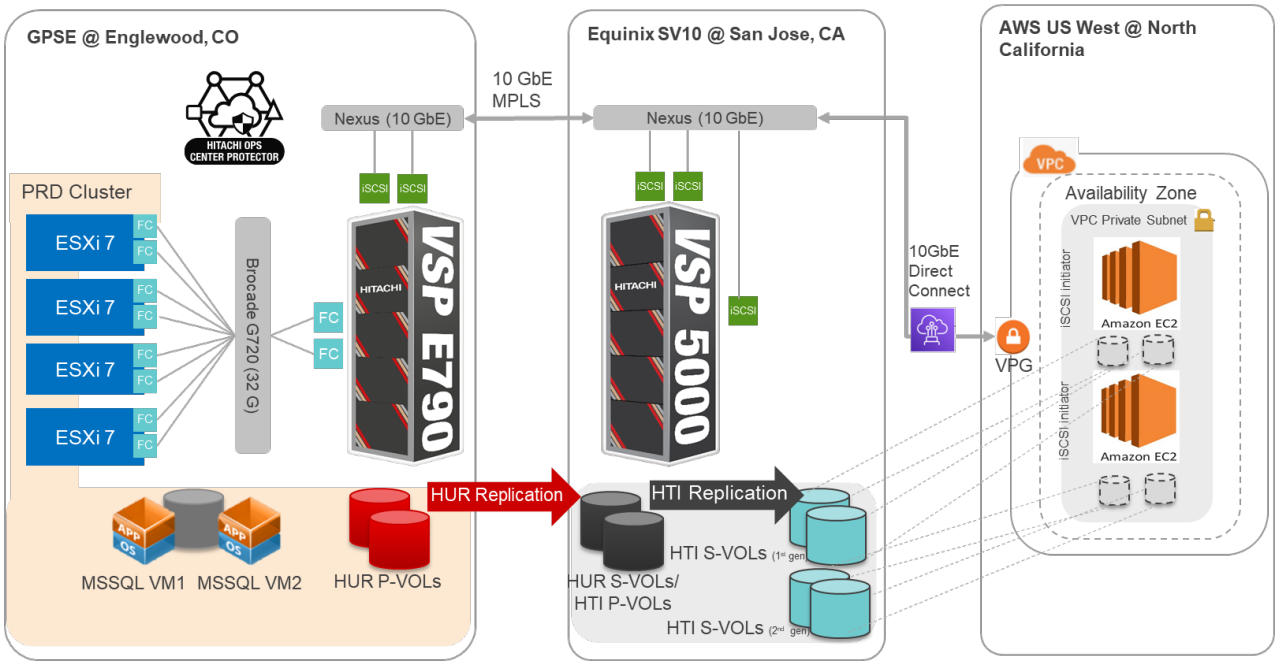

Figure 2. *Test Environment*

Reference Architecture

Hardware and Software

Table 1 provides the hardware specifications used in this validation.

Item	Description	Version	Function
Hitachi VSP E790	768 GB cache (2) 32-core MPUs (3) RAID6 6D-2P parity groups (2) 32 Gbps FC ports (4) 10 GbE iSCSI ports	SVOS RF 9.6 93-06-01-x0/00-M039	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
Quanta D51B-2U	(2) 12-core Intel Xeon E5-2680 v3 @ 2.5 GHz 192 GB cache (1) Emulex LPe32002 HBA (1) Intel 82599 10 Gigabit Ethernet Controller	BMC 3.83.00 BIOS S2B_3B22	4-node primary vSphere cluster (PRD)
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 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 secondary data center. Serviced Direct Connect to AWS.

Table 1. *Hardware Components*

Table 2 provides the software specifications used in this validation.

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	7.3.0.90312-R7.3	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

Item	Version	Function
Microsoft SQL Server Enterprise	15.0.2000.5	Database application used to validate data consistency

Table 2. *Software Components*

Test Scenarios

Table 3 lists the test scenarios performed in the validation.

Test	Description	Success Criteria
1	Prepare the environment: 1. Provision 1x 2 TB LUN on VSP E790 to use as a datastore (non-replicated) and 2x 2 TB LUNs on VSP E790 to use as RDMs (will be replicated with HUR). 2. Provision 1x 400 GB volume on both storage systems to use as HUR journal. 3. Deploy 2x MS SQL 2019 virtual machines on the datastore. 4. Attach RDMs as shared disks to MS SQL 2019 virtual machines. Set up SQL Server failover cluster on virtual machines. Create SQL database on shared disks. 5. Deploy 1x MS SQL 2019 as AWS EC2 virtual machine. 6. Deploy Ops Center and ISM proxy virtual machines. 7. Define HUR remote paths.	Environment is set up per specifications.
2	Create HUR pairs with cascaded HTI snapshots: 1. Create HUR pairs using Ops Center Protector. 2. Create HTI snapshots using Ops Center Protector.	• HUR pairs are created. • HTI snapshots are created.
3	Access snapshots from AWS EC2: 1. Trigger HTI snapshots. 2. Map snapshots to EC2 virtual machine. Bring copy of SQL database online. 3. Write new records in copy of SQL database. 4. Detach database. Delete snapshots.	• Data created at primary site is present on snapshots. • EC2 virtual machine can write to snapshots without affecting primary site.
4	Multigeneration HTI snapshots: 1. Trigger HTI snapshots (1st generation). 2. Map 1st generation snapshots to EC2 virtual machine. Bring copy of SQL database online. 3. Write new records in copy of SQL database. 4. Write new records in SQL database on HUR primary volumes. 5. Trigger HTI snapshots (2nd generation). 6. Map 2nd generation snapshots to different EC2 virtual machine. Bring copy of SQL database online.	SQL database on 2nd generation snapshots contains records written by the primary site but not records from the 1st generation snapshots.
5	Planned outage with data return: 1. Take SQL database offline. Power off virtual machines 2. Perform Swap to make HUR secondary volumes accessible. 3. Map secondary volumes to EC2 virtual machine. Bring SQL database online.	• Snapshots can still be taken while HUR roles are swapped. • Data created at DR site is

Test	Description	Success Criteria
	4. Write new records in SQL database. 5. Trigger HTI snapshots. 6. Map snapshots to different EC2 virtual machine. Bring copy of SQL database online. Write new records in copy of SQL database. 7. Detach copy of SQL database. Delete snapshots. 8. Detach SQL database. 9. Perform Swap to reverse HUR replication to original direction. 10. Power on virtual machines. Bring SQL database online. 11. Verify records were copied back to primary site.	returned to primary site.
6	Unplanned outage with data return: 1. Abruptly disable host connections and remote paths to VSP E790. 2. Perform Swap to make HUR secondary volumes accessible. 3. Map secondary volumes to EC2 virtual machine. Bring SQL database online. 4. Write new records in SQL database. 5. Trigger HTI snapshots. 6. Map snapshots to different EC2 virtual machine. Bring copy of SQL database online. Write new records in copy of SQL database. 7. Detach copy of SQL database. Delete snapshots. 8. Restore host connections and remote paths to VSP E790. 9. Detach SQL database. 10. Perform Swap to reverse HUR replication to original direction 11. Power on virtual machines. Bring SQL database online. 12. Verify records were copied back to primary site.	• Snapshots can still be taken while HUR roles are swapped. • Data created at DR site is returned to primary site.
7	Restore from snapshots: 1. Trigger HTI snapshots. 2. Corrupt SQL database at primary site. 3. Suspend HUR pairs. 4. Restore HTI snapshots to HUR secondary volumes. 5. Copy data from HUR secondary volumes to HUR primary volumes.	• Corrupted database at primary site can be recovered with snapshots at DR site.

Table 3. *Test Scenarios*

Guidelines and Recommendations

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

- The testing performed in this validation is crash consistent. Protector does not support application consistent backup for MS SQL with HUR. Because HUR is asynchronous, Protector cannot guarantee data consistency when a snapshot is triggered at the DR site.
- Protector can be deployed on a Windows Failover Cluster to guard from application failure or corruption. For customers without Windows Failover Cluster software, it is recommended to run 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. For example, if the primary data center goes offline, Protector can issue a “horctakeover” (hard failover) to the remaining HUR secondary storage system.
- HUR remote paths must be created outside of Protector using the embedded storage management interface (Storage Navigator) or raidcom. The raidcom command is: **raidcom add rcu -cu_free 51234 R900 1 -mcu_port CL1-C -rcu_port CL1-D**.
  - **51234 R900 1**: Serial number of the remote storage system, model of the remote storage system, and remote path group ID.
  - **CL1-C**: Local storage port that will be used for this remote path.
  - **CL1-D**: Remote storage port that will be used for this remote path.
- The version of Protector available during testing (7.3.0), had the following limitations:
  - Replication target volumes must be mapped to iSCSI host groups outside of Protector by using Hitachi Ops Center Administrator or another tool.
  - In our testing, creating snapshots for analytics had a dependency on the primary site and could not be triggered when the primary site was unavailable.
- Although performance testing was not in the scope of this validation, some basic performance observations were made:
  - To increase the initial copy throughput, add more HUR remote paths, even if the line speed has not been reached. This works because during HUR Initial Copy, the secondary storage system sends a fixed number of read requests on each remote path to the primary storage system. The secondary storage system waits until data is sent back from the primary storage system before sending out the next batch of read requests. Adding more remote paths increases the total number of outstanding read requests and therefore increases the initial copy throughput.
  - If iSCSI ports are used for HUR remote paths, increase the port Maximum Window Size parameter from the default 64 KB to the maximum of 1024 KB. This provided a 10x increase in Initial Copy throughput in our testing.
  - Although Jumbo frames were not used, we suspect they would also provide a performance boost. The effectiveness of Jumbo frames will be quantified in a future validation.
- Journal volumes should be sized to meet all possible data scenarios based on your business requirements. If the amount of data exceeds capacity, performance problems and suspensions result. In this case, consider a services engagement for a new HUR implementation.

Validation Results

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

Test 1: Prepare the Environment

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

1. Using Hitachi Ops Center Administrator, four volumes were mapped to the PRD cluster. The first two volumes, 45056 (00:B0:00) and 45057 (00:B0:01), are RDMs for shared storage on the SQL cluster. The third volume, 45058 (00:B0:02), is the quorum disk for the Windows Failover Cluster. The fourth volume, 45062 (00:B0:06), is the VMFS datastore that hosts the SQL virtual machines. The following screenshot shows the volumes:

SELECT	ID	VIRTUAL ID	S/N	VIRTUAL S/N	LABEL	POOL ID	ATTRIBUTE	STATUS	CAPACITY SAVI...	TOTAL	USAGE	PROVISIONING...
	45056 (00:B...	45056 (00:B...	611020	611020	D2_RDM_0	0	Compressio...	Normal	Deduplicati...	2.00 TiB	3%	Unmanaged
	45057 (00:B...	45057 (00:B...	611020	611020	D2_RDM_1	0	Compressio...	Normal	Deduplicati...	2.00 TiB	1%	Unmanaged
	45058 (00:B...	45058 (00:B...	611020	611020	D2_Win_Qu...	0	Thin	Normal	No	1.00 GiB	16%	Unmanaged
	45062 (00:B...	45062 (00:B...	611020	611020	D2_VMFS	0	Compressio...	Normal	Deduplicati...	2.00 TiB	3%	Unmanaged

Note: You do not need to pre-create the HUR target volumes on the secondary storage system because Protector creates these volumes. You can pre-create the HUR journal, but Protector can create the HUR journal as well. These steps are described in section Test 2.

Note: Protector creates replication target volumes with **Capacity Saving=No** regardless of the capacity saving setting of the source volumes.

2. Protector is used to view the HUR remote paths. The first of the following screenshots shows the paths from the primary storage system to 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	CL1-G	iSCSI	✔ Normal	-	-	Receiver
1	CL2-C	CL2-H	iSCSI	✔ Normal	-	-	Receiver
2	CL3-C	CL3-G	iSCSI	✔ Normal	-	-	Receiver
3	CL4-C	CL4-H	iSCSI	✔ 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	CL1-G	CL1-C	iSCSI	✔ Normal	-	-	Receiver
1	CL2-H	CL2-C	iSCSI	✔ Normal	-	-	Receiver
2	CL3-G	CL3-C	iSCSI	✔ Normal	-	-	Receiver
3	CL4-H	CL4-C	iSCSI	✔ Normal	-	-	Receiver

Note: HUR remote paths must be created with Storage Navigator or raidcom.

3. In this case, we chose to have Protector create the HUR journal volumes for the primary site and DR site. Also, note that HUR secondary volumes and HTI snapshots will be created automatically by Protector during the pair creation step.
4. The virtual machine was configured with four hard disks: one 80 GB virtual disk, two 2,048 GB RDMs, and one 1 GB RDM. The second SQL virtual machine was configured identically, but instead of three RDMs, it used "Existing Hard Disks" of SQL Node1. The following screenshot shows the settings of the SQL Node1 virtual machine:

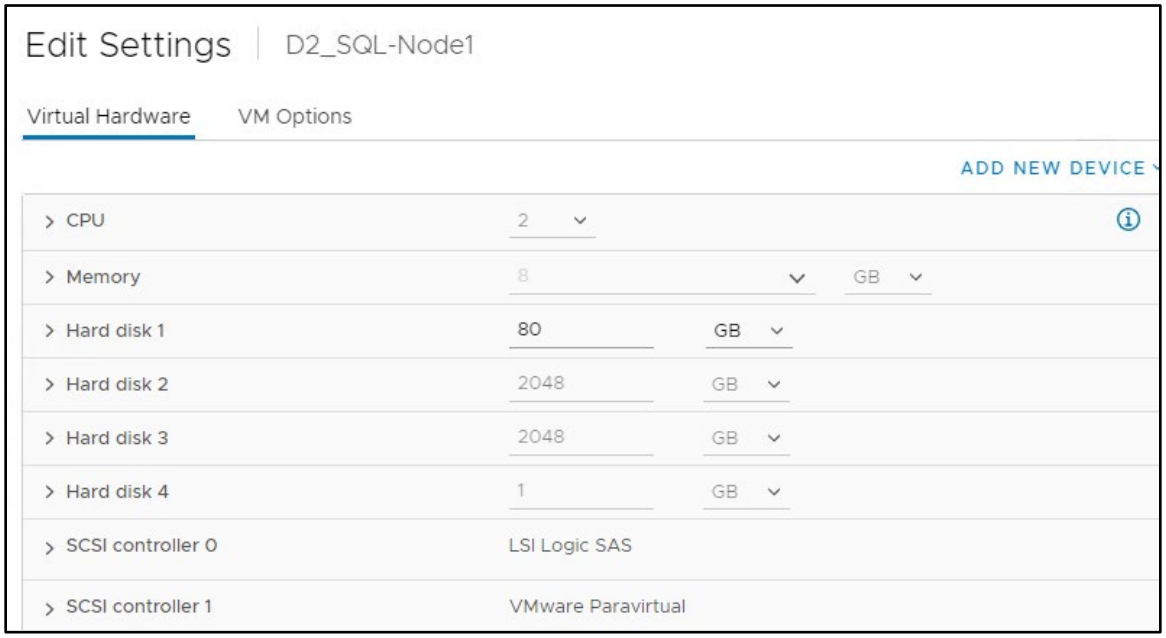

5. Using Windows Failover Cluster Manager from within one of the SQL virtual machines, you can see the Windows Failover Cluster called D2CLU01, which owns three disks. The two 2 TB disks are assigned to the SQL Server role D2SQL. The 1 GB disk is set as Disk Witness in Quorum.

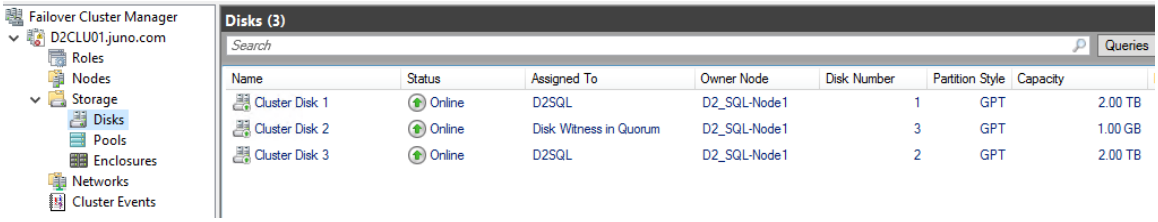

6. The two 2 TB cluster disks are mapped to drive G and drive E as follows:

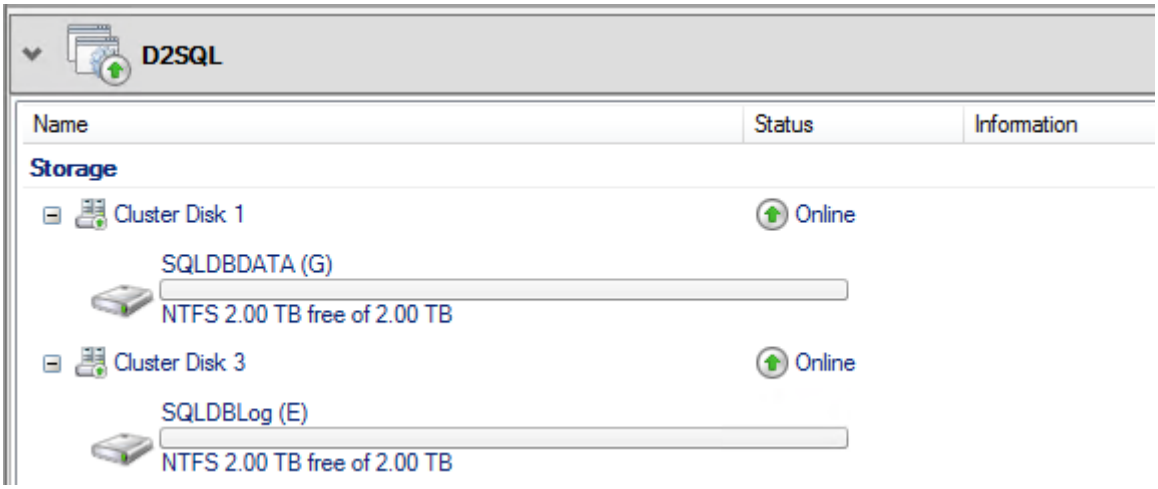

7. To locate the database files registered to this SQL instance, an SQL query was performed. The last two records, circled in red, make up a sample database. The primary database file, D2_DATABASE.mdf, resides on the first of the two cluster disks. The database log file, D2_DATABASE_log.ldf, resides on the second of the two cluster disks.

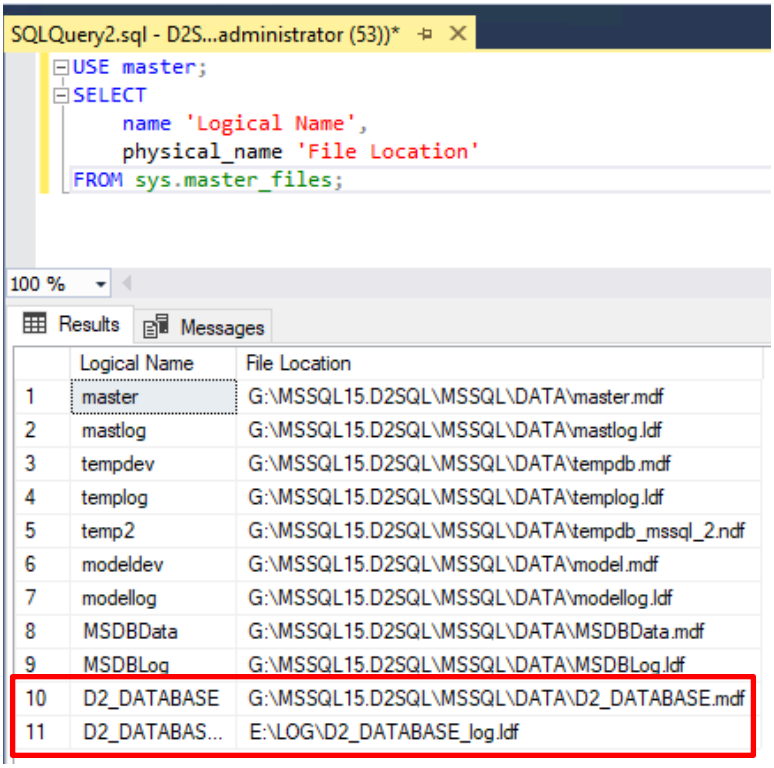

	Logical Name	File Location
1	master	G:\MSSQL15.D2SQL\MSSQL\DATA\master.mdf
2	mastlog	G:\MSSQL15.D2SQL\MSSQL\DATA\mastlog.ldf
3	tempdev	G:\MSSQL15.D2SQL\MSSQL\DATA\tempdb.mdf
4	templog	G:\MSSQL15.D2SQL\MSSQL\DATA\templog.ldf
5	temp2	G:\MSSQL15.D2SQL\MSSQL\DATA\tempdb_mssql_2.ndf
6	modeldev	G:\MSSQL15.D2SQL\MSSQL\DATA\model.mdf
7	modellog	G:\MSSQL15.D2SQL\MSSQL\DATA\modellog.ldf
8	MSDBData	G:\MSSQL15.D2SQL\MSSQL\DATA\MSDBData.mdf
9	MSDBLog	G:\MSSQL15.D2SQL\MSSQL\DATA\MSDBLog.ldf
10	D2_DATABASE	G:\MSSQL15.D2SQL\MSSQL\DATA\D2_DATABASE.mdf
11	D2_DATABAS...	E:\LOG\D2_DATABASE_log.ldf

8. To access data at the DR site, we used an AWS EC2 virtual machine running Windows Server 2019 with a standalone SQL instance. The following screenshot shows the standalone SQL instance:

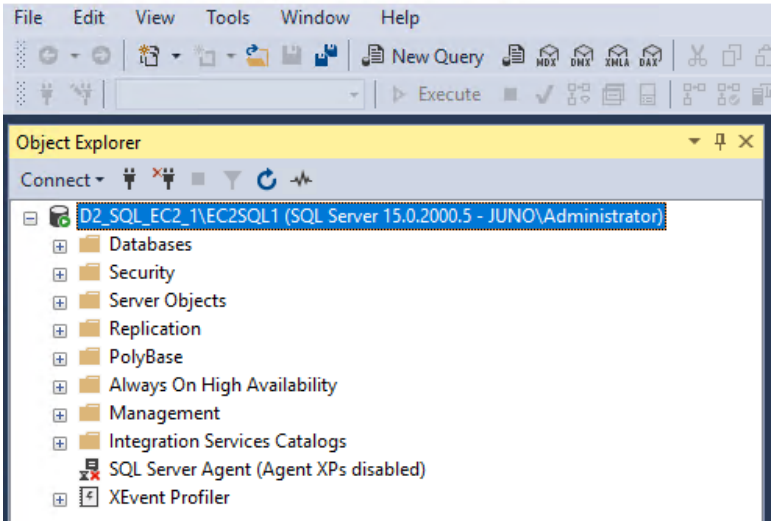

9. Next, we added the AWS EC2 virtual machines as server objects in Ops Center Administrator as follows:
   a. To locate the iSCSI initiator name of the EC2 virtual machine, the PowerShell command **(Get-InitiatorPort).NodeAddress** was used.
   b. In the Administrator dashboard, we selected **Servers > Add Servers**.

c. After clicking the **+** (add) symbol under iSCSI Servers, we entered the EC2 virtual machine name and iSCSI initiator name, then clicked **Submit**.

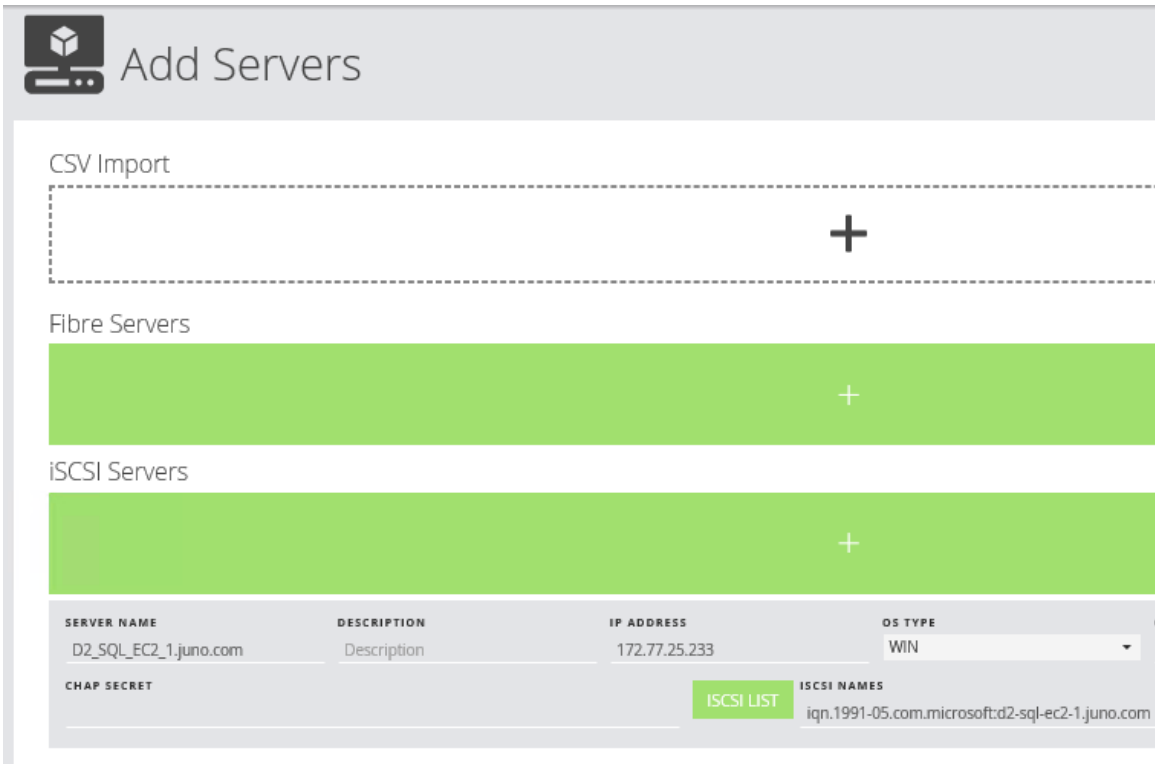

Test 2: Create HUR Pairs with Cascaded HTI Snapshots

This test case describes the process of creating Hitachi Universal Replicator pairs with cascaded Hitachi Thin Image pairs in Protector. The process is summarized in the following workflow.

Register Storage Systems		Define Node for Primary Volumes		Define Policy		Define & Activate Data Flow
Register storage systems in Protector with their respective proxy node (ISM).	➔	Specify which volumes will serve as the replication source.	➔	Specify the type of replication and snapshot frequency that will be used.	➔	Link the source node with the replication target and define associated replication and snapshot parameters.

Note: The deployment of proxy (ISM) nodes is not covered in this document because there are many options between the supported operating systems and physical server versus a virtual machine. For your reference, the ISM Preconfigured Media was used for the validation.

1. To register the primary and secondary storage systems in Protector, complete the following steps:
   a. Click **Nodes** and click the **+** (plus) symbol.
   b. Select **Storage** and then **Hitachi Block Device**. Click **Next**.
   c. Enter a node name and tags (optional) and check the **I confirm…** checkbox. Click **Next**.

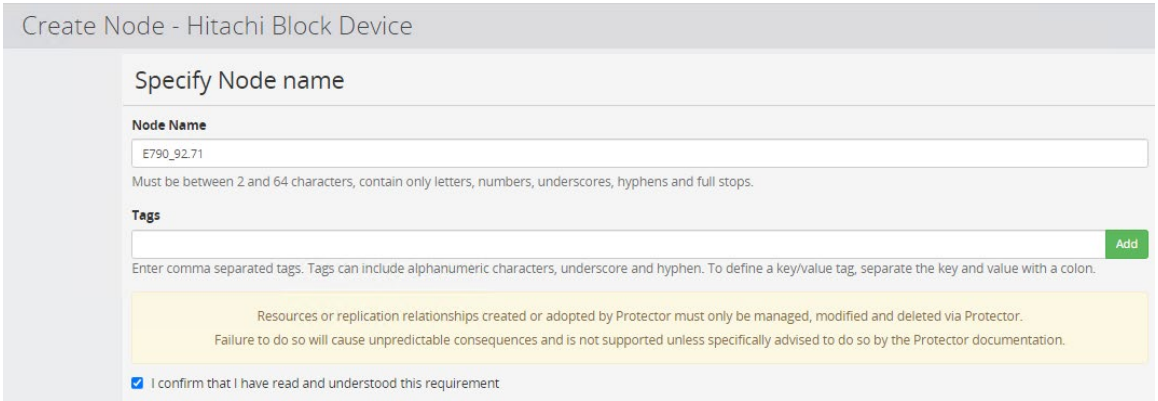

   d. In the following screen, click **Next**.

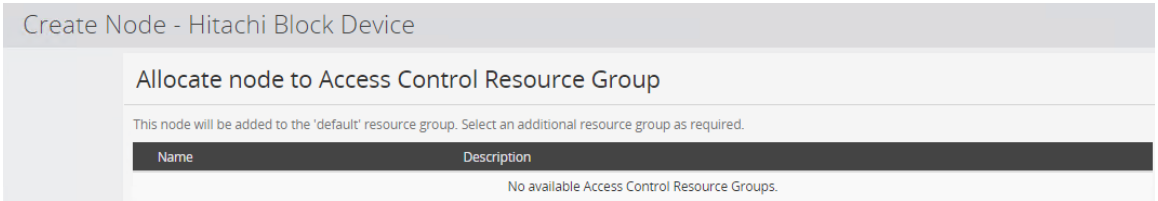

   e. Select the matching proxy node from the dropdown and then click **Next**.

