

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	3
Executive Summary	4
About This Guide	5
Introduction.....	5
Intended Audience.....	5
Document Revisions.....	5
References.....	5
Comments.....	5
Solution Overview	6
Benefits.....	6
Key Components.....	6
Engineering Validation	8
Validation Method.....	8
High Level Diagram.....	8
Hardware and Software.....	9
Test Scenarios.....	10
Guidelines and Recommendations	12
Validation Results	13
Test 1: Prepare the Environment.....	13
Test 2: Create HUR Pairs with Cascaded HTI Snapshots.....	17
Test 3: Access Snapshots from AWS EC2.....	33
Test 4: Multigeneration HTI Snapshots.....	38
Test 5: Planned Outage.....	42
Test 6: Unplanned Outage.....	50
Test 7: Restore from Snapshots.....	59

Notices and Disclaimer

© 2022 Hitachi Vantara LLC. All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

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

Executive Summary

Hitachi 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.