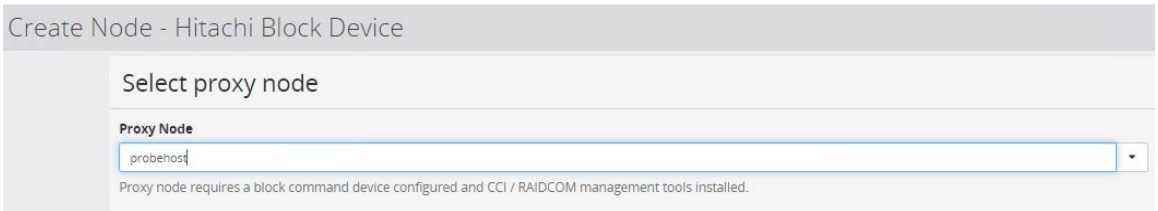

f. Enter a directory for the proxy node to store the metadata and click **Next**.

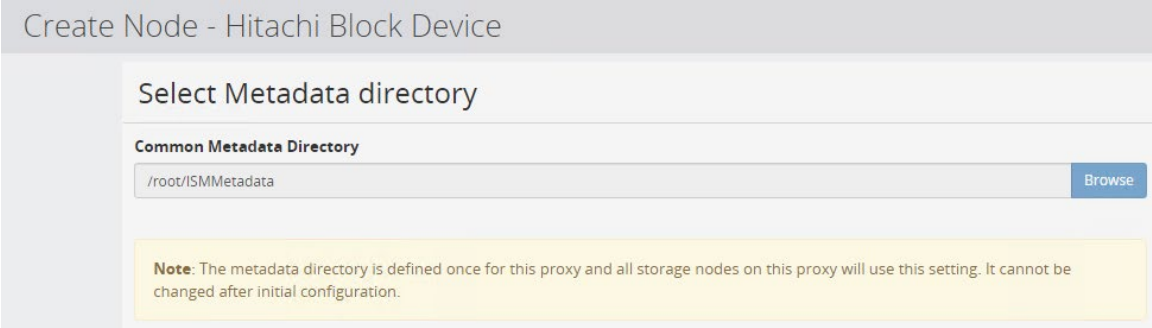

g. Select the command device if using FC (and already mapped to the proxy host) or enter the IP address of the SVP/controller if using IP. Click **Next**.

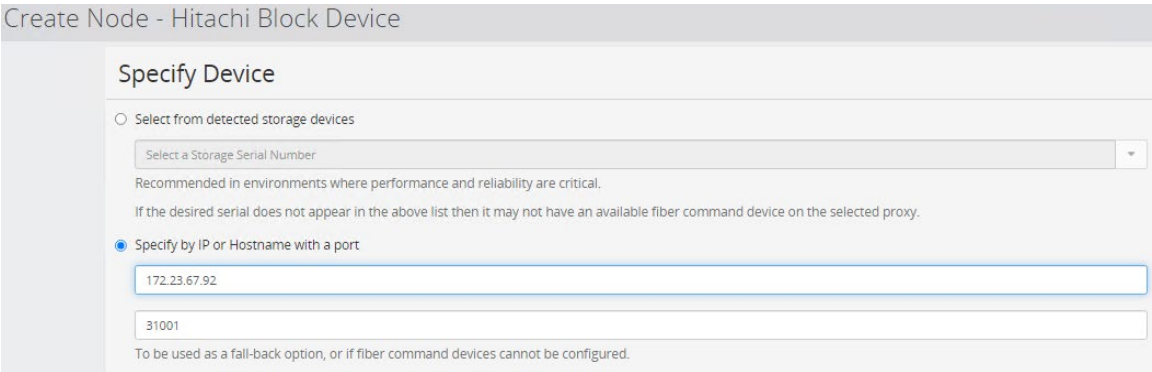

h. Enter the storage system credentials and click **Next**.

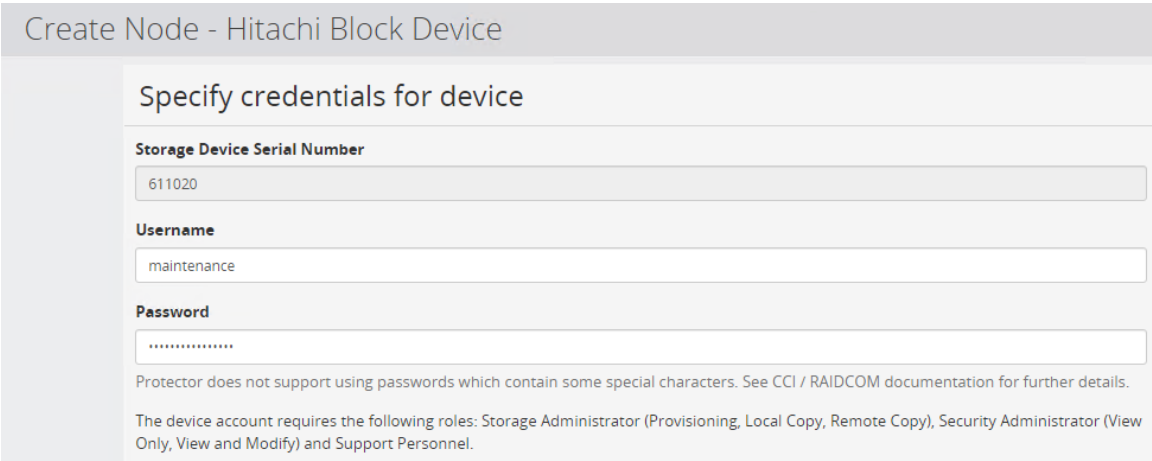

i. In the Specify Configuration for Global Replication Reports screen, click **Next**.

j. Specify **All** for LDEV Range and click **Next**.

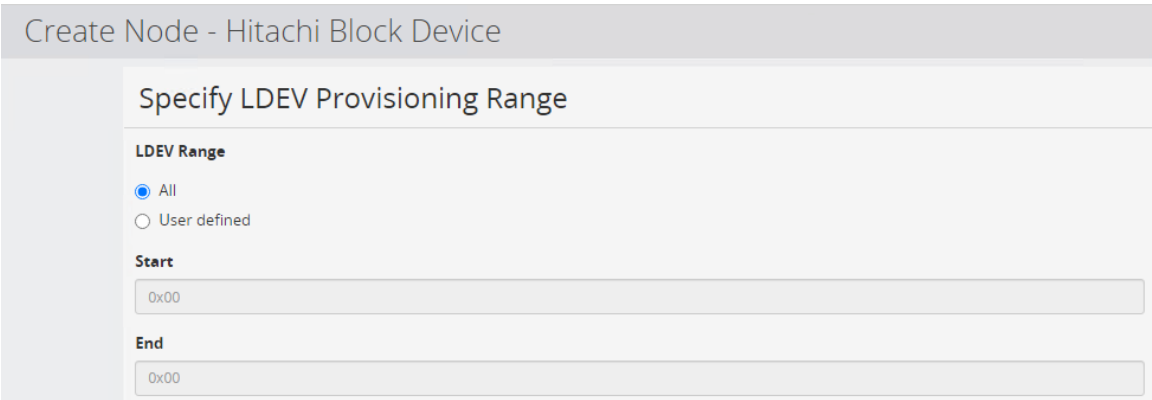

k. Continue to click **Next** until you reach the summary screen.

l. Review the information in the summary screen and then click **Finish**.

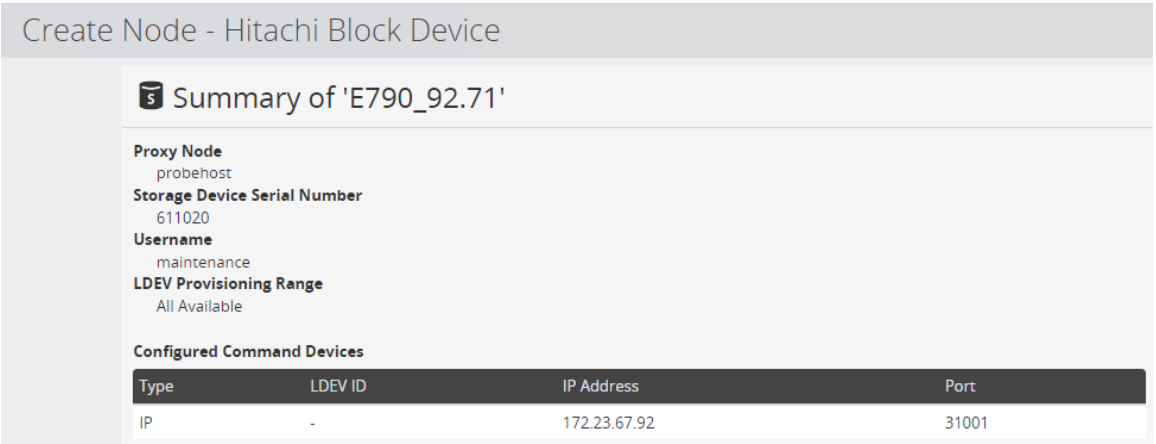

2. Create a Hitachi Block Host node to designate the HUR source volumes.

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

b. Select **Host** and then **Hitachi Block Host**. Click **Next**.

c. Enter a name and tags (optional) and click **Next**.

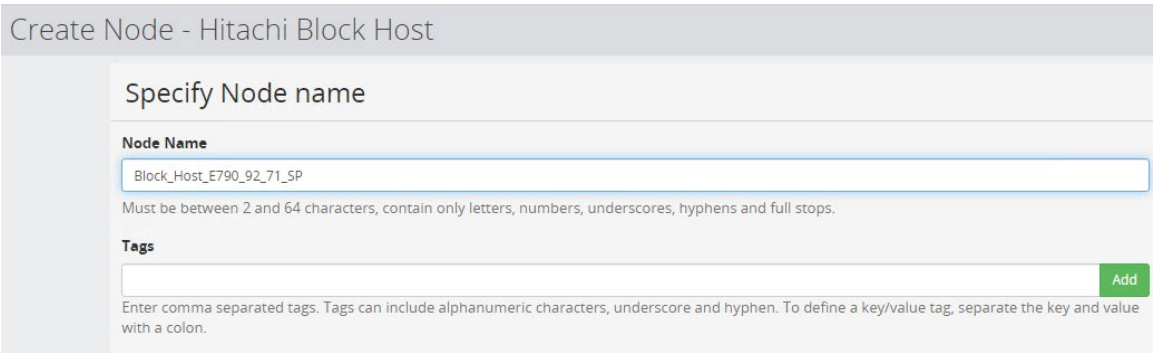

d. In the Allocate Node to Access Control Resource Group screen, click **Next** without changing the default group.

e. Select the source storage system and click **Next**.

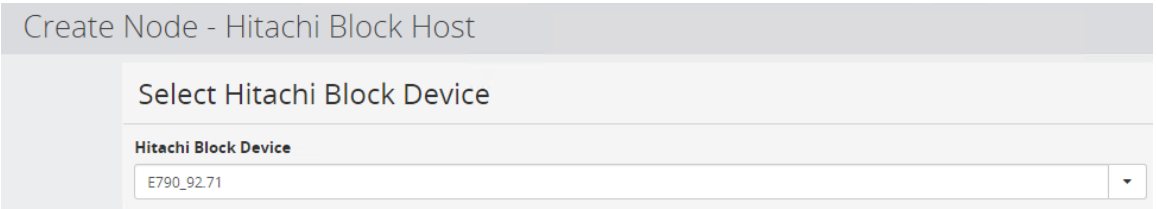

f. Filter for specific volumes by volume ID or host group. For example, we entered the IDs of the three volumes that we want to replicate with HUR. Click **Next**.

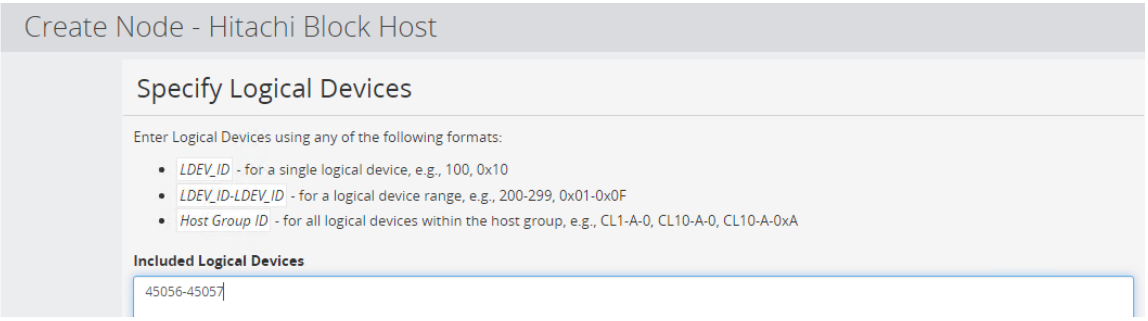

g. Review the information on the summary screen and then click **Finish**.

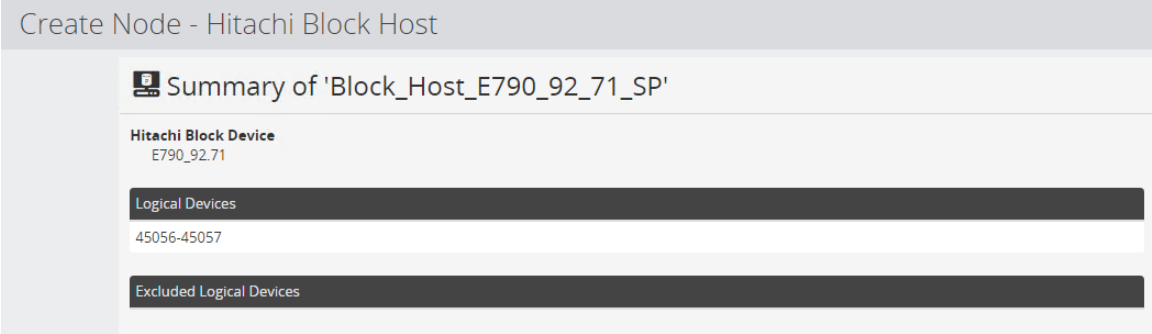

3. The next step is to create a policy.
   a. Click **Policies** and click the **+** (plus) symbol.
   b. Enter a name, description (optional), and tags (optional). Click **Next**.

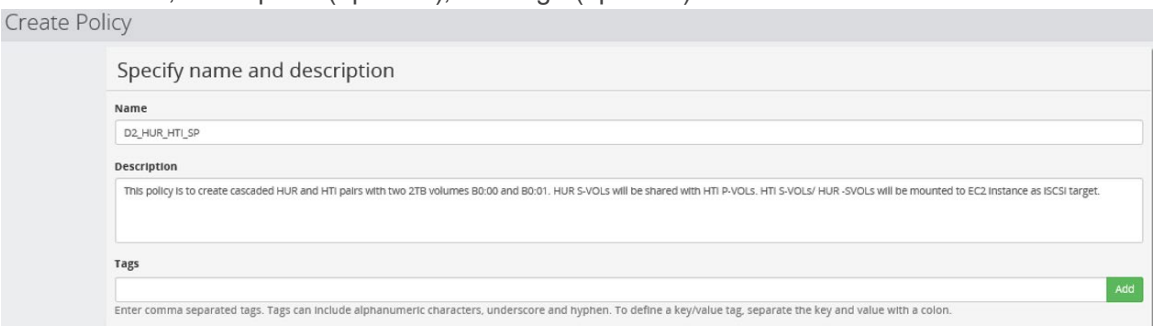

   c. In the Add One or More Classifications screen, click the **+** (plus) symbol.
   d. Click **Physical** and then click **Path**. Click **Next**.
   e. Select **Use Hitachi Block Host selections** and click **Apply**.

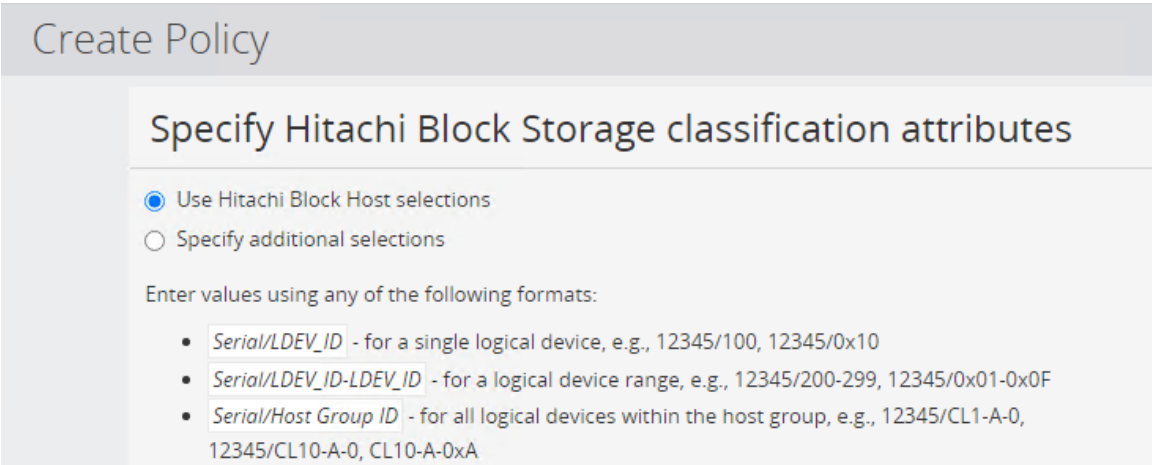

f. In the next window, as shown in the following screenshot, click **Next**.

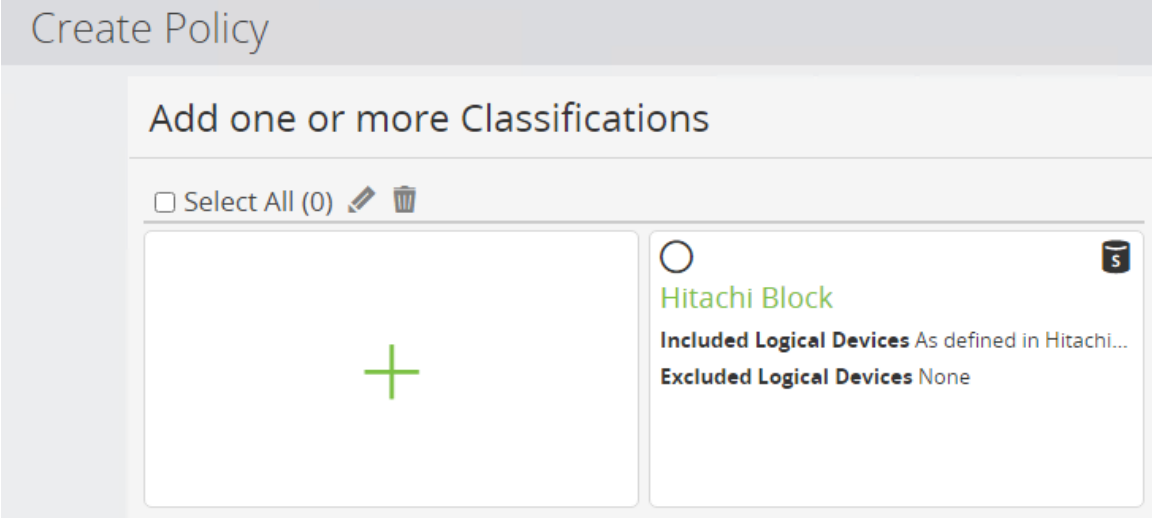

g. In the Add One or More Operations screen, click the **+** (plus) symbol.
h. Select **Replicate** and click **Next**.
i. Deselect **Quiesce…** and then click **Apply**.

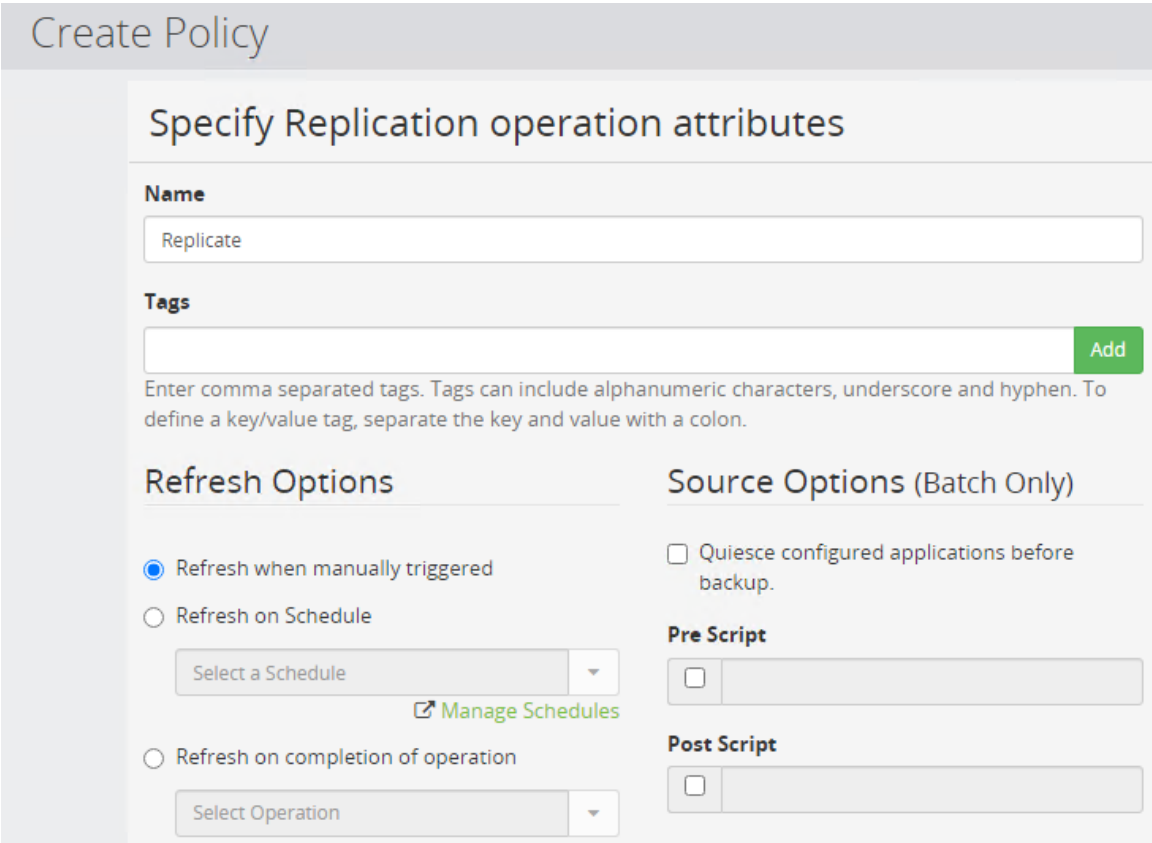

j. In the next window, as shown in the following screenshot, click **Next**.

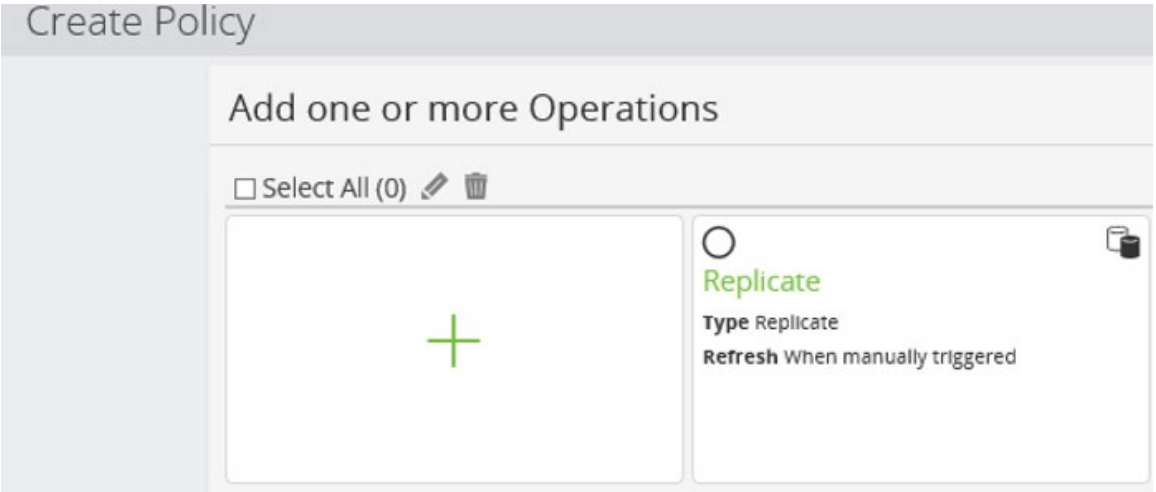

k. In the Add One or More Operations screen, click the **+** (plus) symbol again to add a snapshot operation.

l. Select **Snapshot** and click **Next**.

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

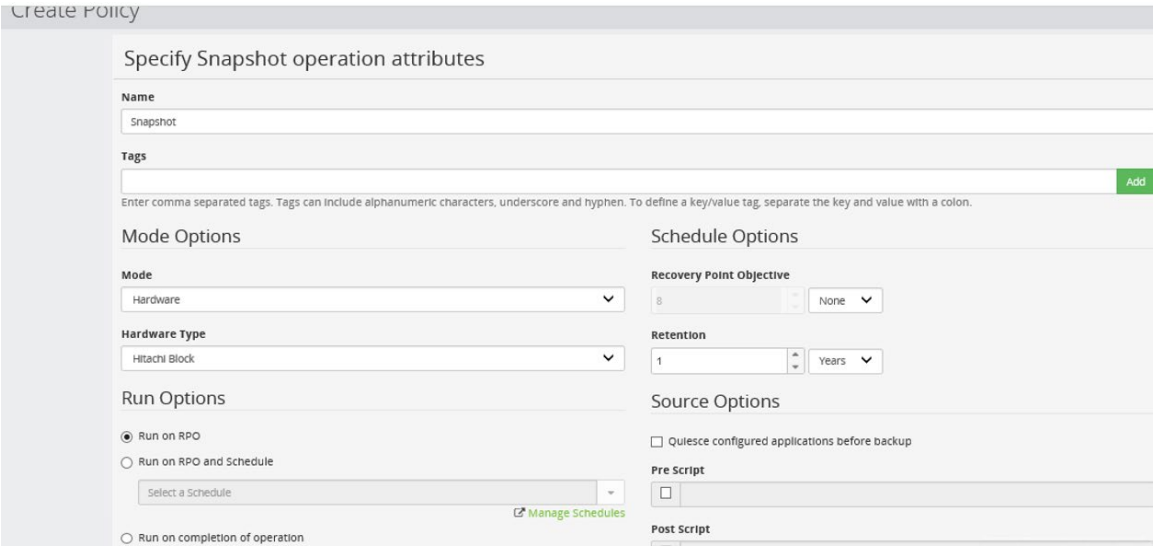

Note: **Recovery Point Objective**=**None** indicates that the snapshot will not generate automatically. It must be triggered manually.

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

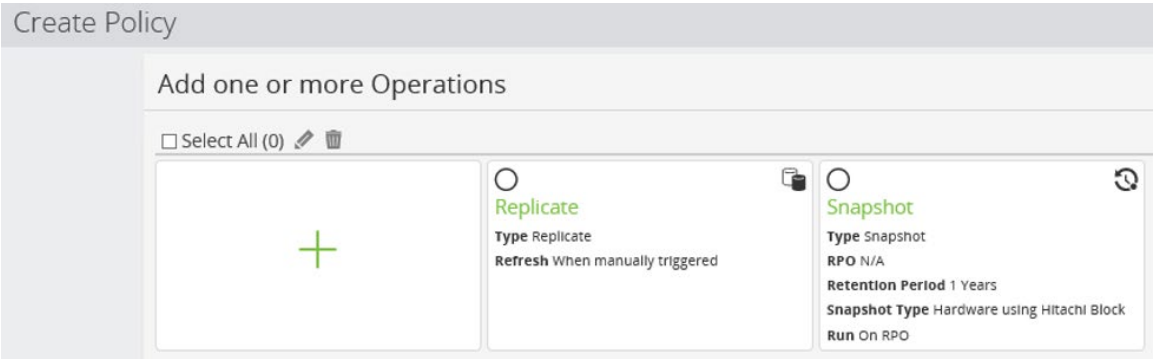

   o. Click **Finish**.

4. The final step is to create a data flow and activate it to initiate HUR 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**.

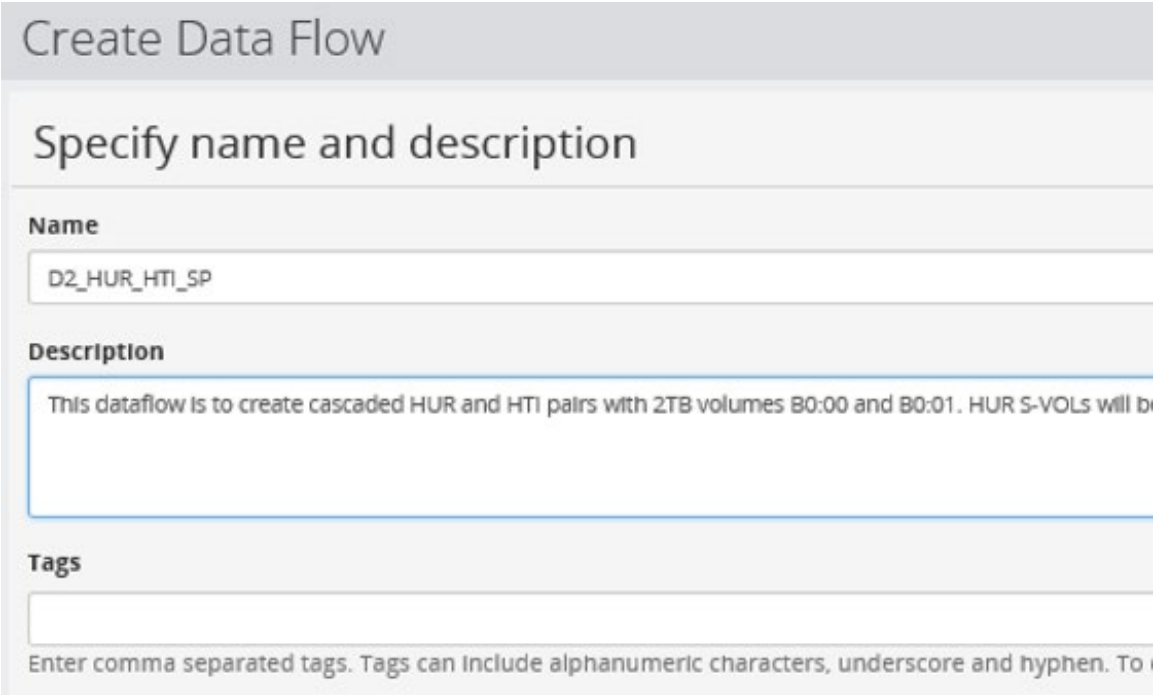

   c. Drag the source Block Host node onto the whiteboard.

   d. Drag the target Block Device node onto the whiteboard over the source Block Host node. This will create an arrow between the two nodes.

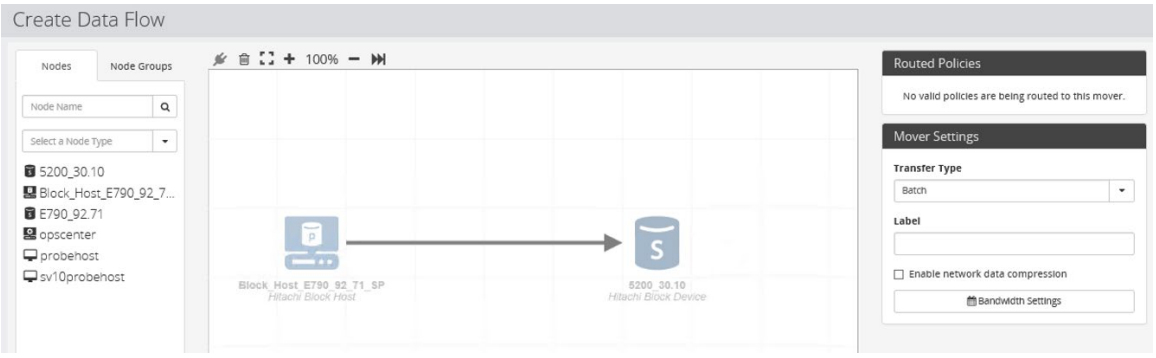

e. Select the arrow and change **Transfer Type** to **Continuous**.

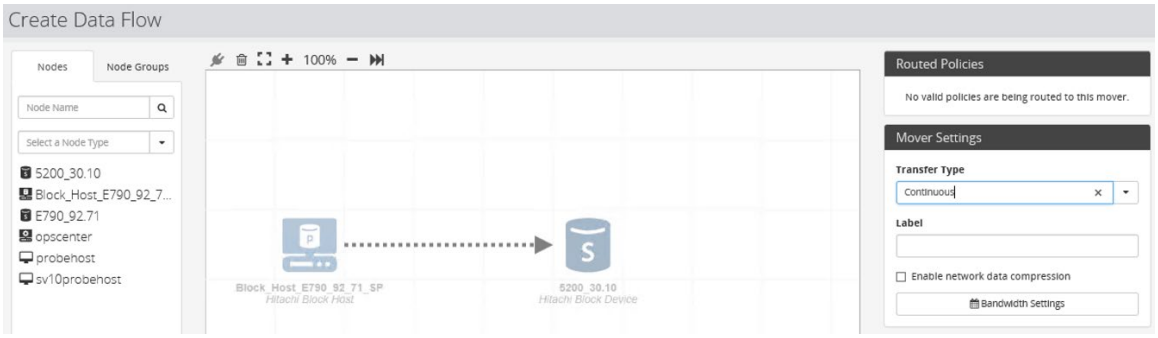

f. Select the source node again and activate the previously created policy from the Policies list.

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

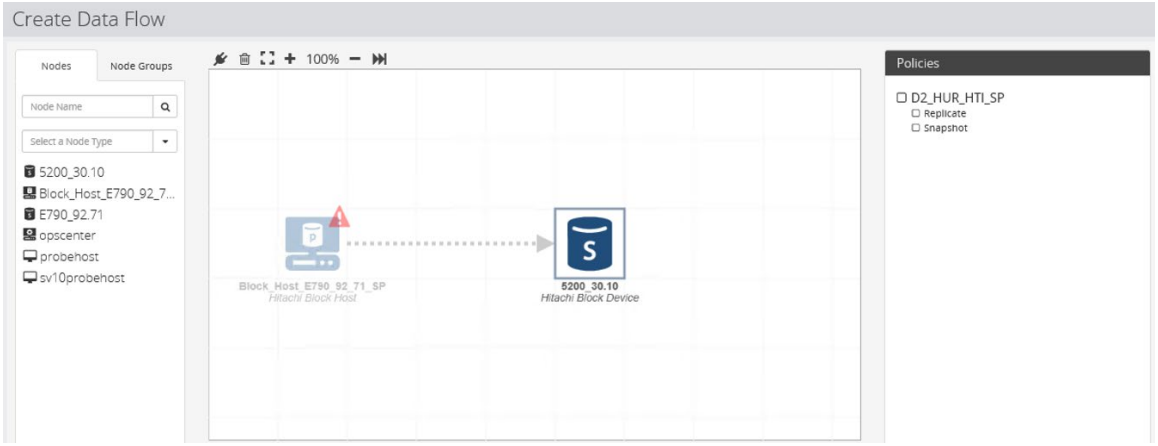

h. Select **Configure new replication** and click **Next**.

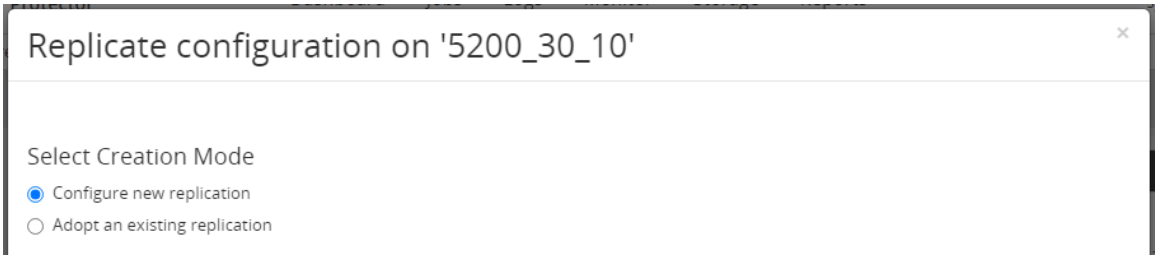

i. Select **Asynchronous Remote Clone (Universal Replicator)** and click **Next**.

j. Select the pool on the secondary storage system where the target volumes will be created and click **Next**.

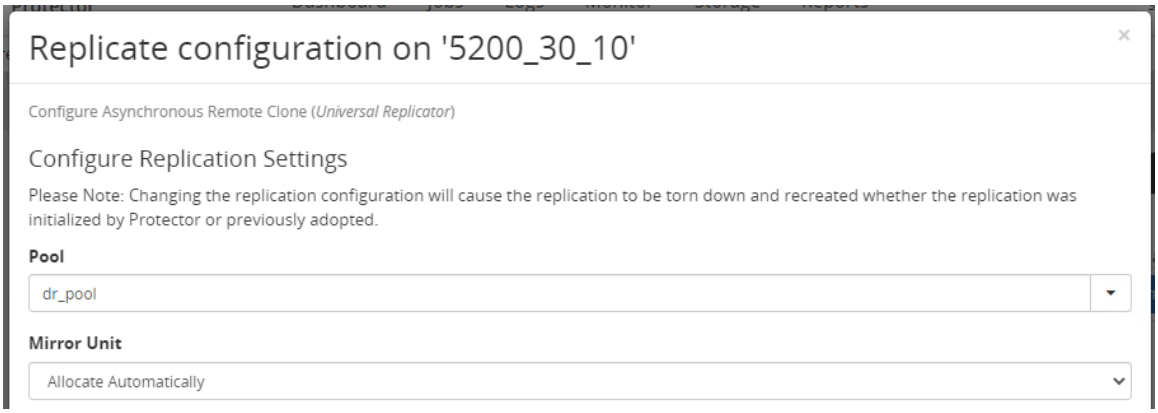

k. Select the appropriate journal option. We have not created journals so **Create new journals** was selected. Click **Next**.

l. The next screen prompts for the Source Journal Pool and Destination Journal Pool. Enter the desired journal size in **Journal Sizes**. Click **Next**.

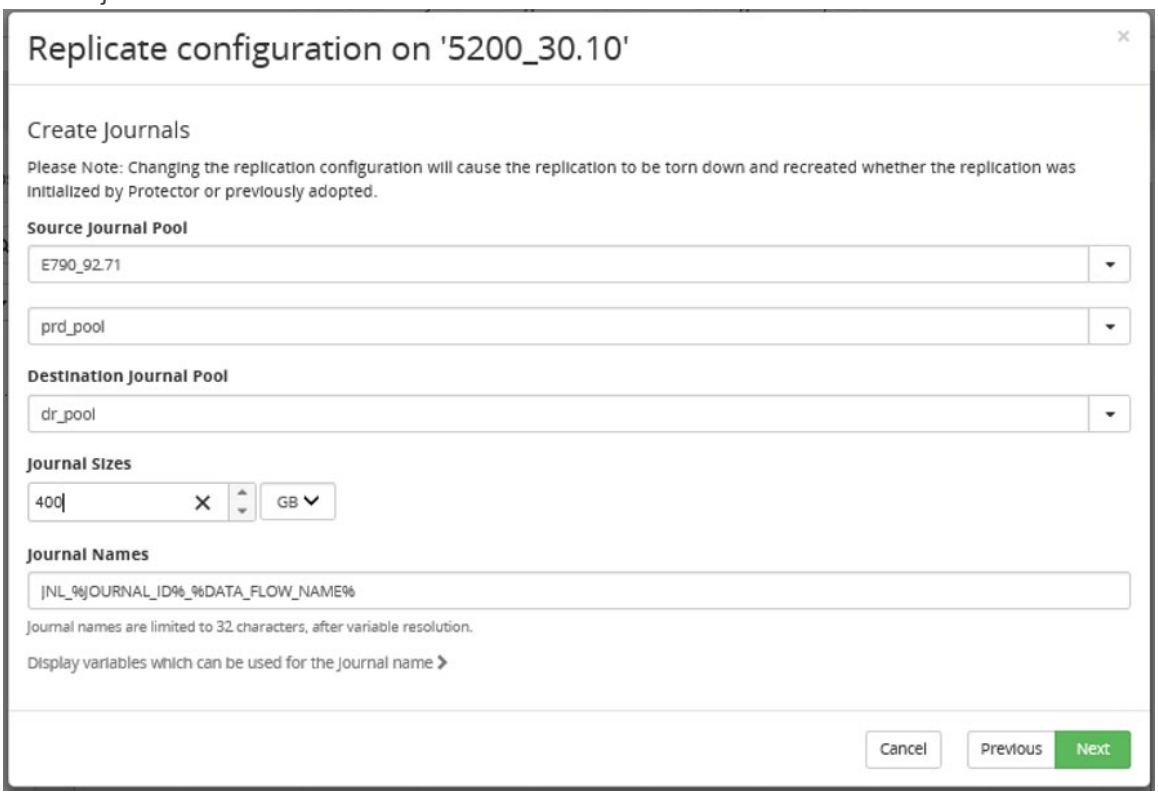

Note: Protector will create a volume in the specified pool and assign it to a journal group. This will be done on both storage systems. Journals created by Protector will have Medium (8) copy pace.

m. Select the remote replication path group and click **Next**.

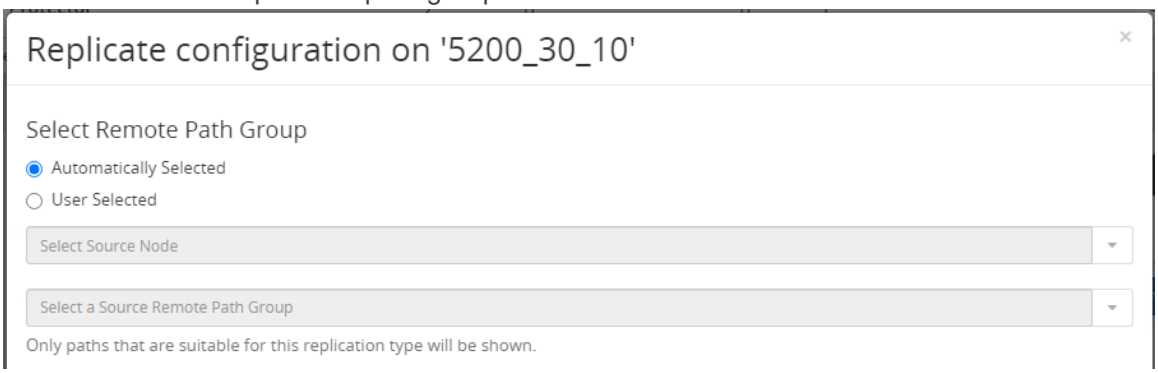

n. Select the resource group and click **Next**.

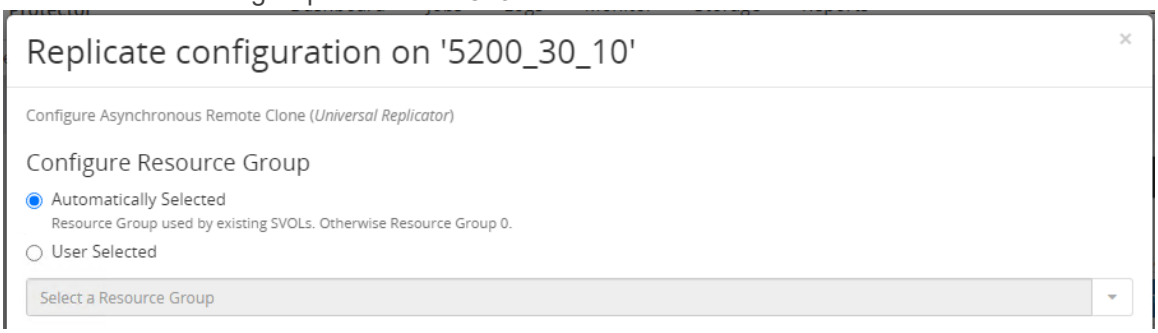

o. Specify the host groups on the secondary storage system where the HUR secondary volumes will be mapped. Click **Next**.

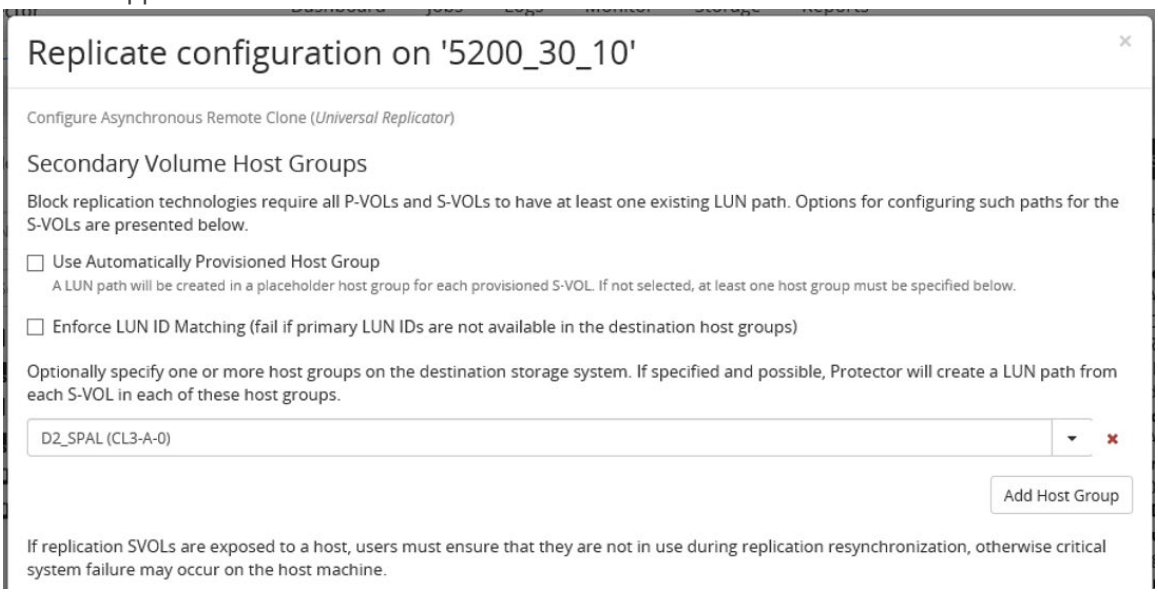

Note: The version of Protector available during testing (7.3.0) does not map volumes to an iSCSI host group. Volumes must be mapped using Ops Center Administrator.

p. Specify the naming option for the target volumes. The default option **Match Origin** will name the volumes with the same name from the source volumes, assuming they are named. Click **Next**.

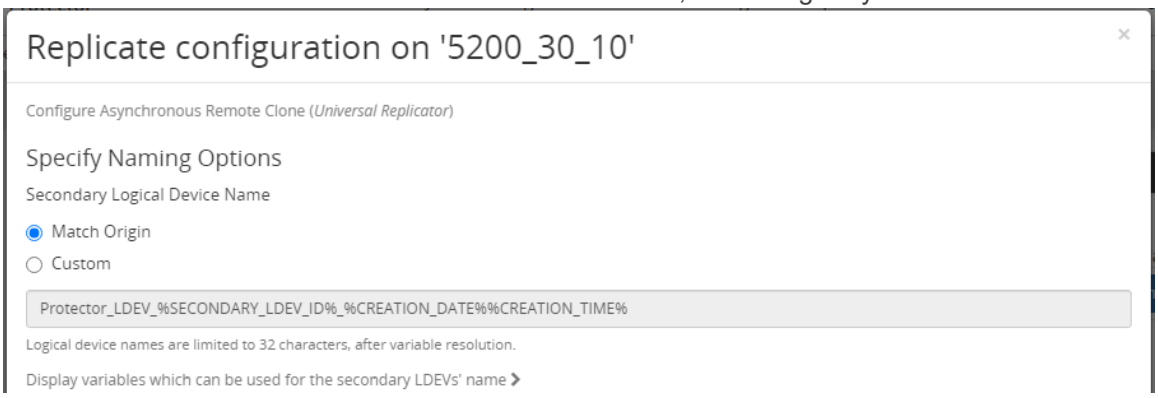

q. Review the information on the summary screen and then click **Finish**.

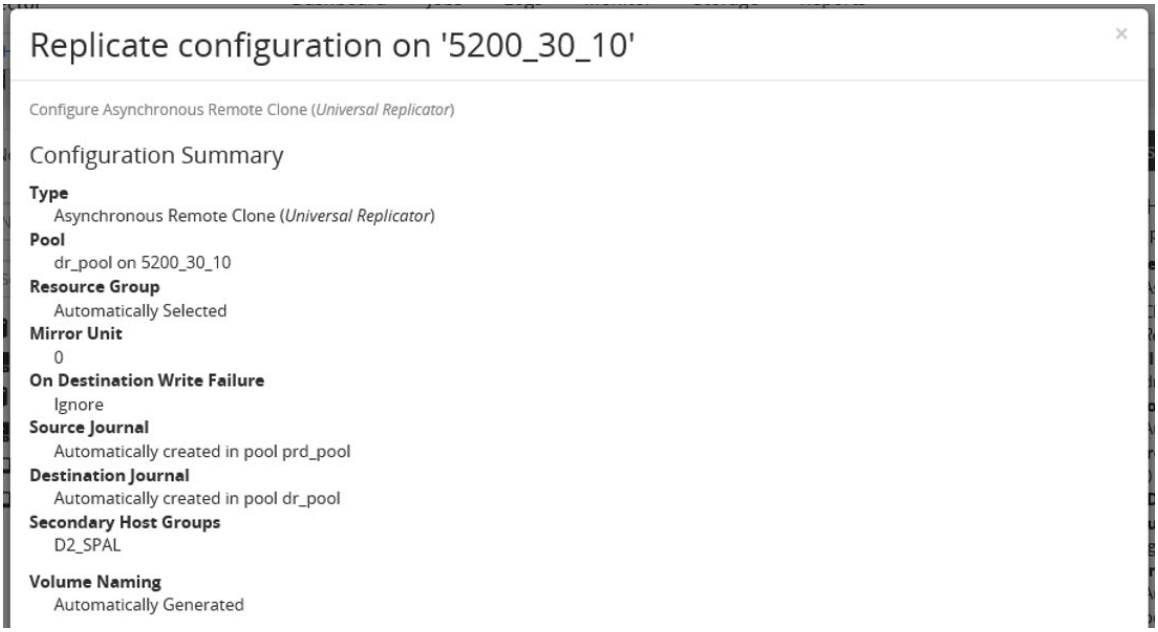

r. Click **Snapshot** to bring up the Snapshot configuration options. Select the snapshot pool. Click **Advanced Configuration** and then click **Next**.

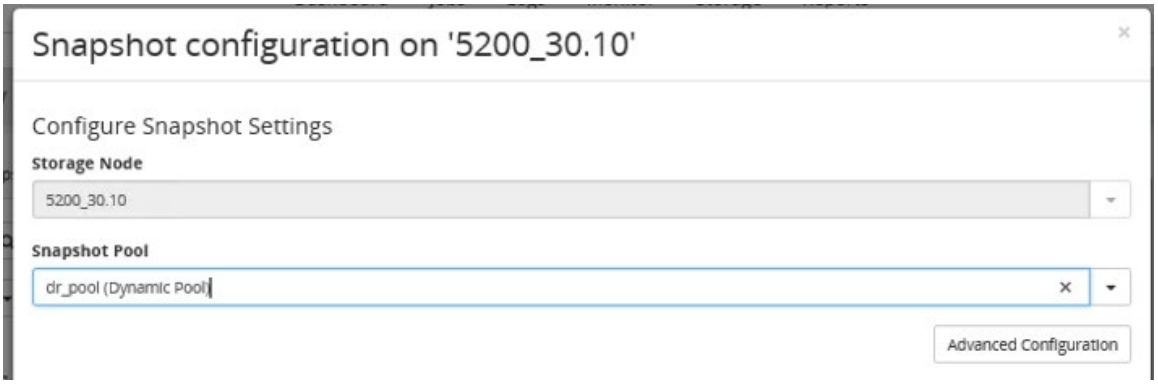

s. Select the appropriate resource group and click **Next**.

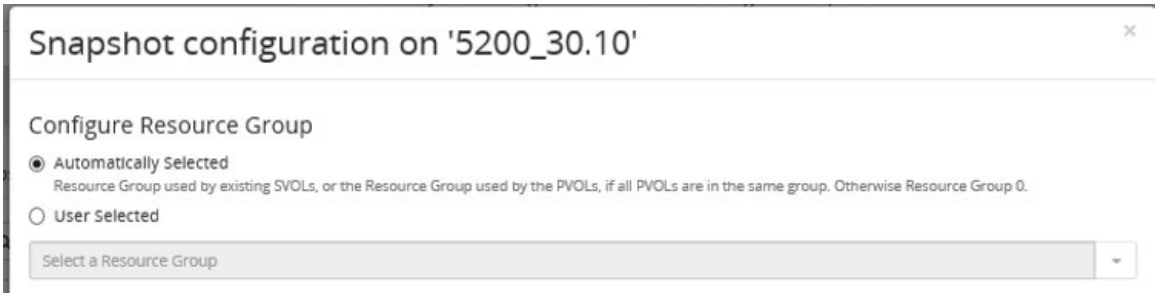

t. Select all three options: **Consistency group**, **Fully provisioned**, and **Cascade mode**. Select the appropriate cascade pool. Click **Next**.

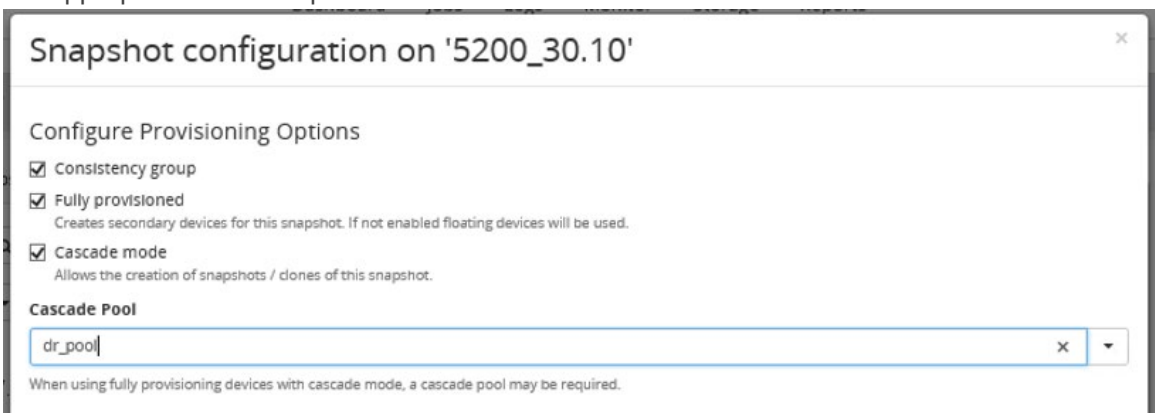

Note: **Fully provisioned** and **Cascaded mode** allow us to quickly access the HTI secondary volumes as needed.

u. In the Specify Naming Options screen, you can choose the naming option for the individual HTI secondary volumes as well as the naming option for the whole snapshot group. Click **Next**.

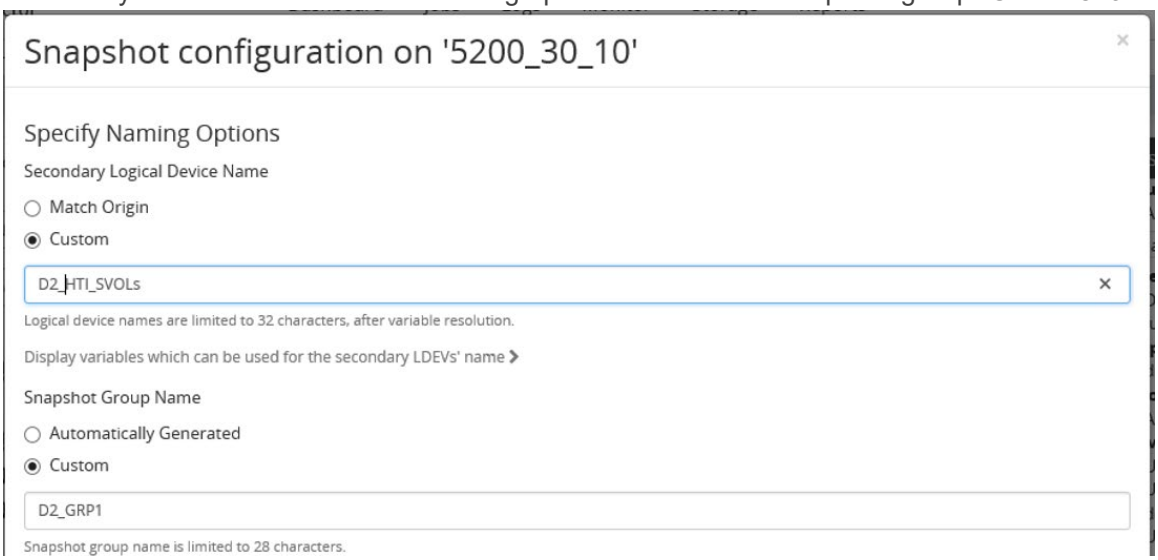

v. For **Protection Type**, select **None**. Click **Next**.

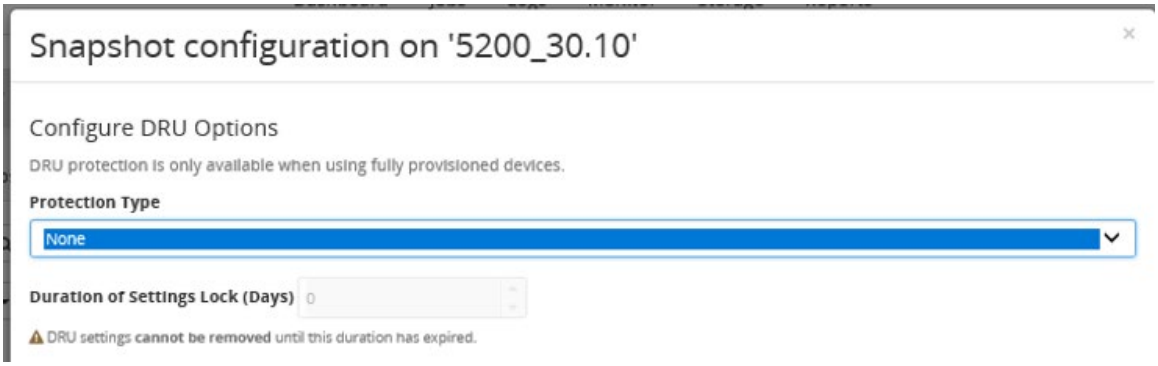

w. Review the information on the summary screen and click **Finish**.

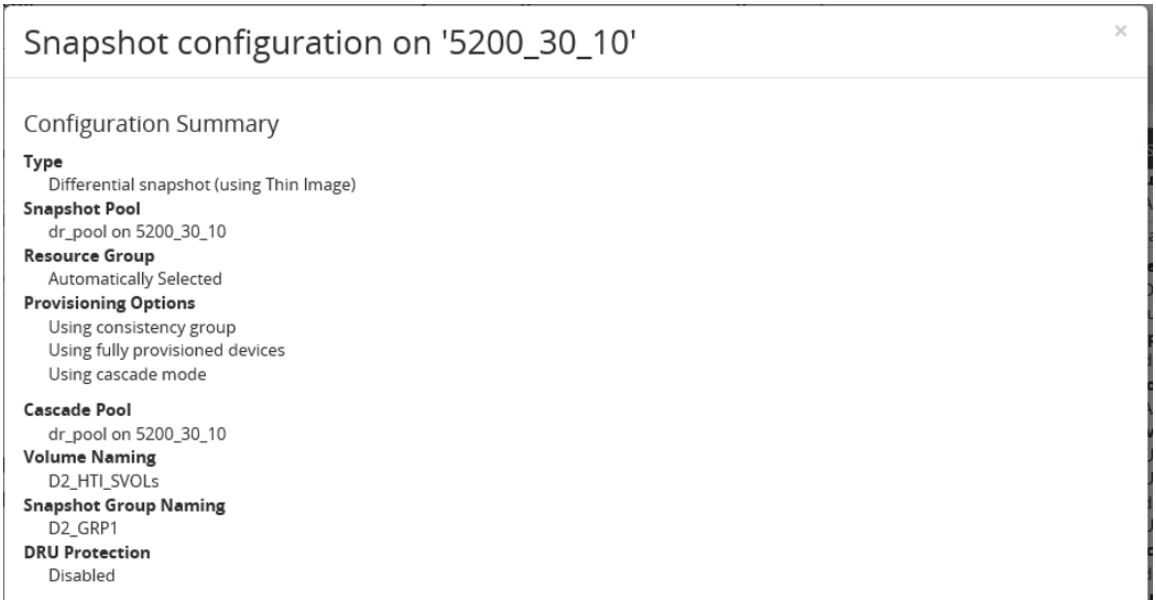

The newly created data flow is shown as follows:

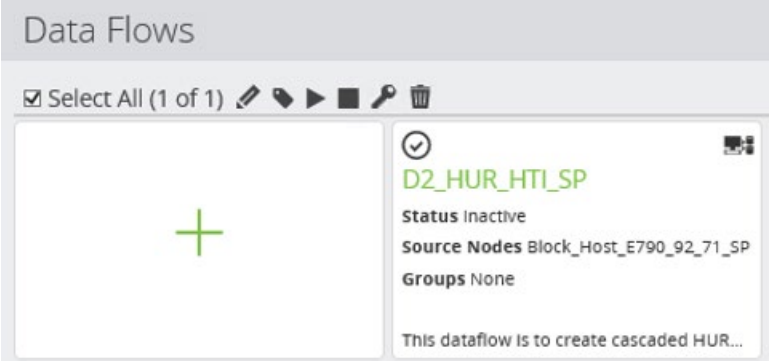

x. Initiate replication by selecting the newly created data flow and clicking **Play**.

y. Review the information in the prompt and then click **Activate**.

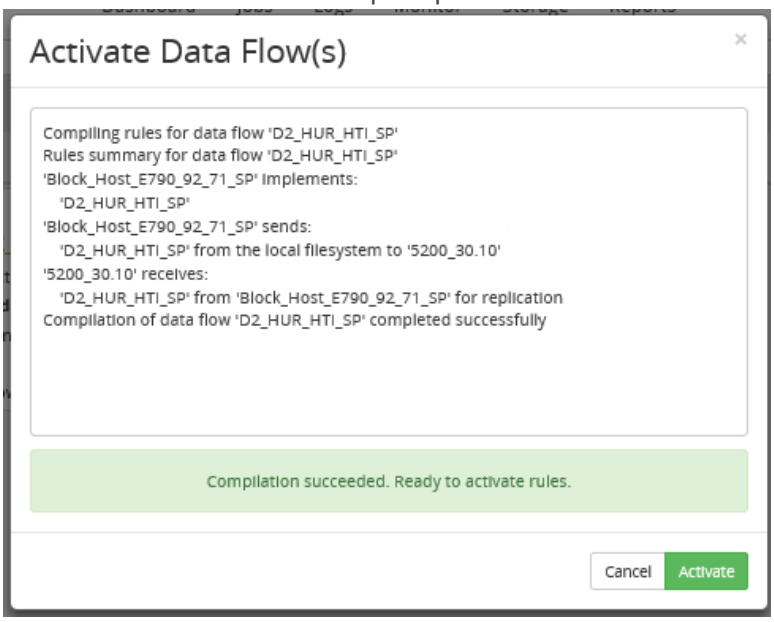

5. Track the progress by clicking on **Monitor** at the top and then drill into the data flow.

a. Click on the arrow between the nodes. The Copy Progress is displayed on the right.

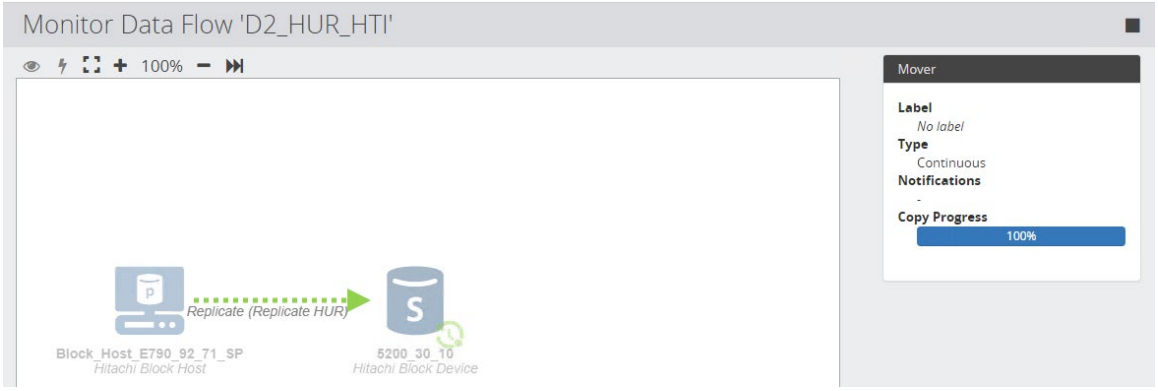

b. After the copy progress is 100%, the initial copy is done.

6. The next step is to trigger the Hitachi Thin Image snapshots.
   a. Click **Monitor** at the top and click into the previously activated data flow. The data flow will appear as follows:

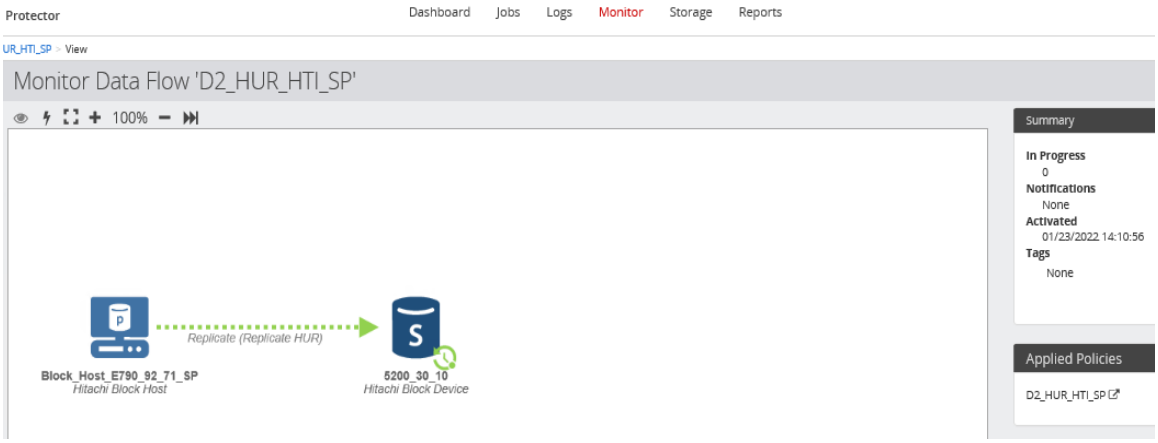

   b. Select the secondary storage system and click the **Trigger** (lightning) symbol. Click **Run Now**.

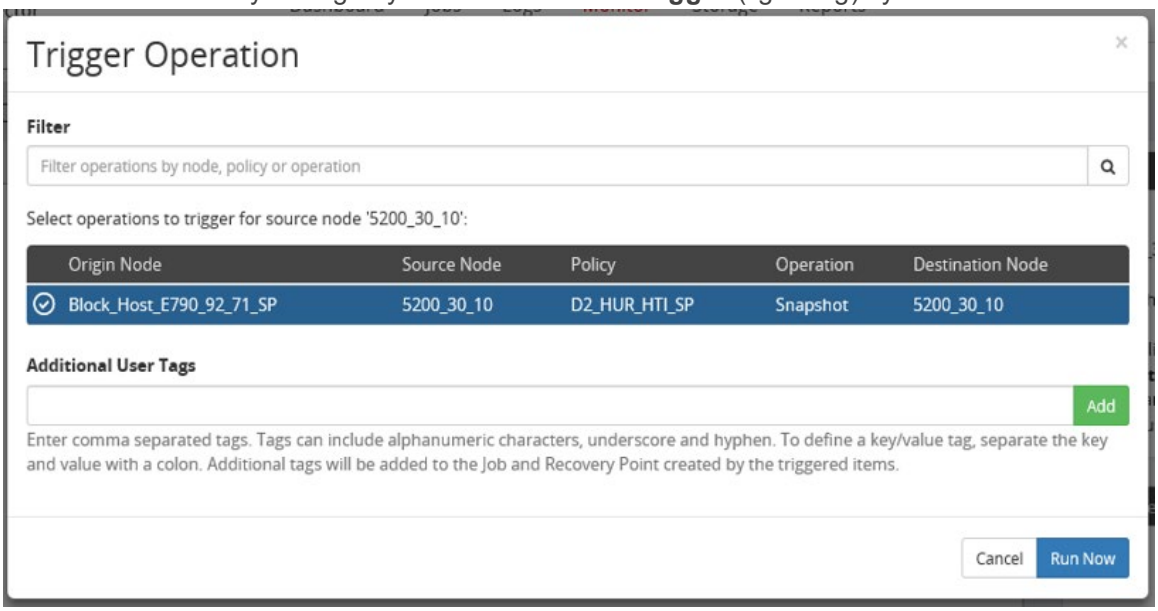

7. To view the snapshot details after the task has completed, click **Storage**, click the secondary storage system, and click **Snapshots**.
   a. Click the previously created snapshot object.

b. Click **Pairs** to display the HTI snapshot details.

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	540028	-	P-VOL	-		-/-	→	0x0006	540028	-	S-VOL	-		-/-
0xb001	540028	-	P-VOL	-		-/-	→	0x0007	540028	-	S-VOL	-		-/-

c. Because we chose **Fully provisioned**, the HTI secondary volumes are automatically mapped to random ports. For example, volume 00:06 was mounted to port CL4-B and volume 00:07 was mounted to port CL5-A.

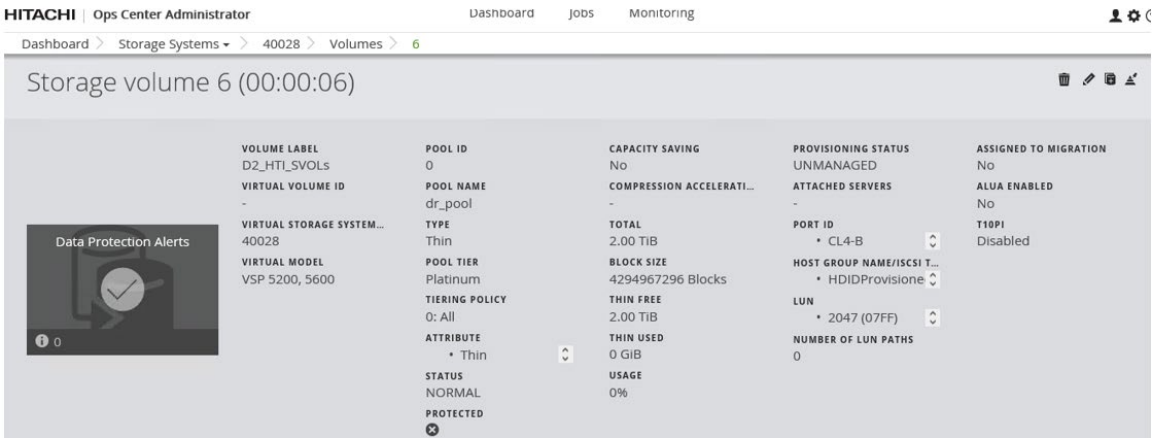

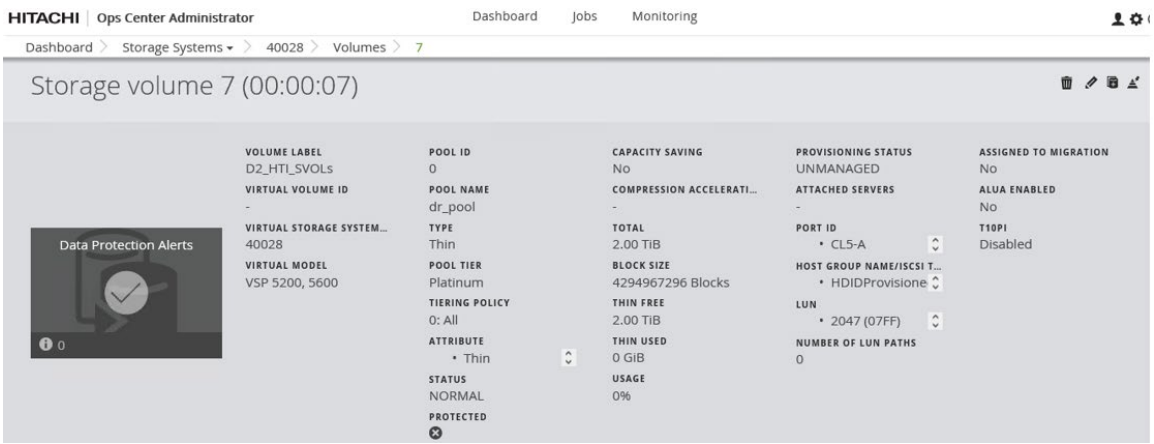

Test 3: Access Snapshots from AWS EC2

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

1. We wrote some new records into our SQL database running on the HUR primary volumes. The Update_On column indicates these records were created on the HUR P-VOLs.

Results | Messages

	Person_ID	Last_name	First_Name	Country	Update_On
1	00000	Wahl	Letha	Croatia	HUR P-VOLs
2	00001	Boucher	Kraig	Ghana	HUR P-VOLs
3	00002	Ludwig	Harlan	Guyana	HUR P-VOLs
4	00003	Craven	Brad	Sri Lanka	HUR P-VOLs
5	00004	Sizemore	Enrique	Sri Lanka	HUR P-VOLs
6	00005	Flowers	Bennie	Swaziland	HUR P-VOLs
7	00006	Conway	Francis	Iraq	HUR P-VOLs
8	00007	Waite	Nickie	Nicaragua	HUR P-VOLs
9	00008	Skaggs	Boyd	Malawi	HUR P-VOLs
10	00009	Peacock	Alonzo	Uruguay	HUR P-VOLs
11	00010	Hwang	Reed	Australia	HUR P-VOLs

Query executed successfully. D2SQLDB\D2S

2. Trigger snapshots using the instructions provided in section Test 2.
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 (00:06 and 00:07 in our case) and click **Attach Volumes**.

ashboard > Storage Systems > 40028 > Volumes

Select All 2 selected

show 200 items « Previous 1 Next »

SELECT	ID	VIRTUAL ID...	S/N	VIRTUAL S/N...	LABEL	POOL ID	ATTRIBUTE	STATUS	CAPACITY SA...	TOTAL	USAGE	PROVISIONI...
	0 (00:00:00)	0 (00:00:00)	40028	40028	backup_ta...	0	Thin	Normal	No	5.00 TiB	0%	Attached
	3 (00:00:03)	3 (00:00:03)	40028	40028	JNL_0_D2_...	0	Thin	Normal	No	400.00 GiB	0%	Unattached
	6 (00:00:06)	6 (00:00:06)	40028	40028	D2_HTI_S...	0	Thin	Normal	No	2.00 TiB	0%	Unmanag...
	7 (00:00:07)	7 (00:00:07)	40028	40028	D2_HTI_S...	0	Thin	Normal	No	2.00 TiB	0%	Unmanag...

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

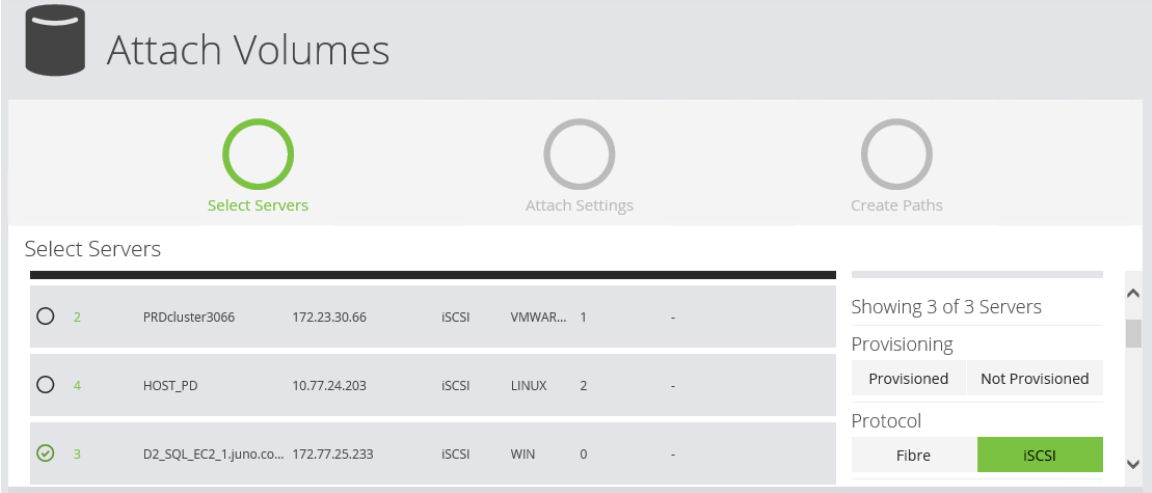

e. Change the Host Mode options if required and click **Next**.

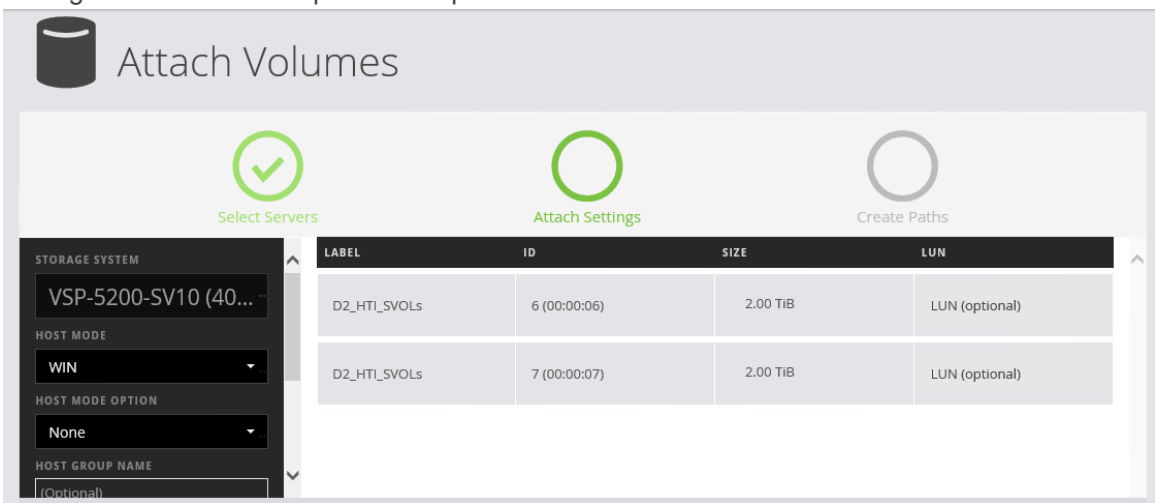

f. Draw a line between the server on the left and the storage ports on the right and click **Submit**.

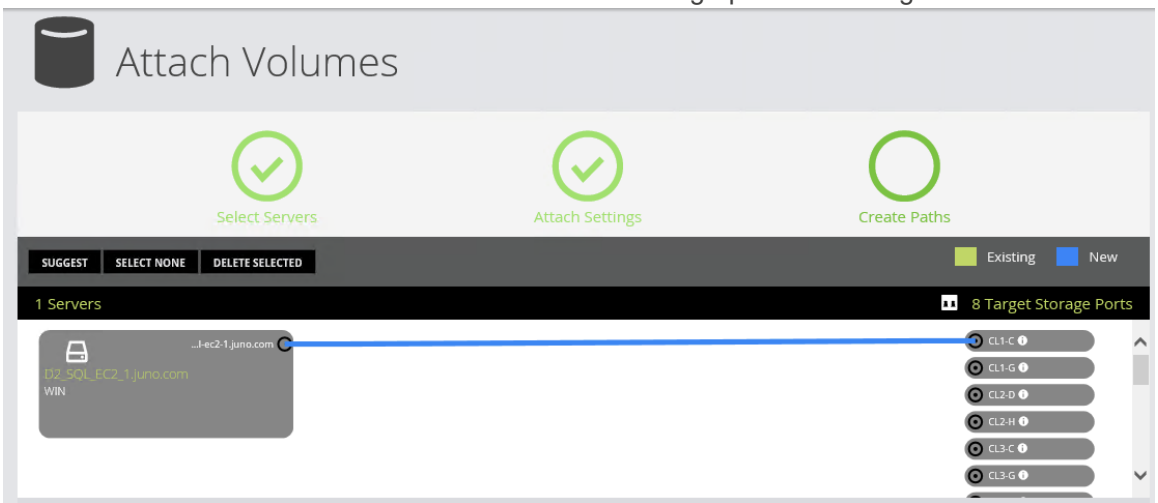

This will create a host group with the EC2 virtual machine initiator name on storage port CL1-C and map the two volumes to this host group.

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.

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

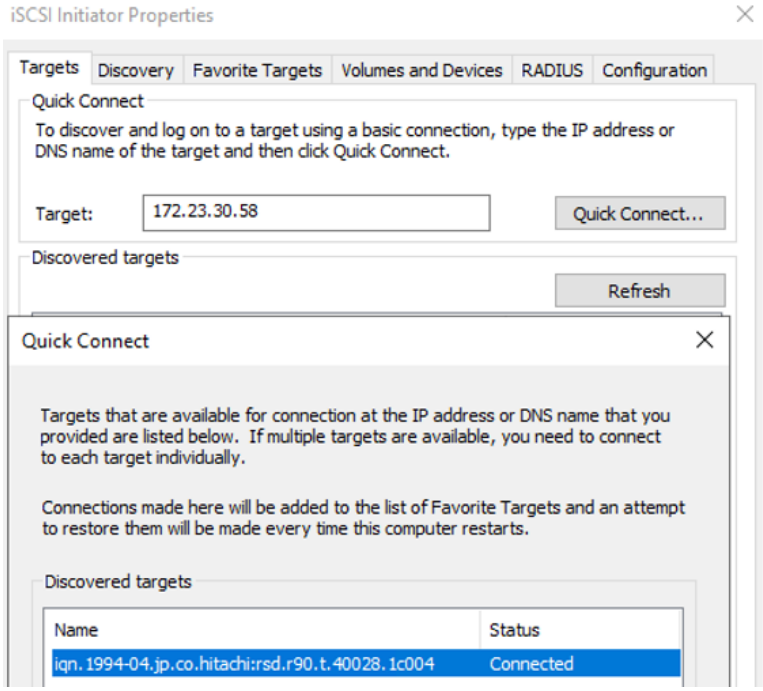

c. Click **Done**.

d. Open **Computer Management > Disk Management**. The two volumes should be listed. Right click on them and select **Online**.

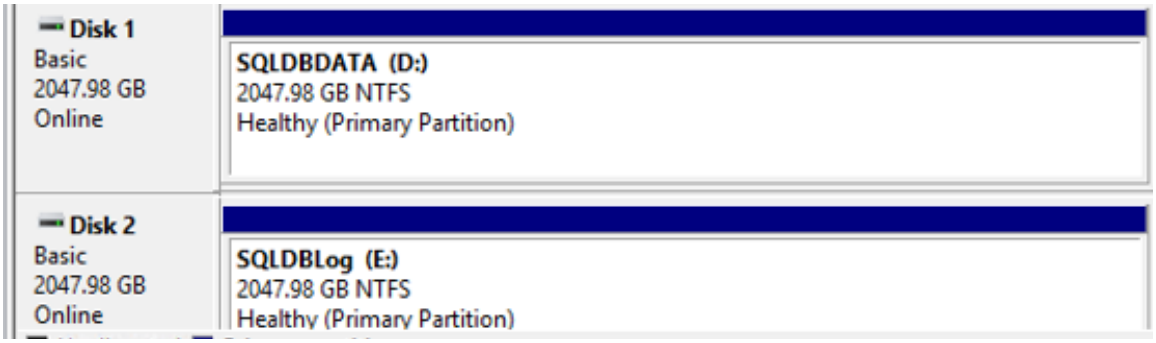

5. Verify that previously written records were replicated correctly and write new records to the SQL copy residing on the HTI snapshots.

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

b. If the database files are not visible in the Attach Databases dialog, set the permissions on the HTI snapshot drives to match the local SQL Server instance directory as follows:

```
PS C:\Users\Administrator.JUNO> Get-Acl "C:\Program Files\Microsoft SQL Server\MSSQL15.EC2SQL1\MSSQL\DATA"| Set-Acl "D:\
MSSQL15.D2SQL\MSSQL\DATA"
PS C:\Users\Administrator.JUNO>
PS C:\Users\Administrator.JUNO>
PS C:\Users\Administrator.JUNO> Get-Acl "C:\Program Files\Microsoft SQL Server\MSSQL15.EC2SQL1\MSSQL\DATA"| Set-Acl "E:\
LOG"
PS C:\Users\Administrator.JUNO>
```

c. Verify that previously written records are present.

Results | Messages

	Person_ID	Last_name	First_Name	Country	Update_On
1	00000	Wahl	Letha	Croatia	HUR P-VOLs
2	00001	Boucher	Kraig	Ghana	HUR P-VOLs
3	00002	Ludwig	Harlan	Guyana	HUR P-VOLs
4	00003	Craven	Brad	Sri Lanka	HUR P-VOLs
5	00004	Sizemore	Enrique	Sri Lanka	HUR P-VOLs
6	00005	Flowers	Bennie	Swaziland	HUR P-VOLs
7	00006	Conway	Francis	Iraq	HUR P-VOLs
8	00007	Waite	Nickie	Nicaragua	HUR P-VOLs
9	00008	Skaggs	Boyd	Malawi	HUR P-VOLs
10	00009	Peacock	Alonzo	Uruguay	HUR P-VOLs
11	00010	Hwang	Reed	Australia	HUR P-VOLs

Query executed successfully.

d. Create a new database on the same HTI snapshot drives. The new database is completely independent from the database running on the HUR primary volumes at the primary site.

Database Properties - D2_Test1.3_HTI_1stGen_DB

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

Script | Help

Database name: D2_Test1.3_HTI_1stGen_DB

Owner: JUNO\Administrator

Use full-text indexing

Database files:

Logical Name	owth / Maxsize	Path	File Na
D2_Test1.3...	MB, Unlimited	D:\MSSQL15.D2SQL\MSSQL\DATA	D2_T
D2_Test1.3...	MB, Limited to 2...	E:\LOG	D2_T

Data is written to the new database.

Results | Messages

	Student_ID	Student_Name	Student_Address	Update_On
1	00000	Letha Bolt	240 Hidden Front Way	HTI_1stGen_S-VOLs
2	00001	Vanita Kelleher	2948 Stonewood Pkwy	HTI_1stGen_S-VOLs
3	00002	Marcos Abraham	44 West Church Ct	HTI_1stGen_S-VOLs
4	00003	Miyoko Mckinney	3778 Edgewood Highway	HTI_1stGen_S-VOLs
5	00004	Adolph Francisco	12 1st Circle	HTI_1stGen_S-VOLs
6	00005	Oma Lawler	3570 Flintwood Parkway	HTI_1stGen_S-VOLs
7	00006	Felix Alba	1790 SW Church Avenue	HTI_1stGen_S-VOLs
8	00007	Eusebia Noland	2014 Social Road	HTI_1stGen_S-VOLs

Note: These 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.

6. Delete the HTI snapshots.
   a. In 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**.

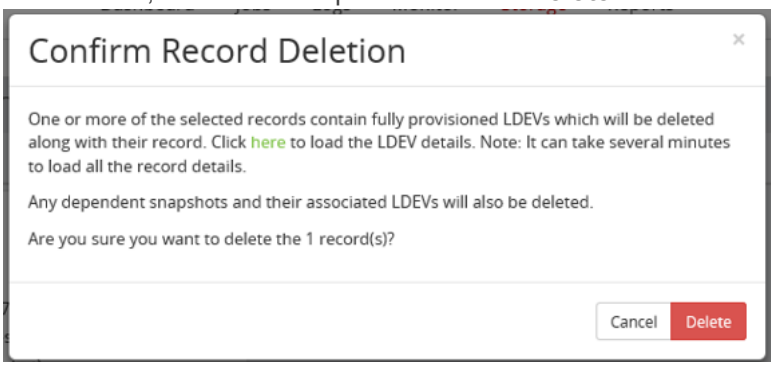

Test 4: Multigeneration HTI Snapshots

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. Generate new records on the HUR primary volumes by creating a database named D2_Test1.4_DB1 and a table named Staff_Table. The Update_On column indicates these records were created on the HUR P-VOLs.

Results | Messages

	Timestamp	Staff_id	First_Name	Last_Name	Email	Update_On
1	2022-01-27 09:45:00.0000000	00000	Isobel	Wahl	Bolt@nowhere.com	1_HUR_P-VOLs
2	2022-01-27 09:45:00.0000000	00001	Magdalen	Boucher	Morehead@example.com	1_HUR_P-VOLs
3	2022-01-27 09:45:00.0000000	00002	Rosio	Ludwig	qpzre411@example.com	1_HUR_P-VOLs
4	2022-01-27 09:45:00.0000000	00003	Enriqueta	Craven	Bray162@example.com	1_HUR_P-VOLs
5	2022-01-27 09:45:00.0000000	00004	Iva	Sizemore	Sidney_Agnew3@nowhere.com	1_HUR_P-VOLs
6	2022-01-27 09:45:00.0000000	00005	Kallie	Flowers	xmhz1900@nowhere.com	1_HUR_P-VOLs
7	2022-01-27 09:45:00.0000000	00006	Breanne	Conway	Monroe.Rinehart@nowhere.com	1_HUR_P-VOLs
8	2022-01-27 09:45:00.0000000	00007	Letisha	Waite	Gutierrez73@example.com	1_HUR_P-VOLs
9	2022-01-27 09:45:00.0000000	00008	Taren	Skaggs	FelipaAbernathy@nowhere.com	1_HUR_P-VOLs
10	2022-01-27 09:45:00.0000000	00009	Zulma	Peacock	Adler@example.com	1_HUR_P-VOLs
11	2022-01-27 09:45:00.0000000	00010	Kalyn	Hwang	uadm55@example.com	1_HUR_P-VOLs

Query executed successfully. D2SQLDB\D2SC

2. Create the first generation of HTI snapshots.
   a. Trigger snapshots as described in section Test 2.
   b. Map the snapshot volumes as described in section Test 3.
   c. In EC2, bring the drives online and attach the database (D2_Test1.4_DB1).
   d. Verify that previously written records are present.

Results | Messages

	Timestamp	Staff_id	First_Name	Last_Name	Email	Update_On
1	2022-01-27 09:45:00.0000000	00000	Isobel	Wahl	Bolt@nowhere.com	1_HUR_P-VOLs
2	2022-01-27 09:45:00.0000000	00001	Magdalen	Boucher	Morehead@example.com	1_HUR_P-VOLs
3	2022-01-27 09:45:00.0000000	00002	Rosio	Ludwig	qpzre411@example.com	1_HUR_P-VOLs
4	2022-01-27 09:45:00.0000000	00003	Enriqueta	Craven	Bray162@example.com	1_HUR_P-VOLs
5	2022-01-27 09:45:00.0000000	00004	Iva	Sizemore	Sidney_Agnew3@nowhere.com	1_HUR_P-VOLs
6	2022-01-27 09:45:00.0000000	00005	Kallie	Flowers	xmhz1900@nowhere.com	1_HUR_P-VOLs
7	2022-01-27 09:45:00.0000000	00006	Breanne	Conway	Monroe.Rinehart@nowhere.com	1_HUR_P-VOLs
8	2022-01-27 09:45:00.0000000	00007	Letisha	Waite	Gutierrez73@example.com	1_HUR_P-VOLs
9	2022-01-27 09:45:00.0000000	00008	Taren	Skaggs	FelipaAbernathy@nowhere.com	1_HUR_P-VOLs
10	2022-01-27 09:45:00.0000000	00009	Zulma	Peacock	Adler@example.com	1_HUR_P-VOLs
11	2022-01-27 09:45:00.0000000	00010	Kalyn	Hwang	uadm55@example.com	1_HUR_P-VOLs

Query executed successfully. D2_SQL_EC2_1\EC2S

e. Create a new database, D2_Test1.4_DB2, on the same HTI snapshot drives.

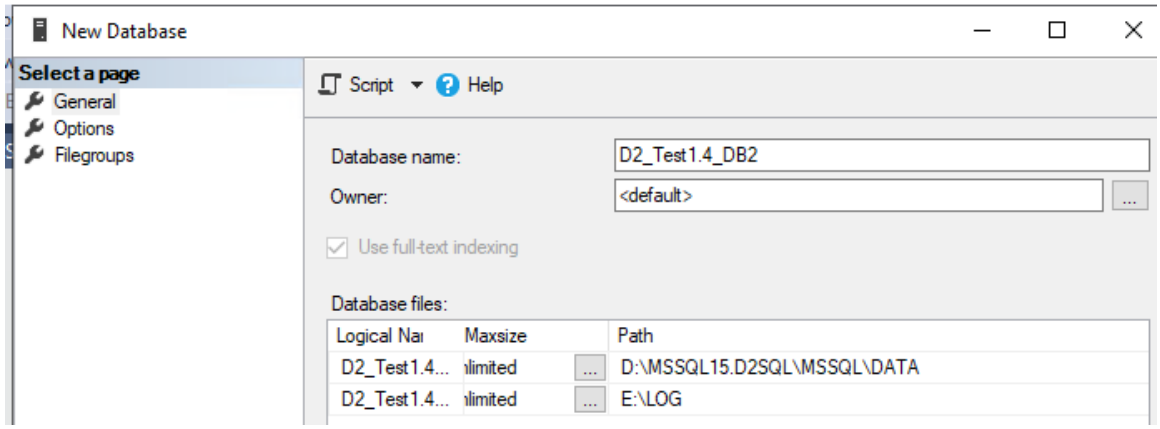

> Note: Because the new database, D2_Test1.4_DB2, was created on the first generation snapshots, it will not be present in subsequent snapshot generations. Snapshot generations are independent, which is part of the design.

3. Generate new records on the HUR primary volumes by creating a database named D2_Test1.4_DB3 and a table named Product_Table. The Update_On column indicates these records were created on the HUR P-VOLs. This database will be used to validate the second generation of HTI snapshots.

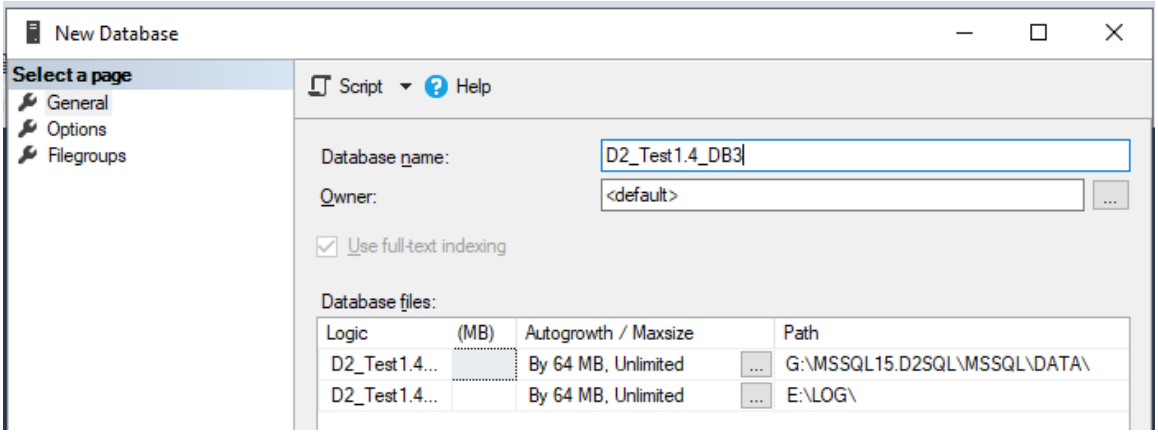

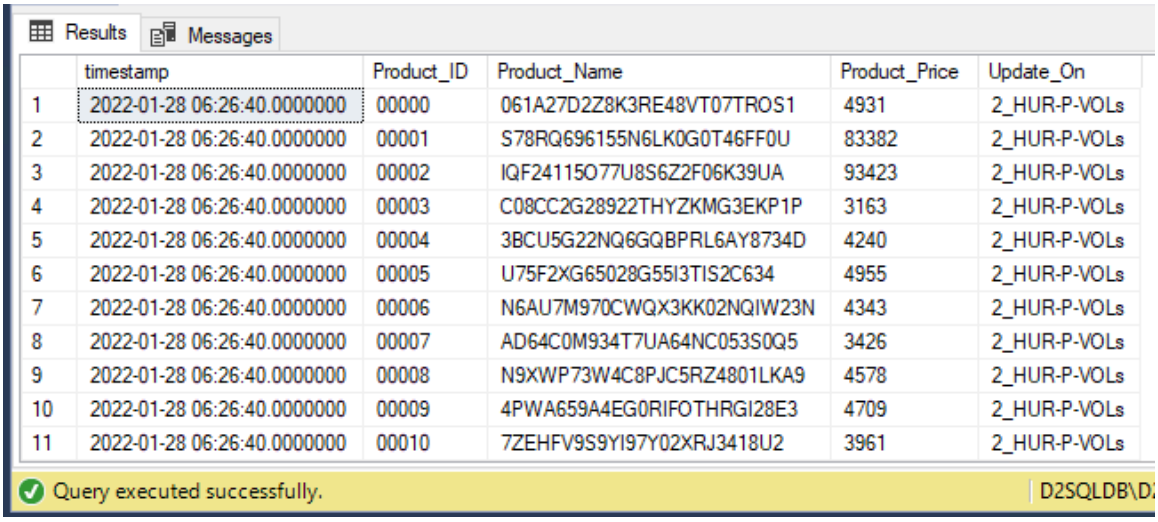

	timestamp	Product_ID	Product_Name	Product_Price	Update_On
1	2022-01-28 06:26:40.0000000	00000	061A27D2Z8K3RE48VT07TROS1	4931	2_HUR-P-VOLs
2	2022-01-28 06:26:40.0000000	00001	S78RQ696155N6LK0G0T46FF0U	83382	2_HUR-P-VOLs
3	2022-01-28 06:26:40.0000000	00002	IQF24115O77U8S6Z2F06K39UA	93423	2_HUR-P-VOLs
4	2022-01-28 06:26:40.0000000	00003	C08CC2G28922THYZKMG3EKP1P	3163	2_HUR-P-VOLs
5	2022-01-28 06:26:40.0000000	00004	3BCU5G22NQ6GQBPRL6AY8734D	4240	2_HUR-P-VOLs
6	2022-01-28 06:26:40.0000000	00005	U75F2XG65028G55I3TIS2C634	4955	2_HUR-P-VOLs
7	2022-01-28 06:26:40.0000000	00006	N6AU7M970CWQX3KK02NQIW23N	4343	2_HUR-P-VOLs
8	2022-01-28 06:26:40.0000000	00007	AD64C0M934T7UA64NC053S0Q5	3426	2_HUR-P-VOLs
9	2022-01-28 06:26:40.0000000	00008	N9XWP73W4C8PJC5RZ4801LKA9	4578	2_HUR-P-VOLs
10	2022-01-28 06:26:40.0000000	00009	4PWA659A4EG0RIFOTHRGI28E3	4709	2_HUR-P-VOLs
11	2022-01-28 06:26:40.0000000	00010	7ZEHFV9S9YI97Y02XRJ3418U2	3961	2_HUR-P-VOLs

4. Create the second generation of HTI snapshots.

a. Trigger snapshots as described in section Test 2. After the operation is done, two snapshot relationships are visible.

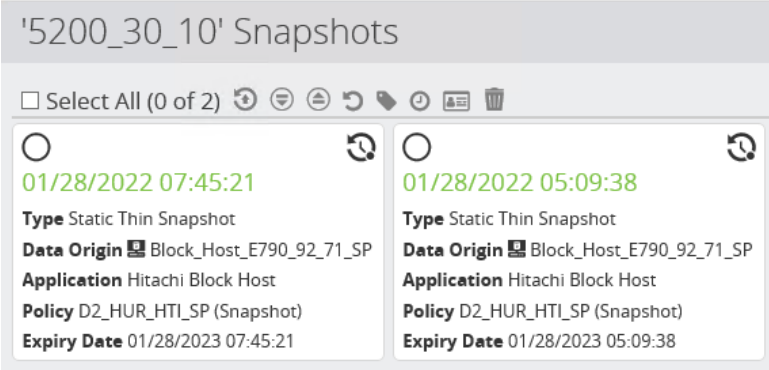

b. In the second snapshot generation, the HTI primary volumes remain the same, but the HTI secondary volumes are different (from the first snapshot generation HTI secondary volumes).

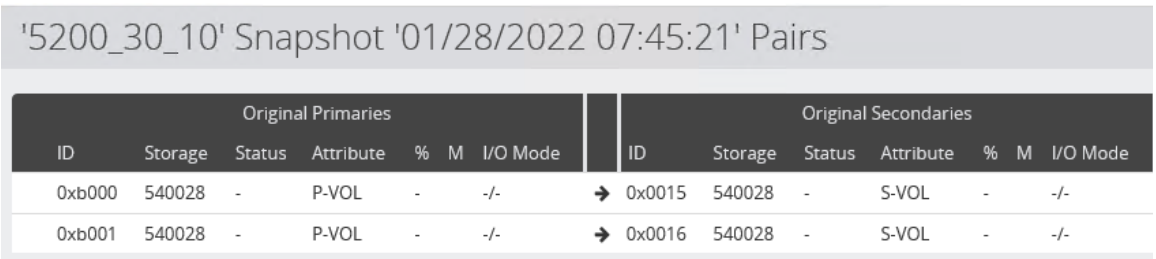

'5200_30_10' Snapshot '01/28/2022 07:45:21' Pairs

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	540028	-	P-VOL	-		-/-	→	0x0015	540028	-	S-VOL	-		-/-
0xb001	540028	-	P-VOL	-		-/-	→	0x0016	540028	-	S-VOL	-		-/-

5. Map the second generation snapshots, as described in section Test 3, to a second EC2 virtual machine.
6. In EC2, bring the drives online and attach the databases: D2_Test1.4_DB1 and D2_Test1.4_DB3. Notice that the database D2_Test1.4_DB2, which was created with the first generation snapshots, is not present.

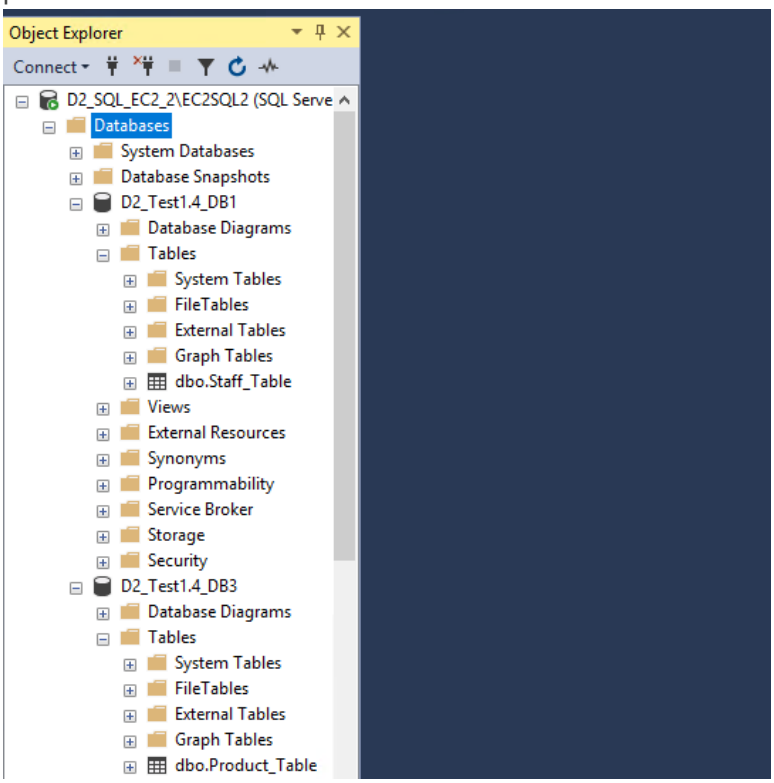

7. Check the content of the databases.

Results | Messages

	Timestamp	Staff_id	First_Name	Last_Name	Email	Update_On
1	2022-01-27 09:45:00.0000000	00000	Isobel	Wahl	Bolt@nowhere.com	1_HUR_P-VOLs
2	2022-01-27 09:45:00.0000000	00001	Magdalen	Boucher	Morehead@example.com	1_HUR_P-VOLs
3	2022-01-27 09:45:00.0000000	00002	Rosio	Ludwig	qpzre411@example.com	1_HUR_P-VOLs
4	2022-01-27 09:45:00.0000000	00003	Enriqueta	Craven	Bray162@example.com	1_HUR_P-VOLs
5	2022-01-27 09:45:00.0000000	00004	Iva	Sizemore	Sidney_Agnew3@nowhere.com	1_HUR_P-VOLs
6	2022-01-27 09:45:00.0000000	00005	Kallie	Flowers	xmhz1900@nowhere.com	1_HUR_P-VOLs
7	2022-01-27 09:45:00.0000000	00006	Breanne	Conway	Monroe.Rinehart@nowhere.com	1_HUR_P-VOLs
8	2022-01-27 09:45:00.0000000	00007	Letisha	Waite	Gutierrez73@example.com	1_HUR_P-VOLs
9	2022-01-27 09:45:00.0000000	00008	Taren	Skaggs	FelipaAbernathy@nowhere.com	1_HUR_P-VOLs
10	2022-01-27 09:45:00.0000000	00009	Zulma	Peacock	Adler@example.com	1_HUR_P-VOLs
11	2022-01-27 09:45:00.0000000	00010	Kalyn	Hwang	uadm55@example.com	1_HUR_P-VOLs

Query executed successfully. | D2_SQL_EC2_2\EC2

Results | Messages

	timestamp	Product_ID	Product_Name	Product_Price	Update_On
1	2022-01-28 06:26:40.0000000	00000	061A27D2Z8K3RE48VT07TROS1	4931	2_HUR-P-VOLs
2	2022-01-28 06:26:40.0000000	00001	S78RQ696155N6LK0G0T46FF0U	83382	2_HUR-P-VOLs
3	2022-01-28 06:26:40.0000000	00002	IQF24115O77U8S6Z2F06K39UA	93423	2_HUR-P-VOLs
4	2022-01-28 06:26:40.0000000	00003	C08CC2G28922THYZKMG3EKP1P	3163	2_HUR-P-VOLs
5	2022-01-28 06:26:40.0000000	00004	3BCU5G22NQ6GQBPRL6AY8734D	4240	2_HUR-P-VOLs
6	2022-01-28 06:26:40.0000000	00005	U75F2XG65028G55I3TIS2C634	4955	2_HUR-P-VOLs
7	2022-01-28 06:26:40.0000000	00006	N6AU7M970CWQX3KK02NQIW23N	4343	2_HUR-P-VOLs
8	2022-01-28 06:26:40.0000000	00007	AD64C0M934T7UA64NC053S0Q5	3426	2_HUR-P-VOLs
9	2022-01-28 06:26:40.0000000	00008	N9XWP73W4C8PJC5RZ4801LKA9	4578	2_HUR-P-VOLs
10	2022-01-28 06:26:40.0000000	00009	4PWA659A4EG0RIFOTHRGI28E3	4709	2_HUR-P-VOLs
11	2022-01-28 06:26:40.0000000	00010	7ZEHFV9S9YI97Y02XRJ3418U2	3961	2_HUR-P-VOLs

Query executed successfully. | D2_SQL_EC2_2\

8. When you are ready to delete the HTI snapshots, follow the instructions in section Test 3.

Test 5: Planned Outage

This test case demonstrates how HUR 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 during the outage. After the outage, data created on the secondary storage system is replicated back to the primary storage system.

1. Generate new records on the HUR primary volumes by creating a database named D2_Test1.5_DB1 and a table named Staff_Table. The Update_On column indicates these records were created on the HUR P-VOLs.

Results | Messages

	timestamp	Staff_ID	First_Name	Last_Name	Email	Update_On
1	2022-01-31 07:36:49.0000000	80460	Letha	Wahl	Bolt@nowhere.com	1_HUR_P-VOLs
2	2022-01-31 07:36:49.0000000	19222	Kraig	Boucher	Morehead@example.com	1_HUR_P-VOLs
3	2022-01-31 07:36:49.0000000	86849	Harlan	Ludwig	qpzre411@example.com	1_HUR_P-VOLs
4	2022-01-31 07:36:49.0000000	08276	Brad	Craven	Bray162@example.com	1_HUR_P-VOLs
5	2022-01-31 07:36:49.0000000	23003	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	1_HUR_P-VOLs
6	2022-01-31 07:36:49.0000000	68767	Bennie	Flowers	xmhz1900@nowhere.com	1_HUR_P-VOLs
7	2022-01-31 07:36:49.0000000	47864	Francis	Conway	Monroe.Rinehart@nowhere.com	1_HUR_P-VOLs
8	2022-01-31 07:36:49.0000000	44963	Nickie	Waite	Gutierrez73@example.com	1_HUR_P-VOLs
9	2022-01-31 07:36:49.0000000	15878	Boyd	Skaggs	FelipaAbernathy@nowhere.com	1_HUR_P-VOLs
10	2022-01-31 07:36:49.0000000	40093	Alonzo	Peacock	Adler@example.com	1_HUR_P-VOLs
11	2022-01-31 07:36:49.0000000	75351	Reed	Hwang	uadm55@example.com	1_HUR_P-VOLs

Query executed successfully. D2SQLDB\D2SQL

2. Initiate the planned outage.
   a. In SSMS, detach the database instance.
   b. In Windows Failover Cluster Manager, stop the SQL role and shut down the cluster.
   c. Power off the virtual machines.
3. Use Protector to suspend HUR replication and make the target volumes accessible.
   a. Click **Storage**, and then click on the secondary storage system.
   b. Click **Replication and Clones**.
   c. Select the HUR relationship. Click on the **Swap** icon (two arrows pointing in opposite directions).

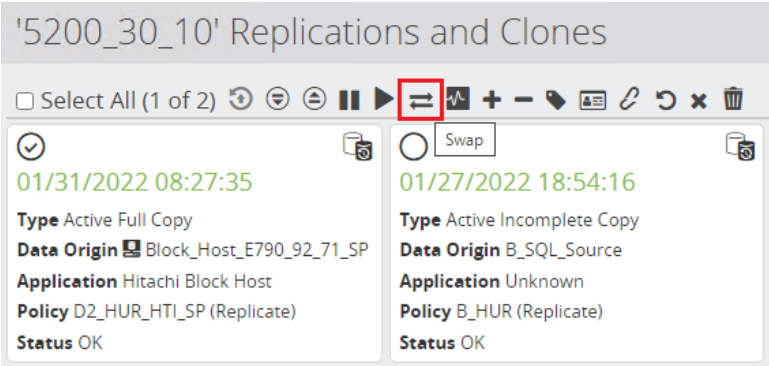

d. Type **SWAP** to confirm the task. In the Direction field, select the direction that is opposite of the current direction and click **Finish**.

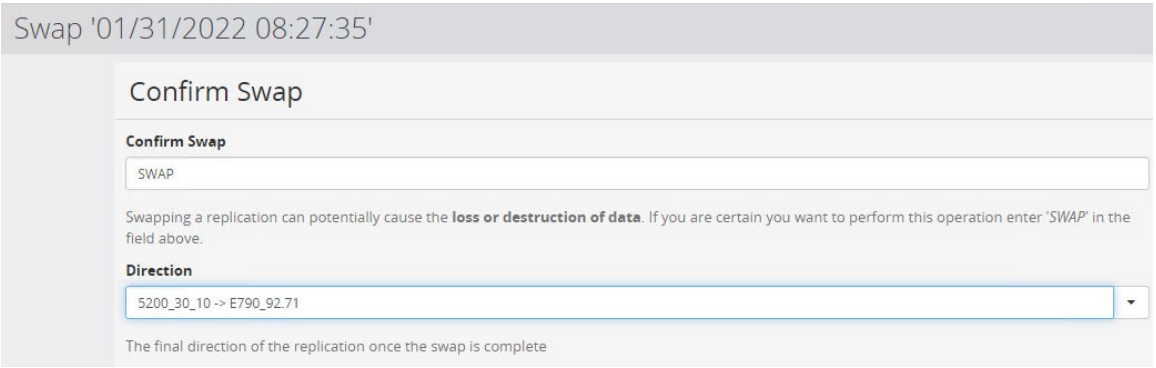

After the operation is complete, the **Status** field is **Swapped**.

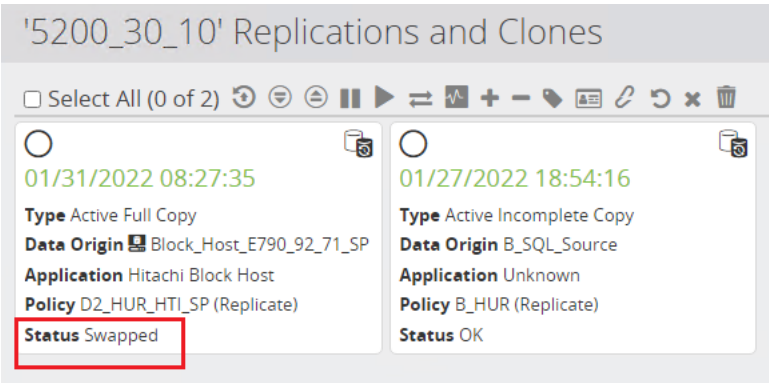

Also, the primary volumes are now S-VOLs, and the secondary volumes are now P-VOLs.

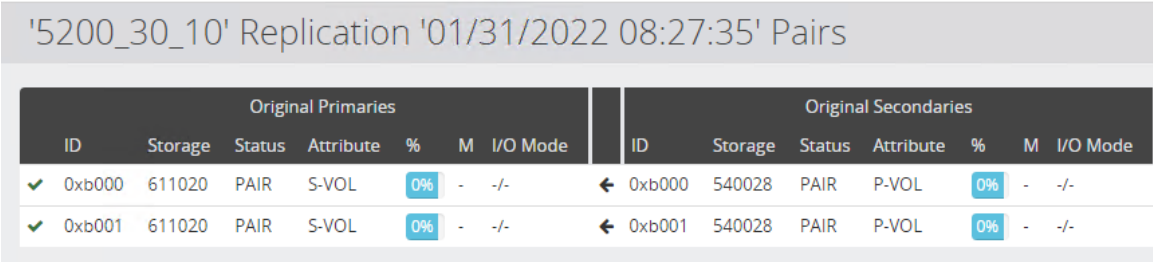

4. Now that the secondary volumes are accessible, bring up SQL on the EC2 virtual machine and write to the database.
   a. In Administrator, **Refresh** the secondary storage system.
   b. Map the volumes to the EC2 server as described in section Test 3.

c. In the EC2 virtual machine, bring the drives online and attach the database. Verify that the previously written records are present.

Results | Messages

	timestamp	Staff_ID	First_Name	Last_Name	Email	Update_On
1	2022-01-31 07:36:49.0000000	80460	Letha	Wahl	Bolt@nowhere.com	1_HUR_P-VOLs
2	2022-01-31 07:36:49.0000000	19222	Kraig	Boucher	Morehead@example.com	1_HUR_P-VOLs
3	2022-01-31 07:36:49.0000000	86849	Harlan	Ludwig	qpzre411@example.com	1_HUR_P-VOLs
4	2022-01-31 07:36:49.0000000	08276	Brad	Craven	Bray162@example.com	1_HUR_P-VOLs
5	2022-01-31 07:36:49.0000000	23003	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	1_HUR_P-VOLs
6	2022-01-31 07:36:49.0000000	68767	Bennie	Flowers	xmhz1900@nowhere.com	1_HUR_P-VOLs
7	2022-01-31 07:36:49.0000000	47864	Francis	Conway	Monroe.Rinehart@nowhere.com	1_HUR_P-VOLs
8	2022-01-31 07:36:49.0000000	44963	Nickie	Waite	Gutierrez73@example.com	1_HUR_P-VOLs
9	2022-01-31 07:36:49.0000000	15878	Boyd	Skaggs	FelipaAbernathy@nowhere.com	1_HUR_P-VOLs
10	2022-01-31 07:36:49.0000000	40093	Alonzo	Peacock	Adler@example.com	1_HUR_P-VOLs
11	2022-01-31 07:36:49.0000000	75351	Reed	Hwang	uadm55@example.com	1_HUR_P-VOLs

Query executed successfully. D2_SQL_EC2_1\EC2S

5. Write new records to SQL while it is attached to the EC2 virtual machine to simulate updates to the database.
   a. Create a new database called D2_Test1.5_DB2.

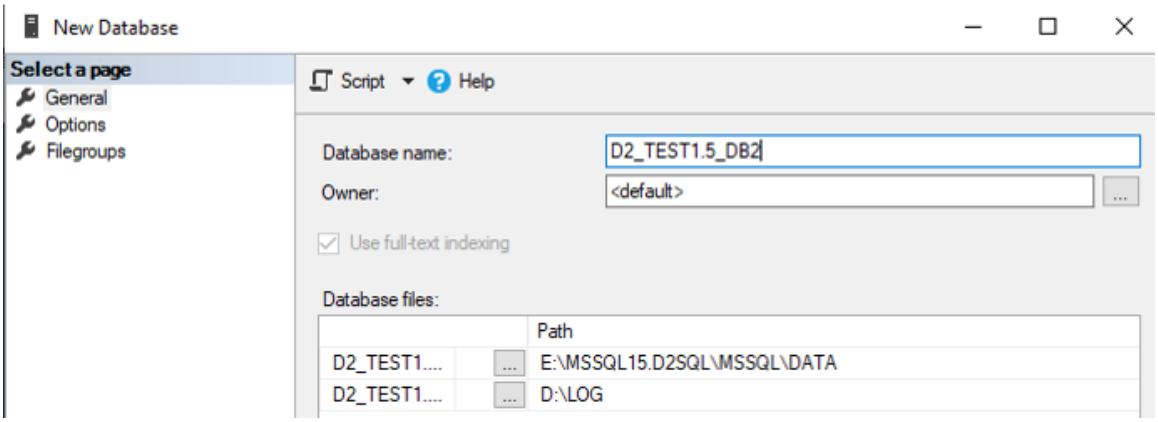

b. Create a new table and ingest new records in it. Notice the information in the **Update_On** column shows where the new records were created.

Results | Messages

	timestamp	Store_ID	Store_Name	City	Update_On
1	2022-01-31 11:34:42.0000000	00000	Gift Store	Boca Raton	HUR_S-VOLs
2	2022-01-31 11:34:42.0000000	00001	Supermarket	University Place	HUR_S-VOLs
3	2022-01-31 11:34:42.0000000	00002	Gift Store	North Royalton	HUR_S-VOLs
4	2022-01-31 11:34:42.0000000	00003	Agent	Atlantic	HUR_S-VOLs
5	2022-01-31 11:34:42.0000000	00004	Wholesaler	Marquette	HUR_S-VOLs
6	2022-01-31 11:34:42.0000000	00005	Corporate	Rochester	HUR_S-VOLs
7	2022-01-31 11:34:42.0000000	00006	Gift Store	Lindsay	HUR_S-VOLs
8	2022-01-31 11:34:42.0000000	00007	Gift Store	Montgomery	HUR_S-VOLs
9	2022-01-31 11:34:42.0000000	00008	Computer Store	Allen Park	HUR_S-VOLs
10	2022-01-31 11:34:42.0000000	00009	Gift Store	Monroe	HUR_S-VOLs
11	2022-01-31 11:34:42.0000000	00010	Agent	Surfside	HUR_S-VOLs

Query executed successfully.

6. In the current swapped state, we can still create HTI snapshots on the secondary storage system. The steps are as follows:
   a. In Protector, click **Monitor** and click into the data flow.
   b. Select the secondary storage system and click **Trigger**.

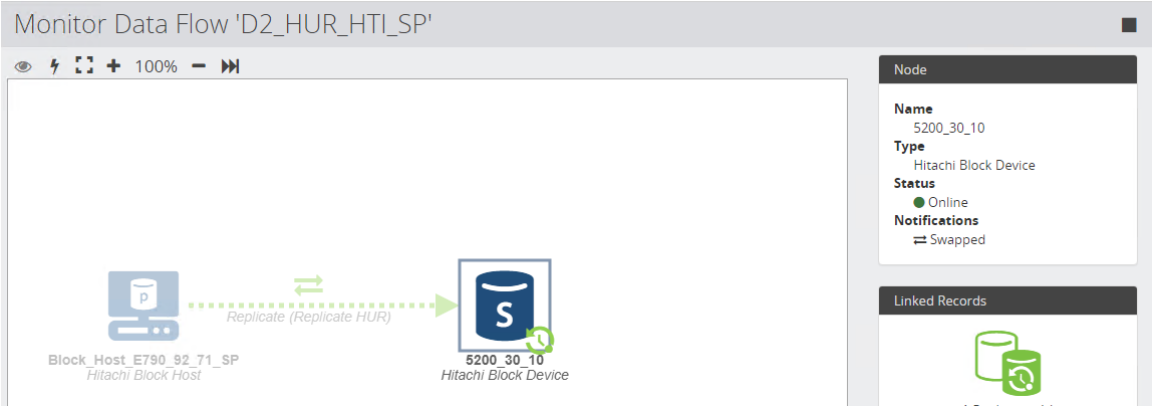

   c. After the snapshots are created, get the HTI snapshot details.

'5200_30_10' Snapshot '01/31/2022 13:42:26' Pairs

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	540028	-	P-VOL	-		-/-	→	0x0019	540028	-	S-VOL	-		-/-
0xb001	540028	-	P-VOL	-		-/-	→	0x001a	540028	-	S-VOL	-		-/-

   d. In Administrator, **Refresh** the secondary storage system.
   e. Map the snapshot volumes, as described in section Test 3, to a second EC2 virtual machine.
   f. In the EC2 virtual machine, bring the drives online and attach the databases: D2_Test1.5_DB1 (created while running on HUR primary volumes) and D2_Test1.5_DB2 (created while running on HUR secondary volumes).
   g. Create a new database and write some updates. This ensures we can read and write to the snapshot volumes during a production outage. The following screenshots show a new database

and a new table that were created on the HTI snapshot volumes: Employee_Table in D2_Test1.5_DB3.

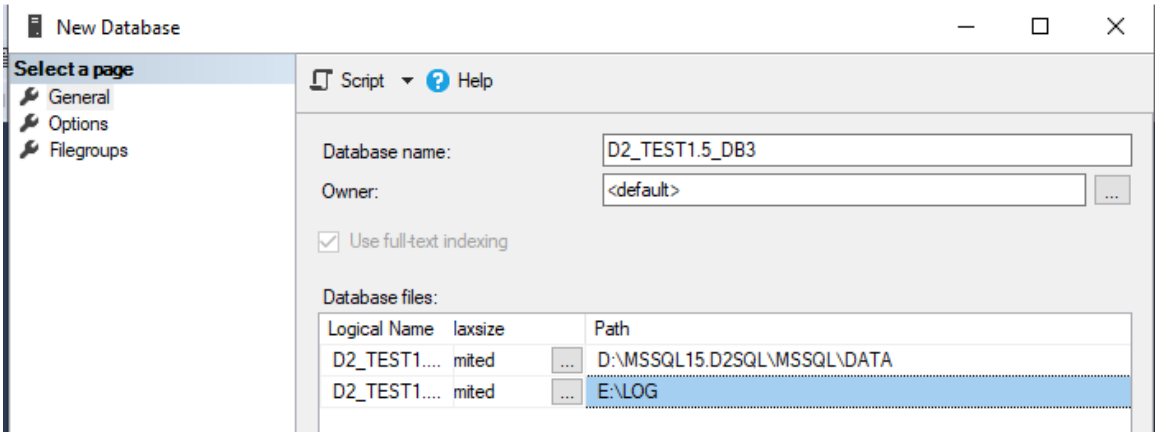

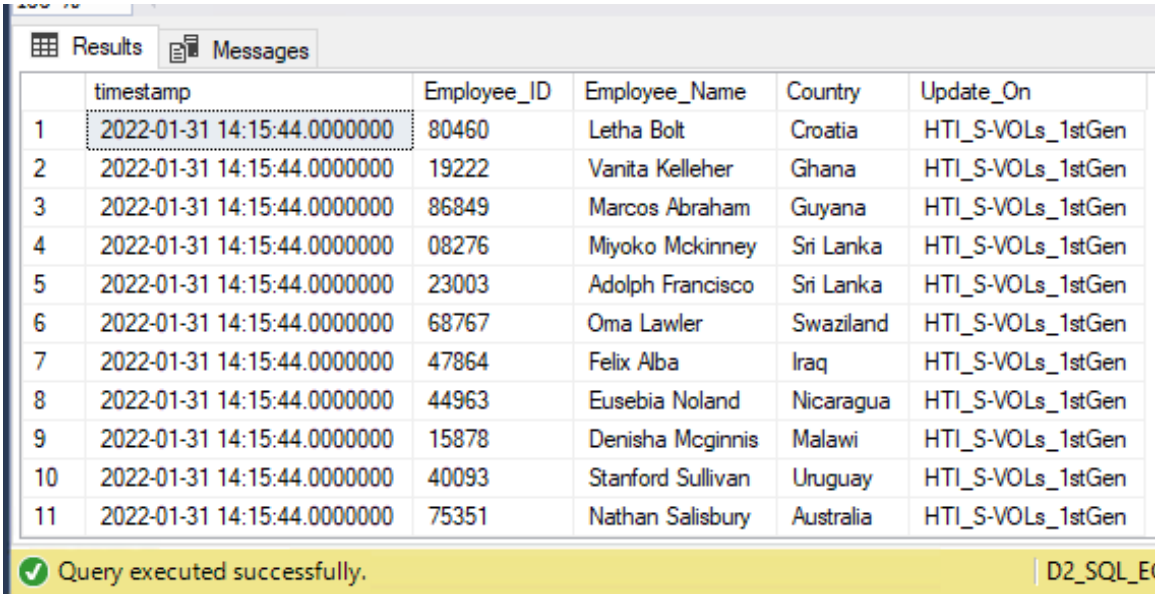

	timestamp	Employee_ID	Employee_Name	Country	Update_On
1	2022-01-31 14:15:44.0000000	80460	Letha Bolt	Croatia	HTI_S-VOLs_1stGen
2	2022-01-31 14:15:44.0000000	19222	Vanita Kelleher	Ghana	HTI_S-VOLs_1stGen
3	2022-01-31 14:15:44.0000000	86849	Marcos Abraham	Guyana	HTI_S-VOLs_1stGen
4	2022-01-31 14:15:44.0000000	08276	Miyoko Mckinney	Sri Lanka	HTI_S-VOLs_1stGen
5	2022-01-31 14:15:44.0000000	23003	Adolph Francisco	Sri Lanka	HTI_S-VOLs_1stGen
6	2022-01-31 14:15:44.0000000	68767	Oma Lawler	Swaziland	HTI_S-VOLs_1stGen
7	2022-01-31 14:15:44.0000000	47864	Felix Alba	Iraq	HTI_S-VOLs_1stGen
8	2022-01-31 14:15:44.0000000	44963	Eusebia Noland	Nicaragua	HTI_S-VOLs_1stGen
9	2022-01-31 14:15:44.0000000	15878	Denisha Mcginnis	Malawi	HTI_S-VOLs_1stGen
10	2022-01-31 14:15:44.0000000	40093	Stanford Sullivan	Uruguay	HTI_S-VOLs_1stGen
11	2022-01-31 14:15:44.0000000	75351	Nathan Salisbury	Australia	HTI_S-VOLs_1stGen

h. When you are ready to delete the HTI snapshots, follow the instructions in section Test 3.

7. When you are ready to restore SQL to the primary site:
   a. In the original EC2 virtual machine SSMS, detach the databases.
   b. In Protector, click **Storage** and then click the secondary storage system.
   c. Click **Replication and Clones**. Note that the current status is **Swapped**.
   d. Select the relationship and click **Swap**.

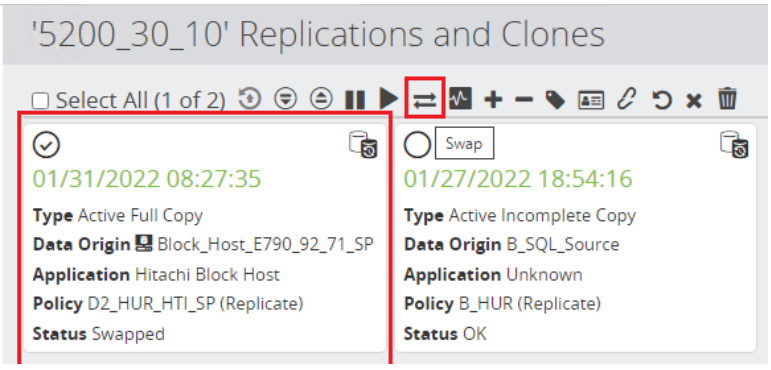

e. Type **SWAP** to confirm the task. In the Direction field, select the original direction and then click **Finish**.

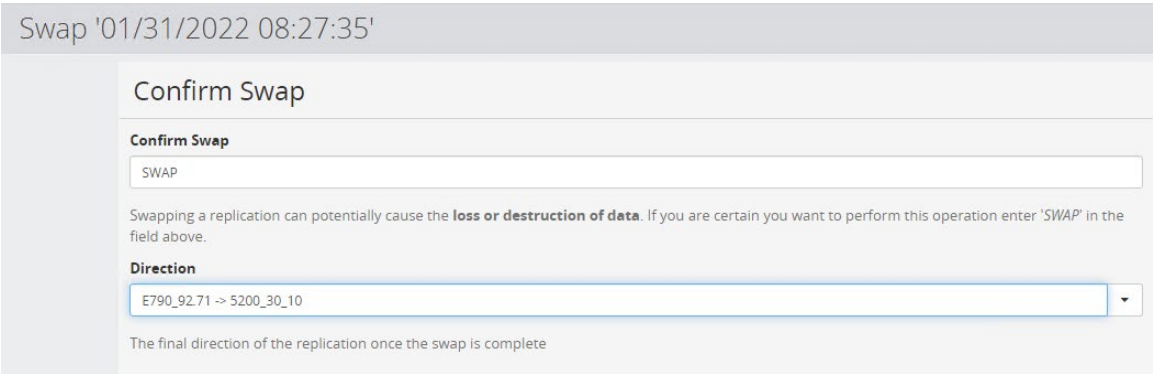

f. After the swap is done, the original primary volumes are now P-VOLs again, and original secondary volumes are S-VOLs again.

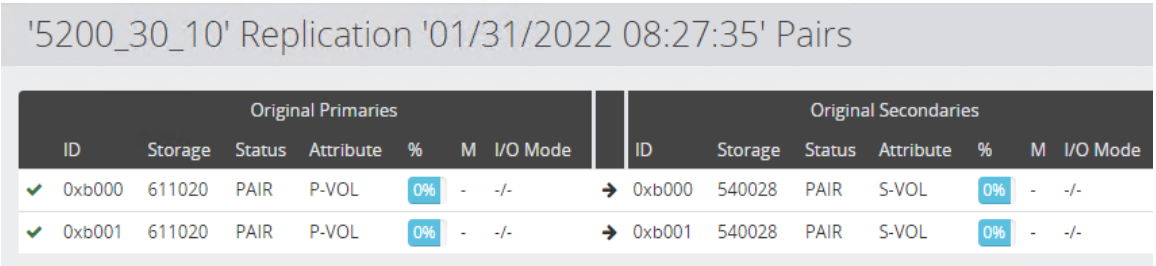

The Status field shows **OK** instead of Swapped now.

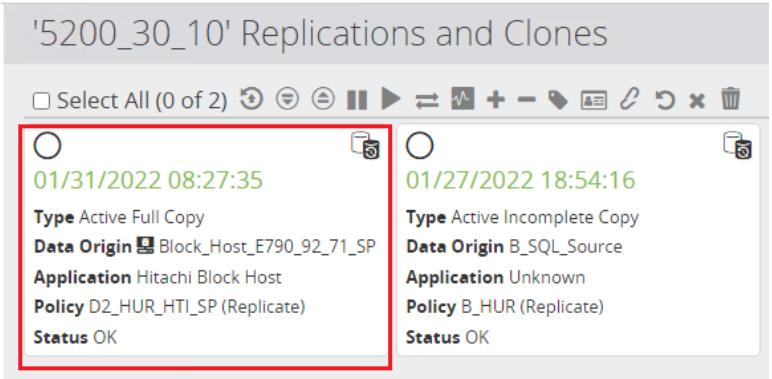

8. Power on the SQL virtual machines at the primary site.

a. After several minutes, check Windows Failover Cluster Manager to ensure the nodes are online.

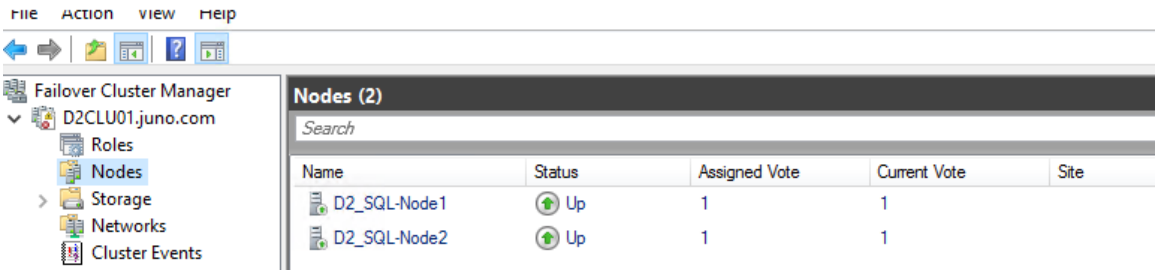

b. Go to **Roles** and start SQL.

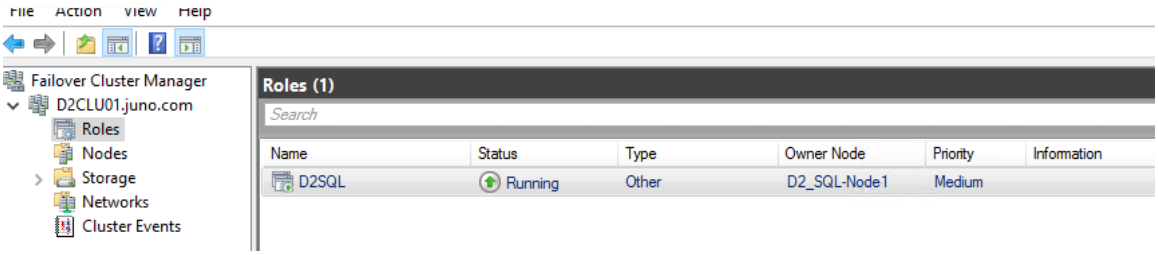

c. Go to **Storage > Disks** and bring the disks online.

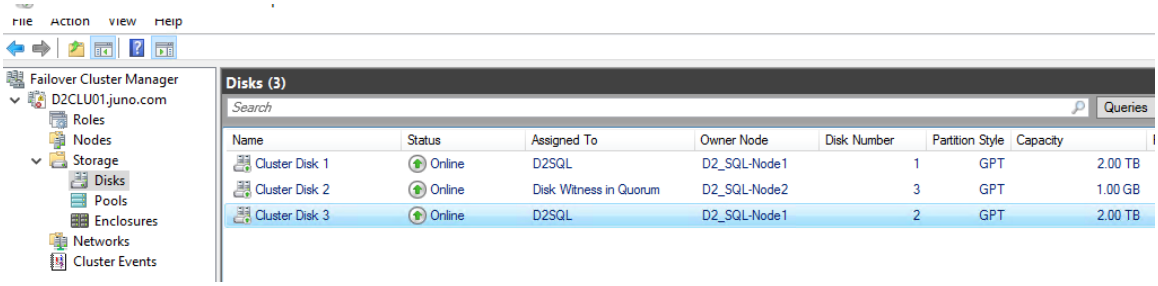

9. Verify that the new records written to SQL at the DR site are present.
   a. In SSMS, import the databases by attaching the database files. There are two databases: D2_Test1.5_DB1 (created while running on HUR primary volumes) and D2_Test1.5_DB2 (created while running on HUR secondary volumes).

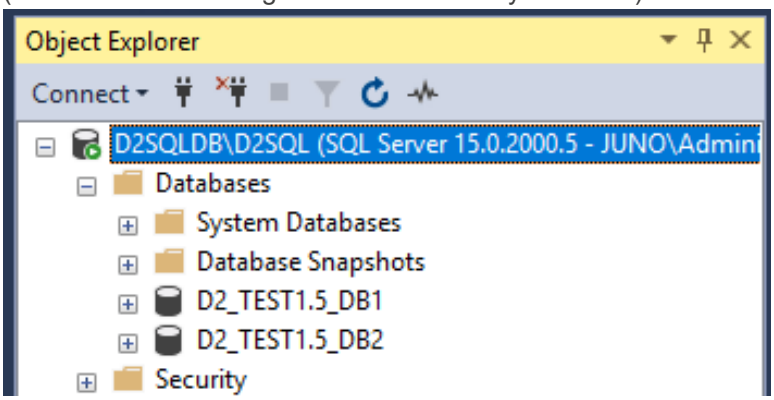

   b. Check the content of the databases.

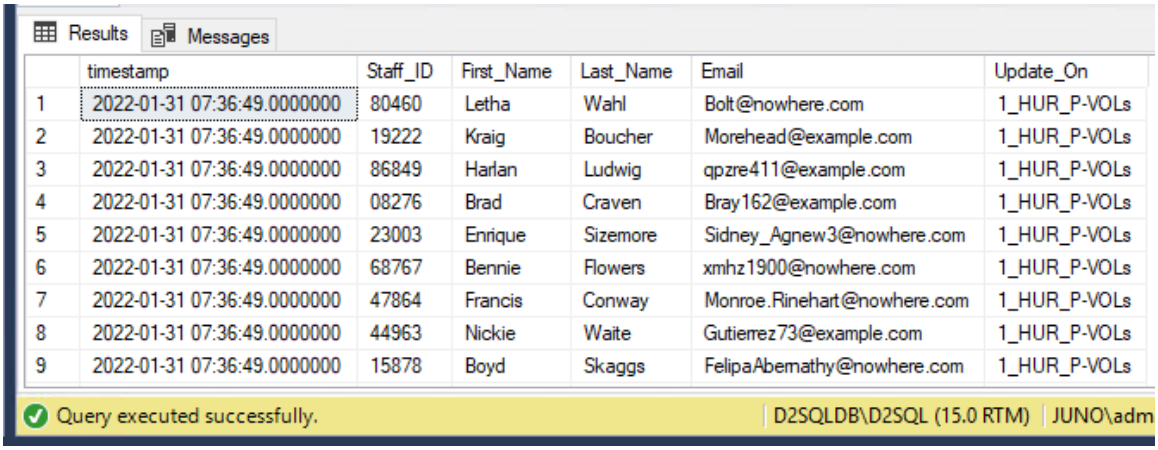

	timestamp	Staff_ID	First_Name	Last_Name	Email	Update_On
1	2022-01-31 07:36:49.0000000	80460	Letha	Wahl	Bolt@nowhere.com	1_HUR_P-VOLs
2	2022-01-31 07:36:49.0000000	19222	Kraig	Boucher	Morehead@example.com	1_HUR_P-VOLs
3	2022-01-31 07:36:49.0000000	86849	Harlan	Ludwig	qpzre411@example.com	1_HUR_P-VOLs
4	2022-01-31 07:36:49.0000000	08276	Brad	Craven	Bray162@example.com	1_HUR_P-VOLs
5	2022-01-31 07:36:49.0000000	23003	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	1_HUR_P-VOLs
6	2022-01-31 07:36:49.0000000	68767	Bennie	Flowers	xmhz1900@nowhere.com	1_HUR_P-VOLs
7	2022-01-31 07:36:49.0000000	47864	Francis	Conway	Monroe.Rinehart@nowhere.com	1_HUR_P-VOLs
8	2022-01-31 07:36:49.0000000	44963	Nickie	Waite	Gutierrez73@example.com	1_HUR_P-VOLs
9	2022-01-31 07:36:49.0000000	15878	Boyd	Skaggs	FelipaAbernathy@nowhere.com	1_HUR_P-VOLs

Results | Messages

	timestamp	Store_ID	Store_Name	City	Update_On
1	2022-01-31 11:34:42.0000000	00000	Gift Store	Boca Raton	HUR_S-VOLs
2	2022-01-31 11:34:42.0000000	00001	Supermarket	University Place	HUR_S-VOLs
3	2022-01-31 11:34:42.0000000	00002	Gift Store	North Royalton	HUR_S-VOLs
4	2022-01-31 11:34:42.0000000	00003	Agent	Atlantic	HUR_S-VOLs
5	2022-01-31 11:34:42.0000000	00004	Wholesaler	Marquette	HUR_S-VOLs
6	2022-01-31 11:34:42.0000000	00005	Corporate	Rochester	HUR_S-VOLs
7	2022-01-31 11:34:42.0000000	00006	Gift Store	Lindsay	HUR_S-VOLs
8	2022-01-31 11:34:42.0000000	00007	Gift Store	Montgomery	HUR_S-VOLs
9	2022-01-31 11:34:42.0000000	00008	Computer Store	Allen Park	HUR_S-VOLs

Query executed successfully. | D2SQLDB\D2SQ

Test 6: Unplanned Outage

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. After the outage, data created on the secondary storage system is replicated back to the primary storage system.

> **Note**: Because Hitachi Universal Replicator is asynchronous, the secondary storage system might not contain all updates from the primary storage system after an unplanned outage. However, HUR journals maintain the sequence of updates so the data that is at the secondary storage system will be consistent.

1. Generate new records on the HUR primary volumes by creating a database named D2_Test1.6_DB1 and a table named Staff_Table. The Update_On column indicates these records were created on the HUR P-VOLs.

Results | Messages

	timestamp	Staff_ID	First_name	Last_Name	Email	Update_On
1	2022-02-06 10:14:33.0000000	80460	Letha	Wahl	Bolt@nowhere.com	HUR_P-VOLs
2	2022-02-06 10:14:33.0000000	19222	Kraig	Boucher	Morehead@example.com	HUR_P-VOLs
3	2022-02-06 10:14:33.0000000	86849	Harlan	Ludwig	qpzre411@example.com	HUR_P-VOLs
4	2022-02-06 10:14:33.0000000	08276	Brad	Craven	Bray162@example.com	HUR_P-VOLs
5	2022-02-06 10:14:33.0000000	23003	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	HUR_P-VOLs
6	2022-02-06 10:14:33.0000000	68767	Bennie	Flowers	xmhz1900@nowhere.com	HUR_P-VOLs
7	2022-02-06 10:14:33.0000000	47864	Francis	Conway	Monroe.Rinehart@nowhere.com	HUR_P-VOLs
8	2022-02-06 10:14:33.0000000	44963	Nickie	Waite	Gutierrez73@example.com	HUR_P-VOLs
9	2022-02-06 10:14:33.0000000	15878	Boyd	Skaggs	FelipaAbernathy@nowhere.com	HUR_P-VOLs

Query executed successfully. | D2SQLDB\D2SQL (15.0 RTM) | JUNO\a

2. Initiate the outage by disabling the FC ports between the VSP E790 and the Production cluster and iSCSI ports used for HUR remote paths.

> Note: The primary ISM virtual machine was also powered off to accurately simulate a site failure. This node must be offline for Protector to perform a Suspend for Swap operation.

   a. Verify that the HUR remote paths are offline.

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-G	CL1-C	iSCSI	Communication Failed	-	-	Receiver
1	CL2-H	CL2-C	iSCSI	Initialization Failed	-	-	Receiver
2	CL3-G	CL3-C	iSCSI	Communication Failed	-	-	Receiver
3	CL4-H	CL4-C	iSCSI	Communication Failed	-	-	Receiver

   b. Verify that the VSP E790 status is showing disconnected because the primary ISM node is unavailable.

Storage

5200_30_10
Type Hitachi Block Device
Proxy sv10probehost
Serial Number 540028

E790_92.71
Type Hitachi Block Device
Proxy probehost
Serial Number 611020

c. In vCenter Server, verify that volumes from the VSP E790 storage system are in Dead or Error state.

Name	LUN	Type	Capacity	Datastore	Operational State
HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000b001)	5	disk	2.00 TB	Not Consumed	Dead or Error
HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000be0...	3	disk	10.00 GB	Not Consumed	Dead or Error
HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000b00...	10	disk	2.00 TB	D_VMFS	Dead or Error
HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000b00...	4	disk	2.00 TB	Not Consumed	Dead or Error

3. Use Protector to promote the HUR secondary volumes and make them accessible:
   a. Click **Storage**, and then click on the replication target storage system.
   b. Click **Replication and Clones**.
   c. Click the HUR relationship. The following screen should display:

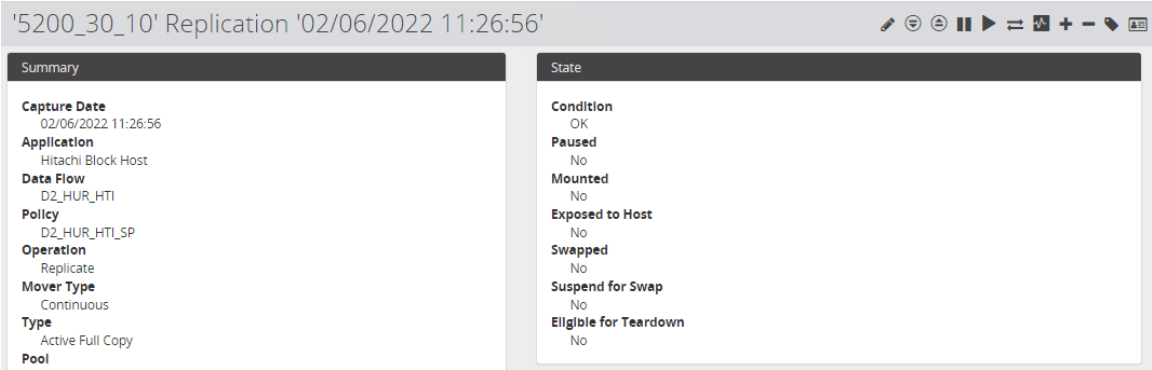

   d. Click on the **Swap** icon (two arrows pointing in opposite directions).
   e. Type **SWAP** to confirm the task. In the Direction field, select the direction that is opposite of the current direction and click **Finish**.

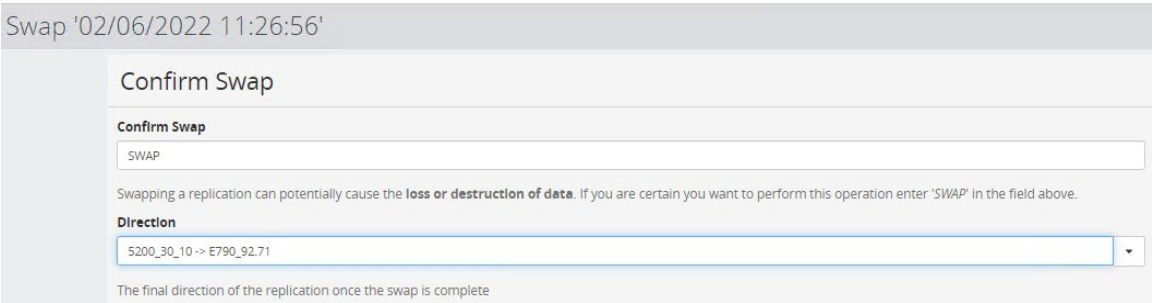

   f. After the swap is done, the Suspend for Swap field is changed to Yes.

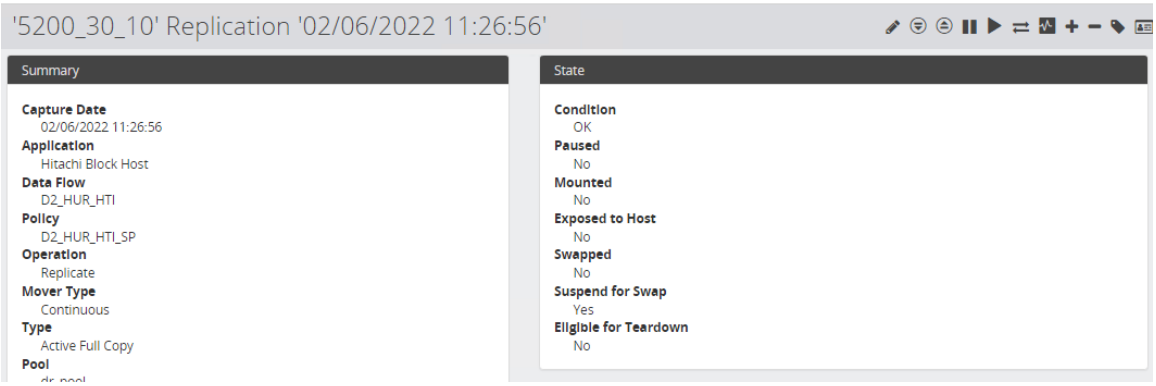

Because the primary ISM is unavailable, the status field still shows OK and the pair roles are unchanged.

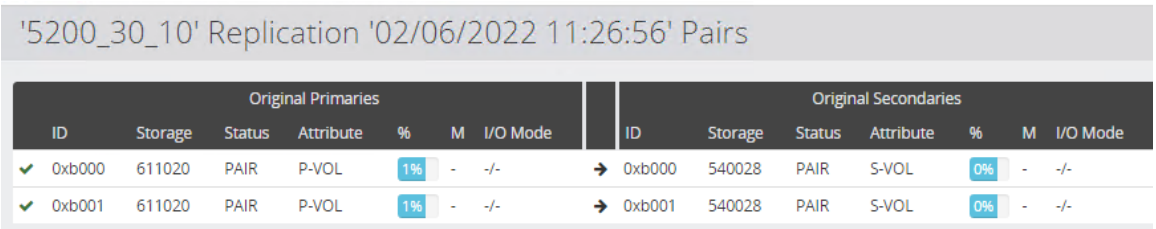

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	611020	PAIR	P-VOL	1%	-	-/-	→	0xb000	540028	PAIR	S-VOL	0%	-	-/-
0xb001	611020	PAIR	P-VOL	1%	-	-/-	→	0xb001	540028	PAIR	S-VOL	0%	-	-/-

4. Now that the secondary volumes are accessible, bring up SQL on the EC2 virtual machine and write to the database.
   a. In Administrator, **Refresh** the secondary storage system.
   b. Map the volumes to the EC2 server as described in section Test 3.
   c. In the EC2 virtual machine, bring the drives online and attach the database. Verify that the previously written records are present.

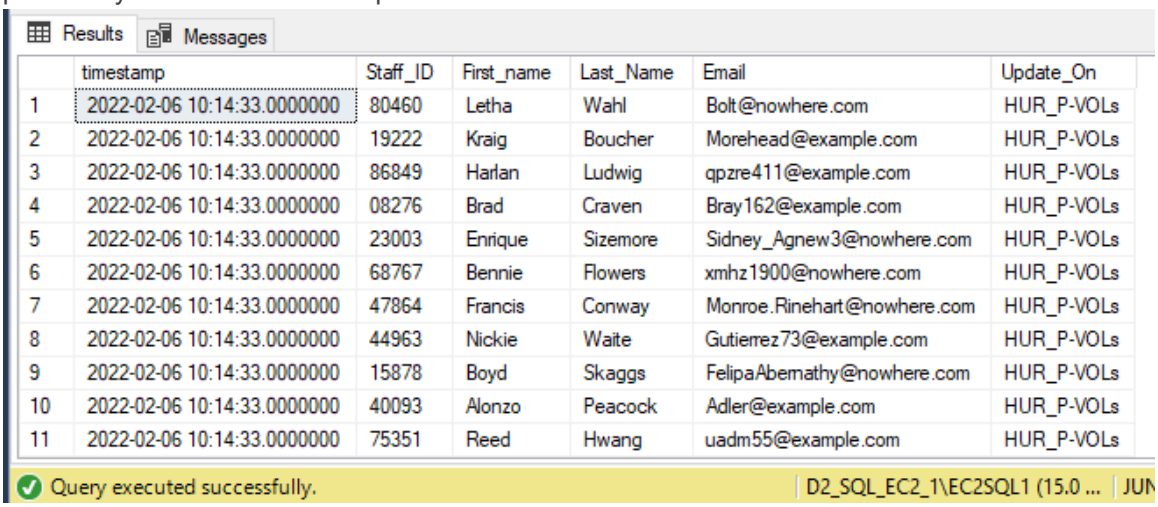

Results | Messages

	timestamp	Staff_ID	First_name	Last_Name	Email	Update_On
1	2022-02-06 10:14:33.0000000	80460	Letha	Wahl	Bolt@nowhere.com	HUR_P-VOLs
2	2022-02-06 10:14:33.0000000	19222	Kraig	Boucher	Morehead@example.com	HUR_P-VOLs
3	2022-02-06 10:14:33.0000000	86849	Harlan	Ludwig	qpzre411@example.com	HUR_P-VOLs
4	2022-02-06 10:14:33.0000000	08276	Brad	Craven	Bray162@example.com	HUR_P-VOLs
5	2022-02-06 10:14:33.0000000	23003	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	HUR_P-VOLs
6	2022-02-06 10:14:33.0000000	68767	Bennie	Flowers	xmhz1900@nowhere.com	HUR_P-VOLs
7	2022-02-06 10:14:33.0000000	47864	Francis	Conway	Monroe.Rinehart@nowhere.com	HUR_P-VOLs
8	2022-02-06 10:14:33.0000000	44963	Nickie	Waite	Gutierrez73@example.com	HUR_P-VOLs
9	2022-02-06 10:14:33.0000000	15878	Boyd	Skaggs	FelipaAbernathy@nowhere.com	HUR_P-VOLs
10	2022-02-06 10:14:33.0000000	40093	Alonzo	Peacock	Adler@example.com	HUR_P-VOLs
11	2022-02-06 10:14:33.0000000	75351	Reed	Hwang	uadm55@example.com	HUR_P-VOLs

Query executed successfully. | D2_SQL_EC2_1\EC2SQL1 (15.0 ... | JUN

5. Write new records to SQL while it is attached to the EC2 virtual machine to simulate updates to the database.

a. Create a new database called D2_Test1.6_DB2.

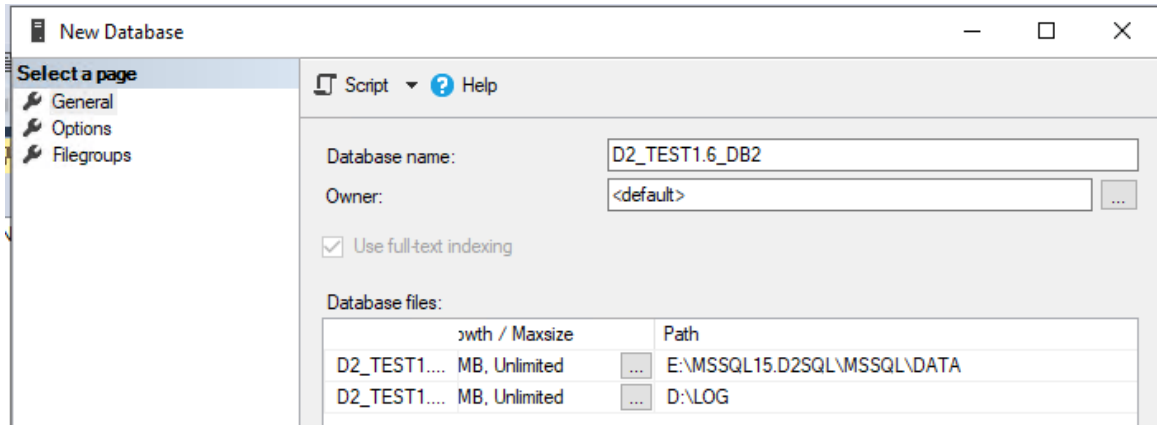

b. Create a new table and ingest new records in it. Notice the information in the **Update_On** column shows where the new records were created.

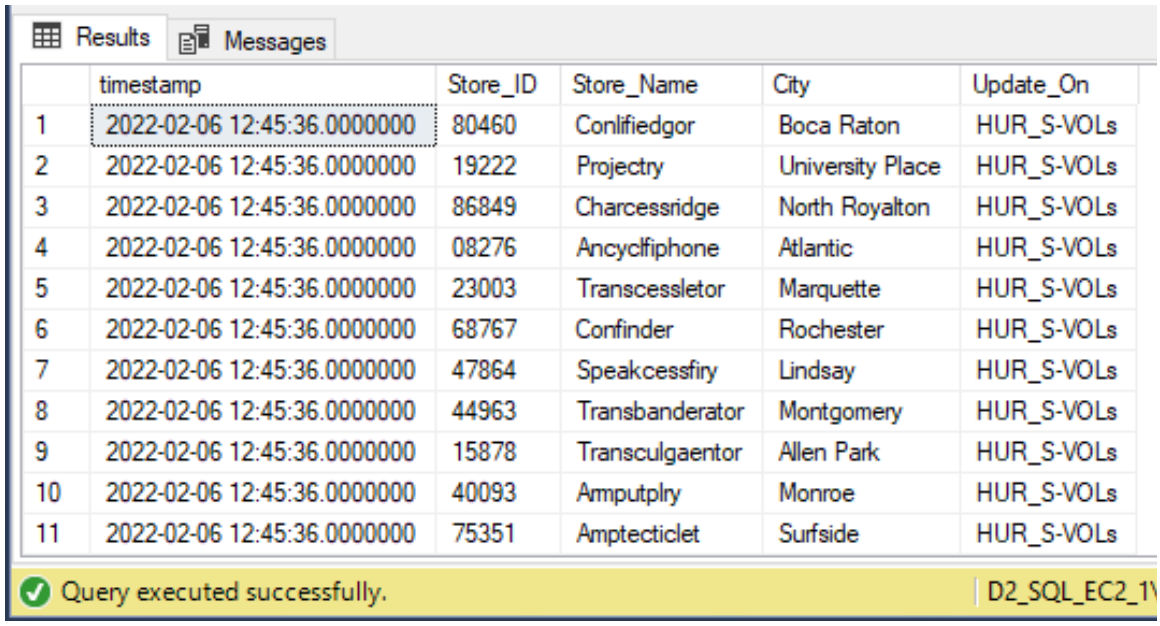

	timestamp	Store_ID	Store_Name	City	Update_On
1	2022-02-06 12:45:36.0000000	80460	Conlifiedgor	Boca Raton	HUR_S-VOLs
2	2022-02-06 12:45:36.0000000	19222	Projectry	University Place	HUR_S-VOLs
3	2022-02-06 12:45:36.0000000	86849	Charcessridge	North Royalton	HUR_S-VOLs
4	2022-02-06 12:45:36.0000000	08276	Ancyclfiphone	Atlantic	HUR_S-VOLs
5	2022-02-06 12:45:36.0000000	23003	Transcessletor	Marquette	HUR_S-VOLs
6	2022-02-06 12:45:36.0000000	68767	Confinder	Rochester	HUR_S-VOLs
7	2022-02-06 12:45:36.0000000	47864	Speakcessfiry	Lindsay	HUR_S-VOLs
8	2022-02-06 12:45:36.0000000	44963	Transbanderator	Montgomery	HUR_S-VOLs
9	2022-02-06 12:45:36.0000000	15878	Transculgaentor	Allen Park	HUR_S-VOLs
10	2022-02-06 12:45:36.0000000	40093	Amputplry	Monroe	HUR_S-VOLs
11	2022-02-06 12:45:36.0000000	75351	Amptecticlet	Surfside	HUR_S-VOLs

6. In the current swapped state, we can still create HTI snapshots on the secondary storage system as follows:

> Note: Snapshots cannot be triggered in the original data flow while the primary ISM is unavailable. A new data flow that uses the secondary ISM is needed.

a. Create a Block Host specifying the HUR secondary volumes from the VSP 5200 storage system.
b. Create a Policy with the **Hitachi Block** classification and **Snapshot** operation.

c. Create a data flow with the new policy.

d. Go to **Monitor** and click into the data flow.

e. Select the block host and click **Trigger**. Click **Run Now**.

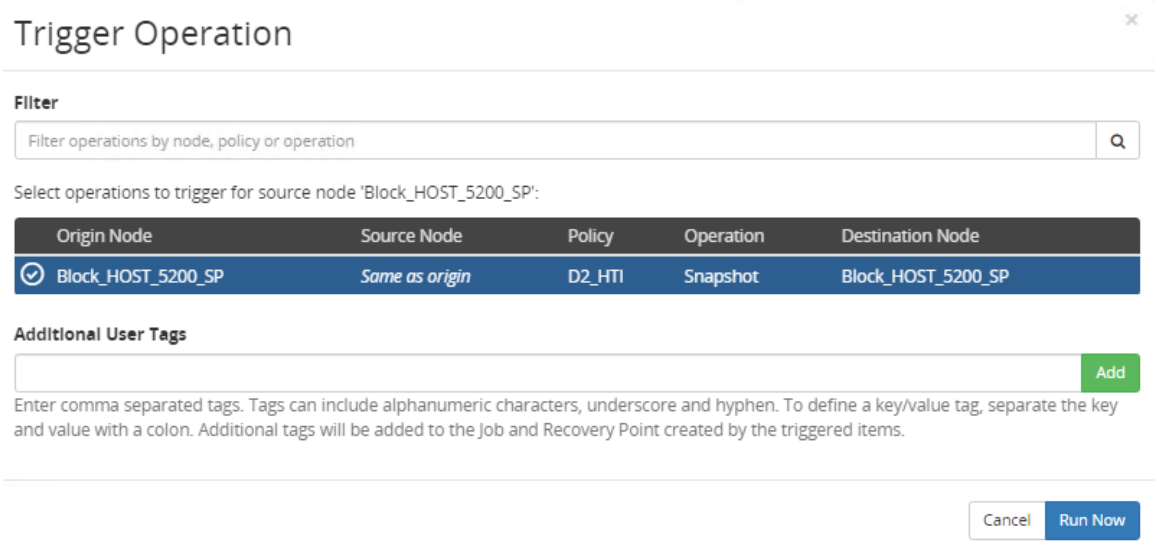

f. After the snapshots are created, get the HTI snapshot details.

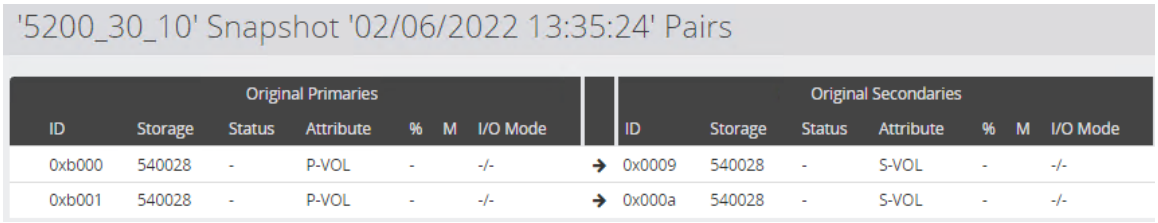

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	540028	-	P-VOL	-		-/-	→	0x0009	540028	-	S-VOL	-		-/-
0xb001	540028	-	P-VOL	-		-/-	→	0x000a	540028	-	S-VOL	-		-/-

g. In Administrator, **Refresh** the secondary storage system.

h. Map the snapshot volumes, as described in section Test 3, to a second EC2 virtual machine.

i. In the second EC2 virtual machine, bring the drives online and attach the databases: D2_Test1.6_DB1 (created while running on HUR primary volumes) and D2_Test1.6_DB2 (created while running on HUR secondary volumes).

j. Create a new database and write some updates. This ensures we can read and write to the snapshot volumes during a production outage.

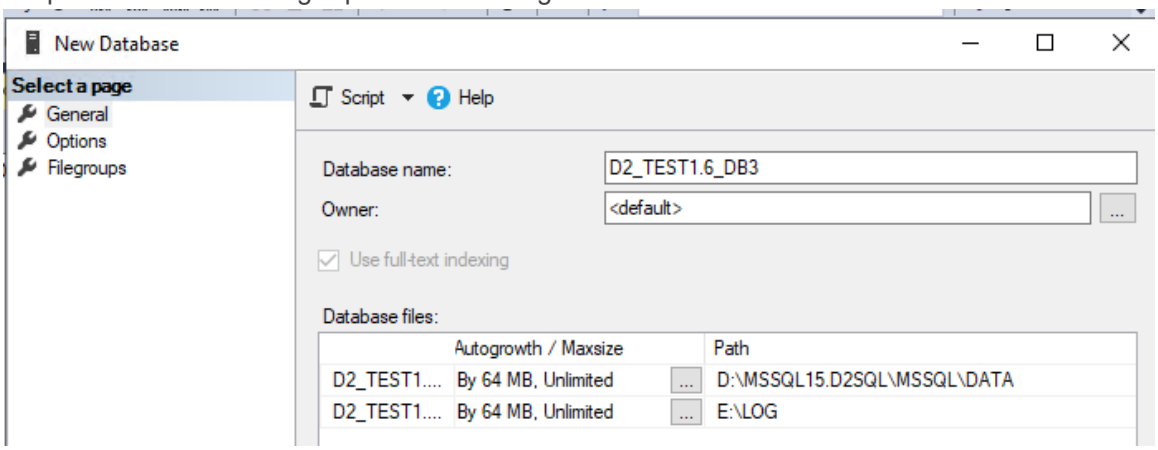

The following screenshot shows a new table and new database were created on the HTI snapshot volumes: Employee_Table in D2_Test1.6_DB3.

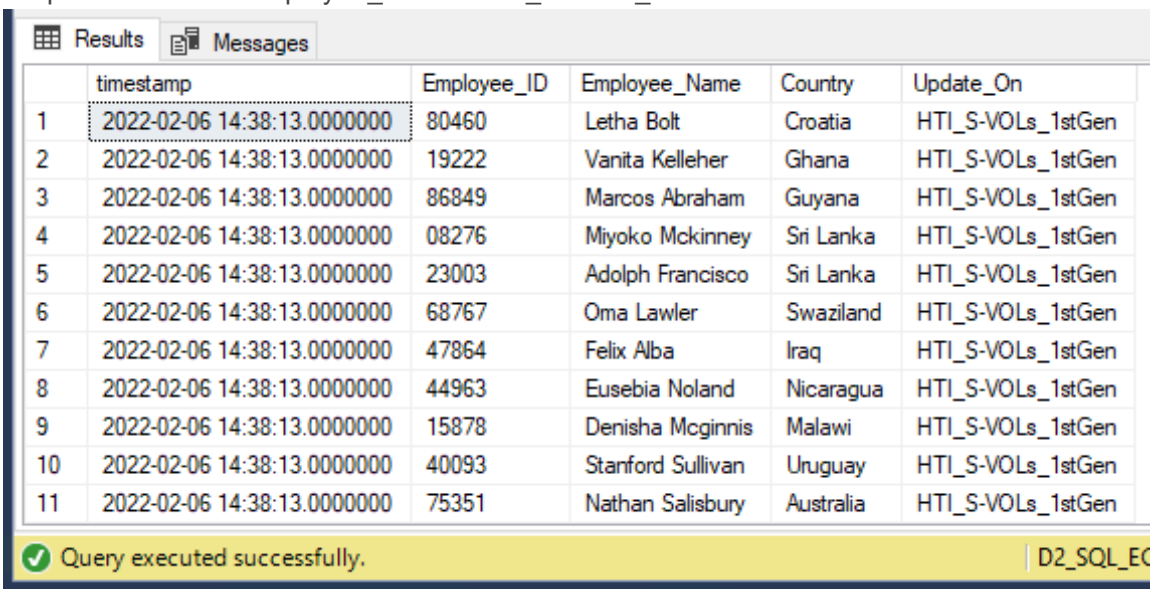

Results | Messages

	timestamp	Employee_ID	Employee_Name	Country	Update_On
1	2022-02-06 14:38:13.0000000	80460	Letha Bolt	Croatia	HTI_S-VOLs_1stGen
2	2022-02-06 14:38:13.0000000	19222	Vanita Kelleher	Ghana	HTI_S-VOLs_1stGen
3	2022-02-06 14:38:13.0000000	86849	Marcos Abraham	Guyana	HTI_S-VOLs_1stGen
4	2022-02-06 14:38:13.0000000	08276	Miyoko Mckinney	Sri Lanka	HTI_S-VOLs_1stGen
5	2022-02-06 14:38:13.0000000	23003	Adolph Francisco	Sri Lanka	HTI_S-VOLs_1stGen
6	2022-02-06 14:38:13.0000000	68767	Oma Lawler	Swaziland	HTI_S-VOLs_1stGen
7	2022-02-06 14:38:13.0000000	47864	Felix Alba	Iraq	HTI_S-VOLs_1stGen
8	2022-02-06 14:38:13.0000000	44963	Eusebia Noland	Nicaragua	HTI_S-VOLs_1stGen
9	2022-02-06 14:38:13.0000000	15878	Denisha Mcginnis	Malawi	HTI_S-VOLs_1stGen
10	2022-02-06 14:38:13.0000000	40093	Stanford Sullivan	Uruguay	HTI_S-VOLs_1stGen
11	2022-02-06 14:38:13.0000000	75351	Nathan Salisbury	Australia	HTI_S-VOLs_1stGen

Query executed successfully. D2_SQL_EC

k. When you are ready to delete the HTI snapshots, follow the instructions in section Test 3.

7. Bring the primary data center back online by re-enabling the FC ports between the VSP E790 storage system and the production cluster and the iSCSI ports used for HUR remote paths. Also, power on the primary ISM.
8. After the primary data center is back online, complete the swap process to begin replication (from secondary storage system to primary storage system).
   a. In Protector, refresh the cache of the two storage systems by going to **Storage**, clicking a storage system, and then clicking the **Refresh** icon on the top right. After refreshing, the HUR primary volumes should show as PSUE and HUR secondary volumes should show as SSWS.

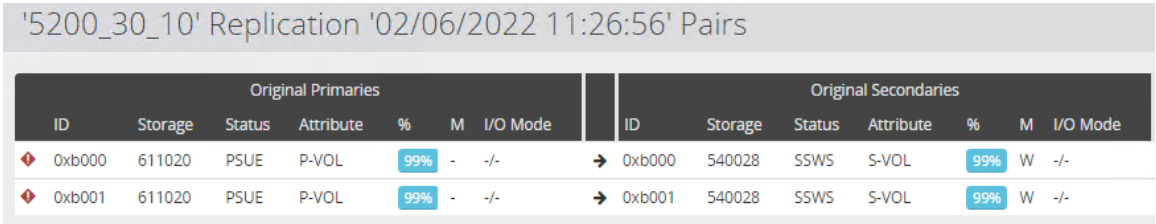

'5200_30_10' Replication '02/06/2022 11:26:56' Pairs

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	611020	PSUE	P-VOL	99%	-	-/-	→	0xb000	540028	SSWS	S-VOL	99%	W	-/-
0xb001	611020	PSUE	P-VOL	99%	-	-/-	→	0xb001	540028	SSWS	S-VOL	99%	W	-/-

   b. Reestablish replication by clicking **Swap** (this is the second time; the first time was to make the HUR secondary volumes accessible after the outage). Select the direction that goes from

secondary storage system to primary storage system.

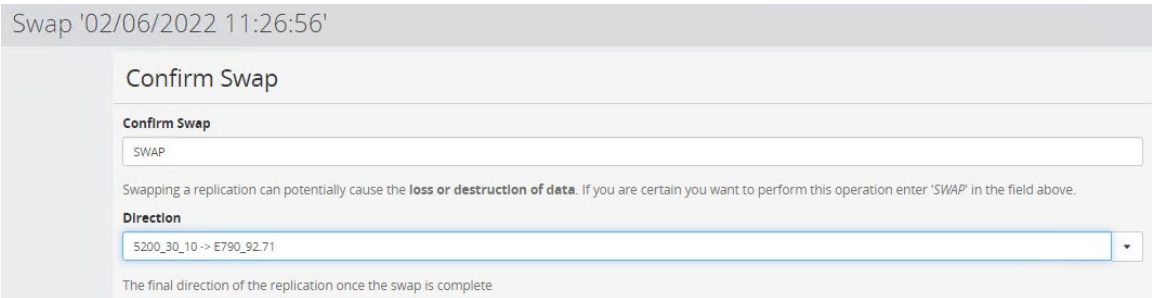

After the Swap operation is done, the roles are swapped. Primary volumes are now S-VOLs and secondary volumes are now P-VOLs.

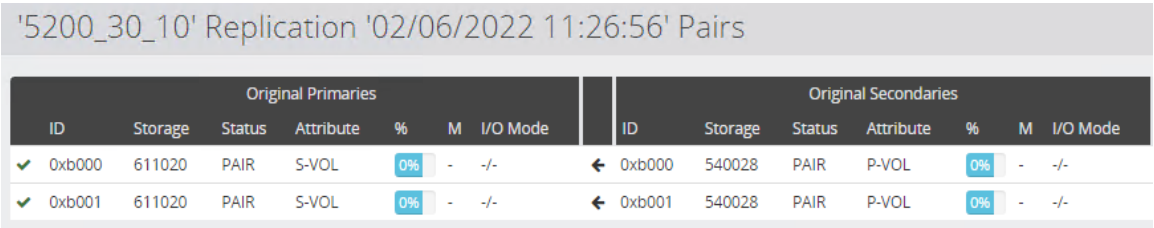

9. When you are ready to restore SQL to the primary site:
   a. In the original EC2 virtual machine SSMS, detach the databases.
   b. In Protector, click **Storage** and then click the secondary storage system.
   c. Click **Replication and Clones**. Note that the current status show is **Swapped**.
   d. Select the relationship and click **Swap**.

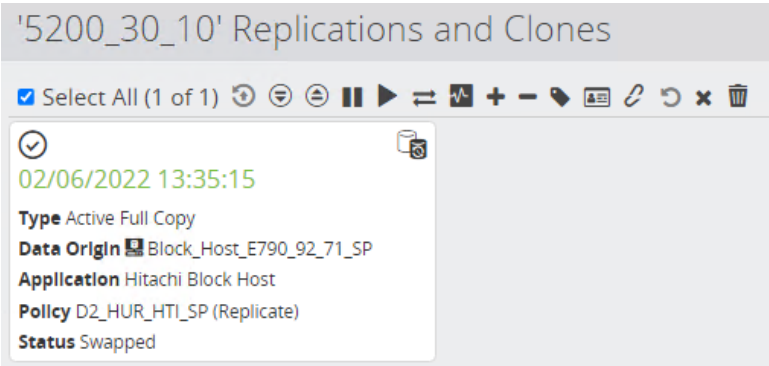

   e. Type **SWAP** to confirm the task. In the Direction field, select the original direction and then click **Finish**.

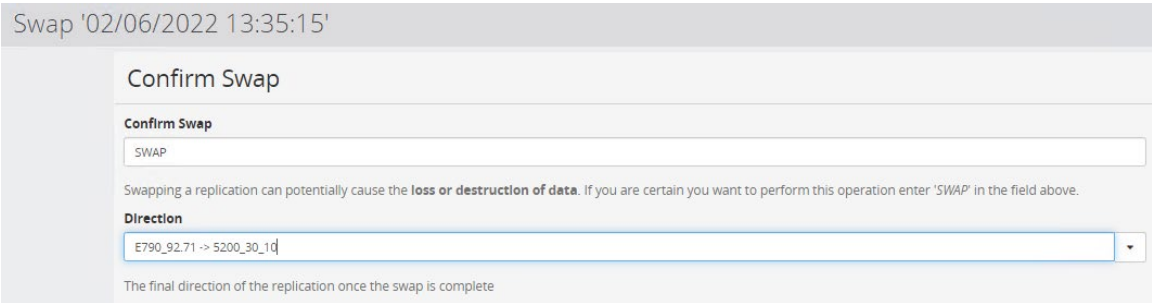

After the swap is done, the original primary volumes are now P-VOLs again, and original secondary volumes are S-VOLs again.

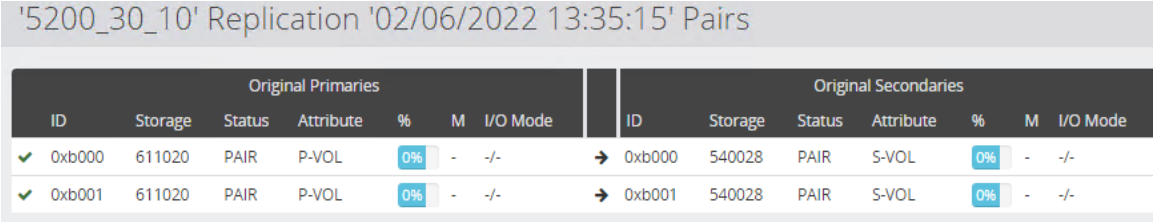

The Status field shows **OK** instead of Swapped now.

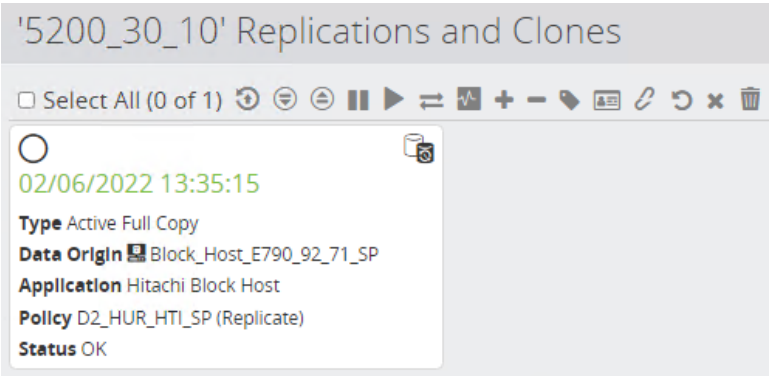

10. Power on the SQL virtual machines at the primary site.
    a. After several minutes, check Windows Failover Cluster Manager to ensure the nodes are online.

    b. Go to **Roles** and start SQL.

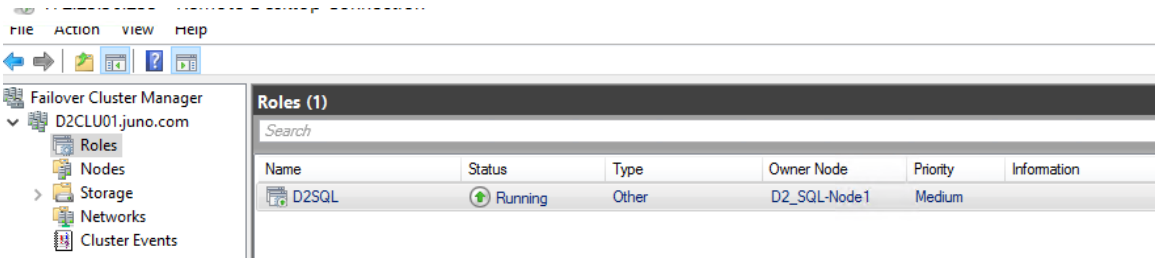

    c. Go to **Storage > Disks** and bring the disks online.

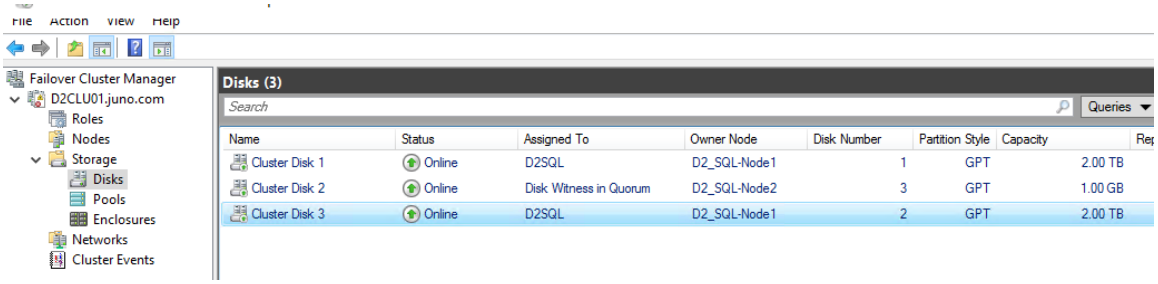

11. Verify that the new records written to SQL at the DR site are present.
    a. In SSMS, import the databases by attaching the database files. There are two databases: D2_Test1.6_DB1 (created while running on HUR primary volumes) and D2_Test1.6_DB2

(created while running on HUR secondary volumes).

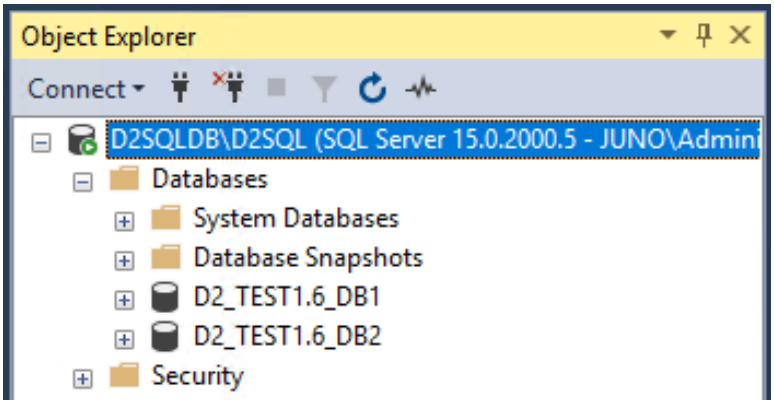

b. Check the content of the databases.

Results | Messages

	timestamp	Staff_ID	First_name	Last_Name	Email	Update_On
1	2022-02-06 10:14:33.0000000	80460	Letha	Wahl	Bolt@nowhere.com	HUR_P-VOLs
2	2022-02-06 10:14:33.0000000	19222	Kraig	Boucher	Morehead@example.com	HUR_P-VOLs
3	2022-02-06 10:14:33.0000000	86849	Harlan	Ludwig	qpzre411@example.com	HUR_P-VOLs
4	2022-02-06 10:14:33.0000000	08276	Brad	Craven	Bray162@example.com	HUR_P-VOLs
5	2022-02-06 10:14:33.0000000	23003	Enrique	Sizemore	Sidney_Agnew3@nowhere.com	HUR_P-VOLs
6	2022-02-06 10:14:33.0000000	68767	Bennie	Flowers	xmhz1900@nowhere.com	HUR_P-VOLs
7	2022-02-06 10:14:33.0000000	47864	Francis	Conway	Monroe.Rinehart@nowhere.com	HUR_P-VOLs
8	2022-02-06 10:14:33.0000000	44963	Nickie	Waite	Gutierrez73@example.com	HUR_P-VOLs
9	2022-02-06 10:14:33.0000000	15878	Boyd	Skaggs	FelipaAbernathy@nowhere.com	HUR_P-VOLs
10	2022-02-06 10:14:33.0000000	40093	Alonzo	Peacock	Adler@example.com	HUR_P-VOLs
11	2022-02-06 10:14:33.0000000	75351	Reed	Hwang	uadm55@example.com	HUR_P-VOLs

Query executed successfully. D2SQLDB\D2S

Results | Messages

	timestamp	Store_ID	Store_Name	City	Update_On
1	2022-02-06 12:45:36.0000000	80460	Conlifiedgor	Boca Raton	HUR_S-VOLs
2	2022-02-06 12:45:36.0000000	19222	Projectry	University Place	HUR_S-VOLs
3	2022-02-06 12:45:36.0000000	86849	Charcessridge	North Royalton	HUR_S-VOLs
4	2022-02-06 12:45:36.0000000	08276	Ancyclfiphone	Atlantic	HUR_S-VOLs
5	2022-02-06 12:45:36.0000000	23003	Transcessletor	Marquette	HUR_S-VOLs
6	2022-02-06 12:45:36.0000000	68767	Confinder	Rochester	HUR_S-VOLs
7	2022-02-06 12:45:36.0000000	47864	Speakcessfiry	Lindsay	HUR_S-VOLs
8	2022-02-06 12:45:36.0000000	44963	Transbanderator	Montgomery	HUR_S-VOLs
9	2022-02-06 12:45:36.0000000	15878	Transculgaentor	Allen Park	HUR_S-VOLs
10	2022-02-06 12:45:36.0000000	40093	Amputplry	Monroe	HUR_S-VOLs
11	2022-02-06 12:45:36.0000000	75351	Amptecticlet	Surfside	HUR_S-VOLs

Query executed successfully.

Test 7: Restore from Snapshots

This test case demonstrates how HTI snapshots can be used to restore data to primary volumes, even across sites.

1. Generate new records on the HUR primary volumes by creating a database named D2_Test1.7_DB1 and a table named Product_Table. The Update_On column indicates these records were created on the HUR P-VOLs.

Results | Messages

	timestamp	Product_ID	Product_Name	Product_Price	Update_On
1	2022-02-03 06:10:03.0000000	00000	Conputor	859179	HUR_P-VOLs
2	2022-02-03 06:10:03.0000000	00001	Charwoofaqury	869131	HUR_P-VOLs
3	2022-02-03 06:10:03.0000000	00002	Translictor	522031	HUR_P-VOLs
4	2022-02-03 06:10:03.0000000	00003	Tabtectar	281539	HUR_P-VOLs
5	2022-02-03 06:10:03.0000000	00004	Computar	285000	HUR_P-VOLs
6	2022-02-03 06:10:03.0000000	00005	Propickedscope	968010	HUR_P-VOLs
7	2022-02-03 06:10:03.0000000	00006	Stereowoofaner	470718	HUR_P-VOLs
8	2022-02-03 06:10:03.0000000	00007	Micbander	380548	HUR_P-VOLs
9	2022-02-03 06:10:03.0000000	00008	Speakleletentor	638867	HUR_P-VOLs
10	2022-02-03 06:10:03.0000000	00009	Stereowooflet	42193	HUR_P-VOLs
11	2022-02-03 06:10:03.0000000	00010	Cleanholdar	411675	HUR_P-VOLs
12	2022-02-03 06:10:03.0000000	00011	Armlifiry	354067	HUR_P-VOLs
13	2022-02-03 06:10:03.0000000	00012	Tweettectollar	648948	HUR_P-VOLs
14	2022-02-03 06:10:03.0000000	00013	Comcycler	787808	HUR_P-VOLs
15	2022-02-03 06:10:03.0000000	00014	Stereocordexry	578809	HUR_P-VOLs

Query executed successfully.

2. Confirm that the HUR pairs are in PAIR state.

'5200_30_10' Replication '02/03/2022 08:35:55' Pairs

Original Primaries								Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode		ID	Storage	Status	Attribute	%	M	I/O Mode
0xb000	611020	PAIR	P-VOL	0%	-	-/-	→	0xb000	540028	PAIR	S-VOL	0%	-	-/-
0xb001	611020	PAIR	P-VOL	0%	-	-/-	→	0xb001	540028	PAIR	S-VOL	0%	-	-/-

3. Trigger HTI snapshots.
   a. In Protector, click **Monitor** and click the data flow.
   b. Select the secondary storage system and click **Trigger**.
   c. After the snapshots are created, get the HTI snapshot details.
   d. In Administrator, **Refresh** the secondary storage system.
   e. Map the snapshot volumes as described in section Test 3.
4. In EC2, bring the drives online and attach the existing database, D2_Test1.7_DB1, that was created on the HUR primary volumes.

5. Create a new database and ingest some records. The following screenshot shows a new table and new database were created on the HTI snapshot volumes: Customer_Table in D2_Test1.7_DB2.

Results | Messages

	timestamp	Customer_Name	Customer_Address	Update_On
1	2022-02-03 11:47:13.0000000	Letha Bolt	240 Hidden Front Way	HTI_S-VOLs
2	2022-02-03 11:47:13.0000000	Vanita Kelleher	2948 Stonewood Pkwy	HTI_S-VOLs
3	2022-02-03 11:47:13.0000000	Marcos Abraham	44 West Church Ct	HTI_S-VOLs
4	2022-02-03 11:47:13.0000000	Miyoko Mckinney	3778 Edgewood Highway	HTI_S-VOLs
5	2022-02-03 11:47:13.0000000	Adolph Francisco	12 1st Circle	HTI_S-VOLs
6	2022-02-03 11:47:13.0000000	Oma Lawler	3570 Flintwood Parkway	HTI_S-VOLs
7	2022-02-03 11:47:13.0000000	Felix Alba	1790 SW Church Avenue	HTI_S-VOLs
8	2022-02-03 11:47:13.0000000	Eusebia Noland	2014 Social Road	HTI_S-VOLs
9	2022-02-03 11:47:13.0000000	Denisha Mcginnis	34 Riddle Hill Pkwy	HTI_S-VOLs
10	2022-02-03 11:47:13.0000000	Stanford Sullivan	65 Waterview Avenue	HTI_S-VOLs
11	2022-02-03 11:47:13.0000000	Nathan Salisbury	3760 West Social Parkway	HTI_S-VOLs
12	2022-02-03 11:47:13.0000000	Burton Bunnell	1254 New Cedar Tree Rd	HTI_S-VOLs
13	2022-02-03 11:47:13.0000000	Adah Jameson	2354 N Cedar Tree Hwy	HTI_S-VOLs
14	2022-02-03 11:47:13.0000000	Adam Geer	95 North Church Road	HTI_S-VOLs
15	2022-02-03 11:47:13.0000000	Freeman Selby	47 W Ashwood Avenue	HTI_S-VOLs
16	2022-02-03 11:47:13.0000000	Laura Barger	2455 Buttonwood Ln	HTI_S-VOLs
17	2022-02-03 11:47:13.0000000	Ariel Hermann	10 Social Lane	HTI_S-VOLs
18	2022-02-03 11:47:13.0000000	Adam Maddox	633 New Parkwood St	HTI_S-VOLs

Query executed successfully. D2_SQL_EC2_1

6. Corrupt the database at the primary site as follows:
   a. In SSMS, delete the database instance (D2_Test1.7_DB1).
   b. In Windows Failover Cluster Manager, stop SQL and shut down the cluster.
   c. Power off the virtual machines.
7. Now we will use Protector to restore the primary database with the HTI snapshots.
   a. The first step is to suspend the HUR pairs and swap their roles.
      i. In Protector, click **Storage**, and then click the secondary storage system.
      ii. Click **Replication and Clones**.
      iii. Select the HUR relationship and click **Swap**.
      iv. Type **SWAP** to confirm the task. In the Direction field, select the direction that is opposite of the current direction and click **Finish**.

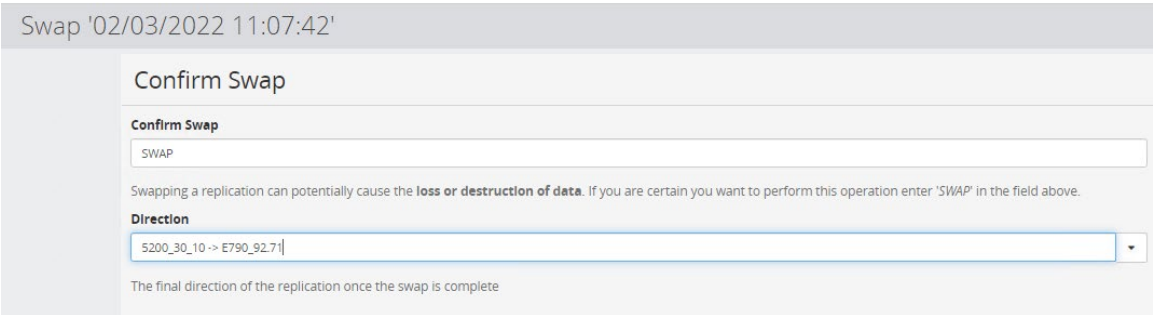

After the operation has completed, the **Status** field is **Swapped**.

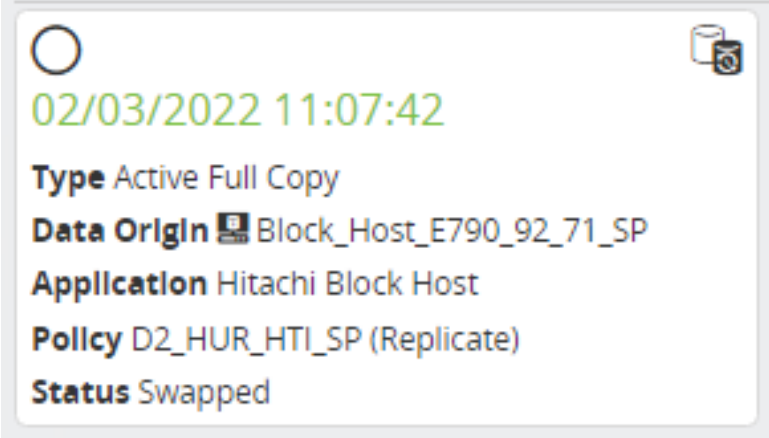

v. Now **Pause** the HUR replication.

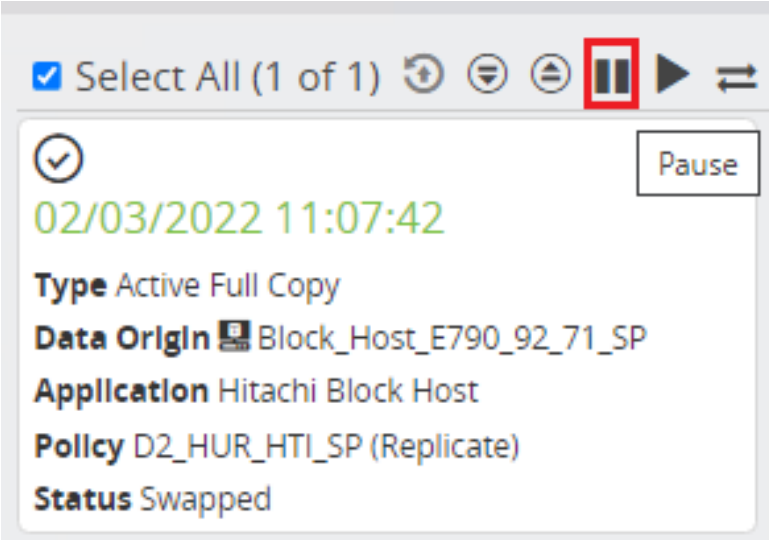

After the operation has completed, the **Status** field is **Paused, Swapped**.

vi. Verify that the HUR primary volumes are now in SSUS state and secondary volumes are in PSUS state.

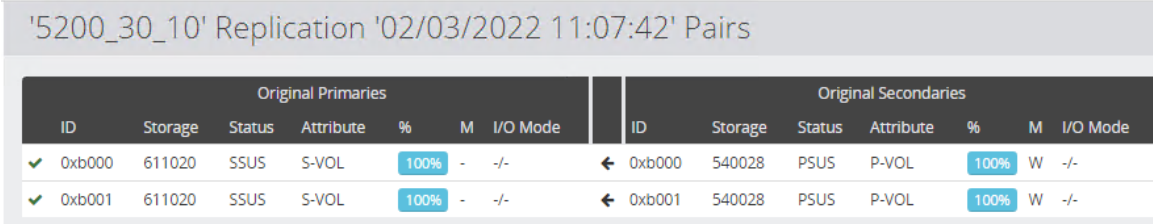

b. In the EC2 virtual machine, detach the databases and offline the volumes.

c. In Administrator, un-map the volumes from the EC2 virtual machine by selecting the volumes and clicking **Detach**.

d. Begin the process of restoring data from the HTI snapshots to the HUR pairs by reverting the snapshots.

i. In Protector, select the HTI relationship and click **Revert**.

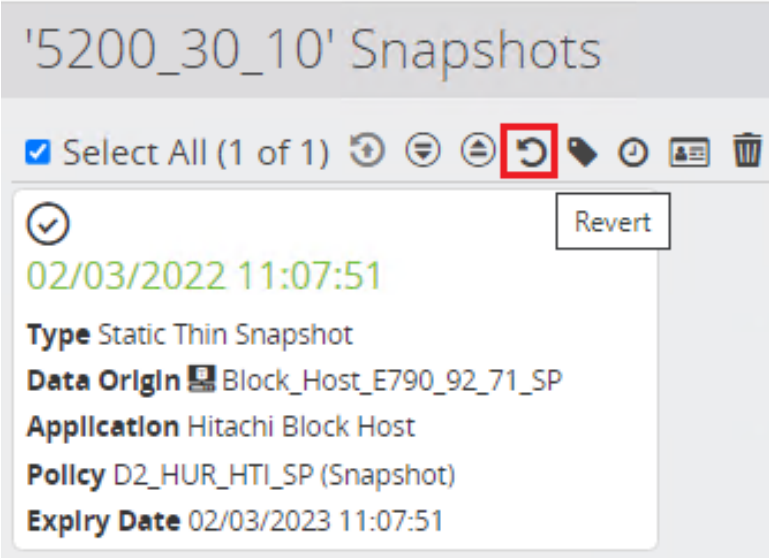

ii. Type **REVERT** to confirm.

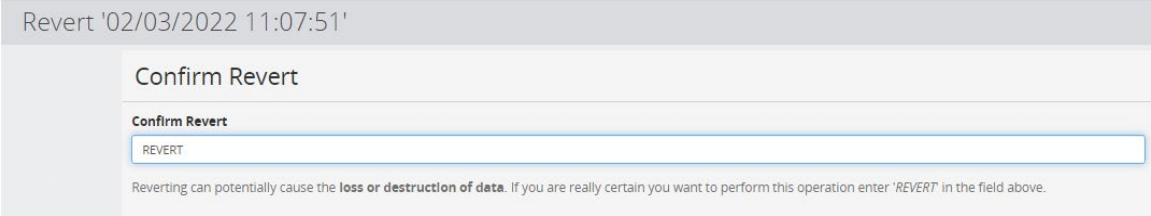

iii. Confirm the operation completed by verifying that Status shows Completed Successfully.

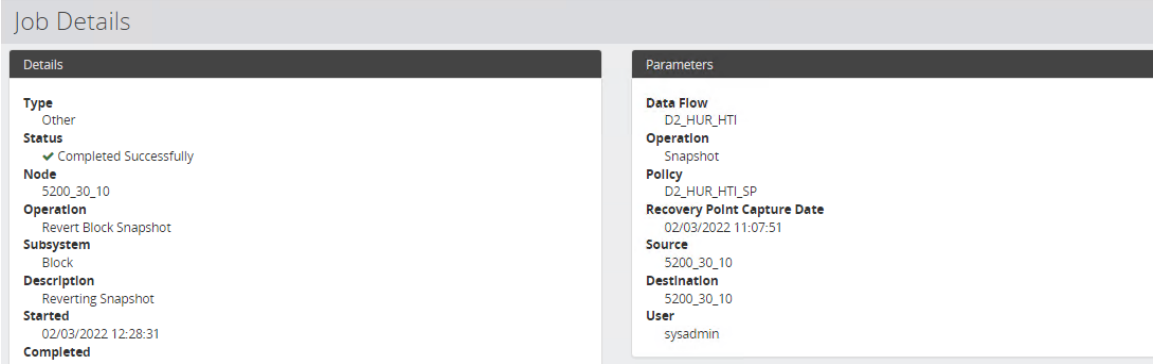

e. The final step is to resume HUR replication and reverse HUR roles to their original designation.

i. Select the HUR relationship and click **Resume**.

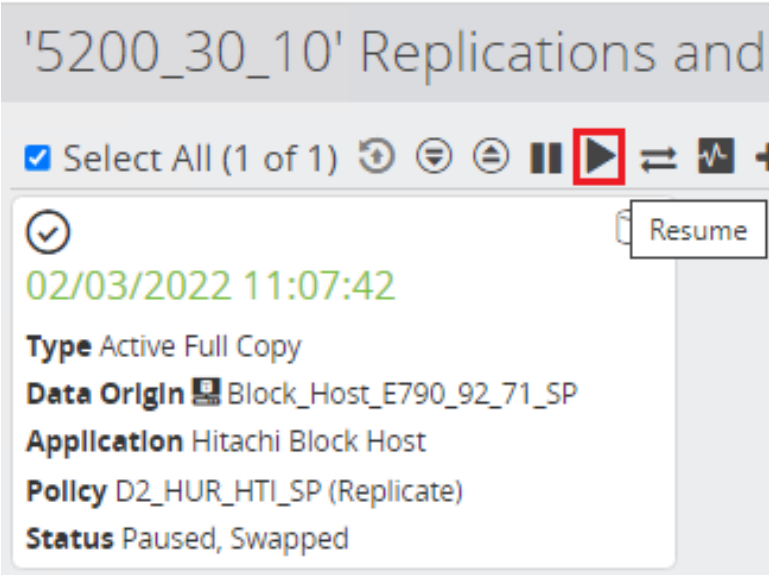

ii. After the resume operation is done, **Swap** the HUR pairs to the original direction, and after the swap is done, the original primary volumes are now P-VOLs again, and original secondary volumes are S-VOLs again.

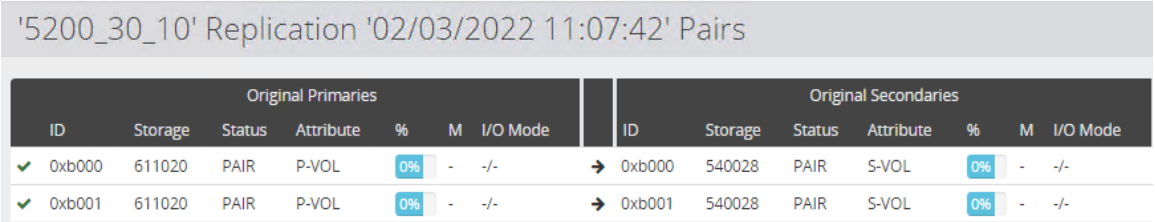

8. Power on the SQL virtual machines at the primary site.
   a. After several minutes, check Windows Failover Cluster Manager to ensure the nodes are online.

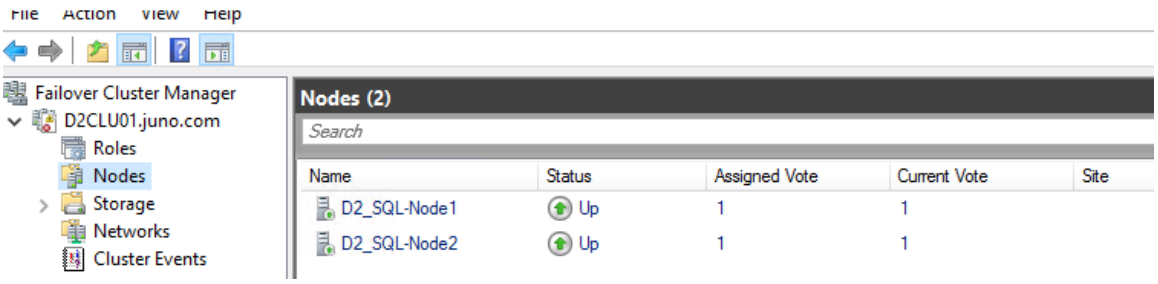

   b. Go to **Roles** and start SQL.

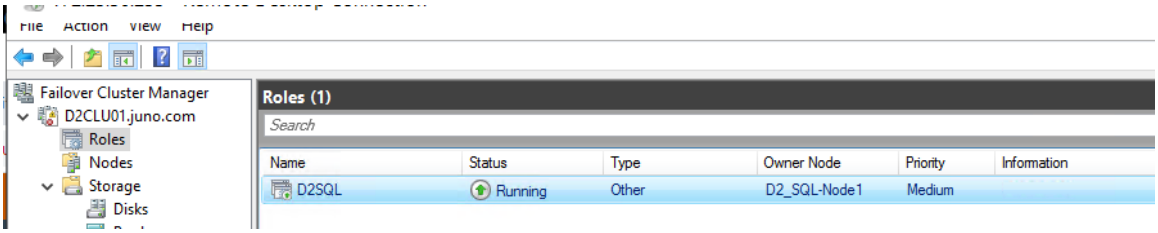

   c. Go to **Storage > Disks** and bring the disks online
9. Verify that the previously deleted database, D2_Test1.7_DB1, is back and the new table, D2_Test1.7_DB2, created on EC2 is also present.

a. In SSMS, import the databases by attaching the database files. The following screenshot shows both databases are present.

10	D2_TEST1.7_DB1	G:\MSSQL15.D2SQL\MSSQL\DATA\D2_TEST1.7_DB1.mdf
11	D2_TEST1.7_DB1_log	E:\LOG\D2_TEST1.7_DB1_log.ldf
12	D2_TEST1.7_DB2	G:\MSSQL15.D2SQL\MSSQL\DATA\D2_TEST1.7_DB2.mdf
13	D2_TEST1.7_DB2_log	E:\LOG\D2_TEST1.7_DB2_log.ldf

b. Check the content of the databases.

Results | Messages

	timestamp	Product_ID	Product_Name	Product_Price	Update_On
1	2022-02-03 06:10:03.0000000	00000	Conputor	859179	HUR_P-VOLs
2	2022-02-03 06:10:03.0000000	00001	Charwoofaqury	869131	HUR_P-VOLs
3	2022-02-03 06:10:03.0000000	00002	Translictor	522031	HUR_P-VOLs
4	2022-02-03 06:10:03.0000000	00003	Tabtectar	281539	HUR_P-VOLs
5	2022-02-03 06:10:03.0000000	00004	Computar	285000	HUR_P-VOLs
6	2022-02-03 06:10:03.0000000	00005	Propickedscope	968010	HUR_P-VOLs
7	2022-02-03 06:10:03.0000000	00006	Stereowoofaner	470718	HUR_P-VOLs
8	2022-02-03 06:10:03.0000000	00007	Micbander	380548	HUR_P-VOLs
9	2022-02-03 06:10:03.0000000	00008	Speakleletentor	638867	HUR_P-VOLs
10	2022-02-03 06:10:03.0000000	00009	Stereowooflet	42193	HUR_P-VOLs
11	2022-02-03 06:10:03.0000000	00010	Cleanholdar	411675	HUR_P-VOLs

Query executed successfully.

Results | Messages

	timestamp	Customer_Name	Customer_Address	Update_On
1	2022-02-03 11:47:13.0000000	Letha Bolt	240 Hidden Front Way	HTI_S-VOLs
2	2022-02-03 11:47:13.0000000	Vanita Kelleher	2948 Stonewood Pkwy	HTI_S-VOLs
3	2022-02-03 11:47:13.0000000	Marcos Abraham	44 West Church Ct	HTI_S-VOLs
4	2022-02-03 11:47:13.0000000	Miyoko Mckinney	3778 Edgewood Highway	HTI_S-VOLs
5	2022-02-03 11:47:13.0000000	Adolph Francisco	12 1st Circle	HTI_S-VOLs
6	2022-02-03 11:47:13.0000000	Oma Lawler	3570 Flintwood Parkway	HTI_S-VOLs
7	2022-02-03 11:47:13.0000000	Felix Alba	1790 SW Church Avenue	HTI_S-VOLs
8	2022-02-03 11:47:13.0000000	Eusebia Noland	2014 Social Road	HTI_S-VOLs
9	2022-02-03 11:47:13.0000000	Denisha Mcginnis	34 Riddle Hill Pkwy	HTI_S-VOLs
10	2022-02-03 11:47:13.0000000	Stanford Sullivan	65 Waterview Avenue	HTI_S-VOLs
11	2022-02-03 11:47:13.0000000	Nathan Salisbury	3760 West Social Parkway	HTI_S-VOLs

Query executed successfully.

Hitachi Vantara

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

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

Author: D/Luong, S/Pal, March 2022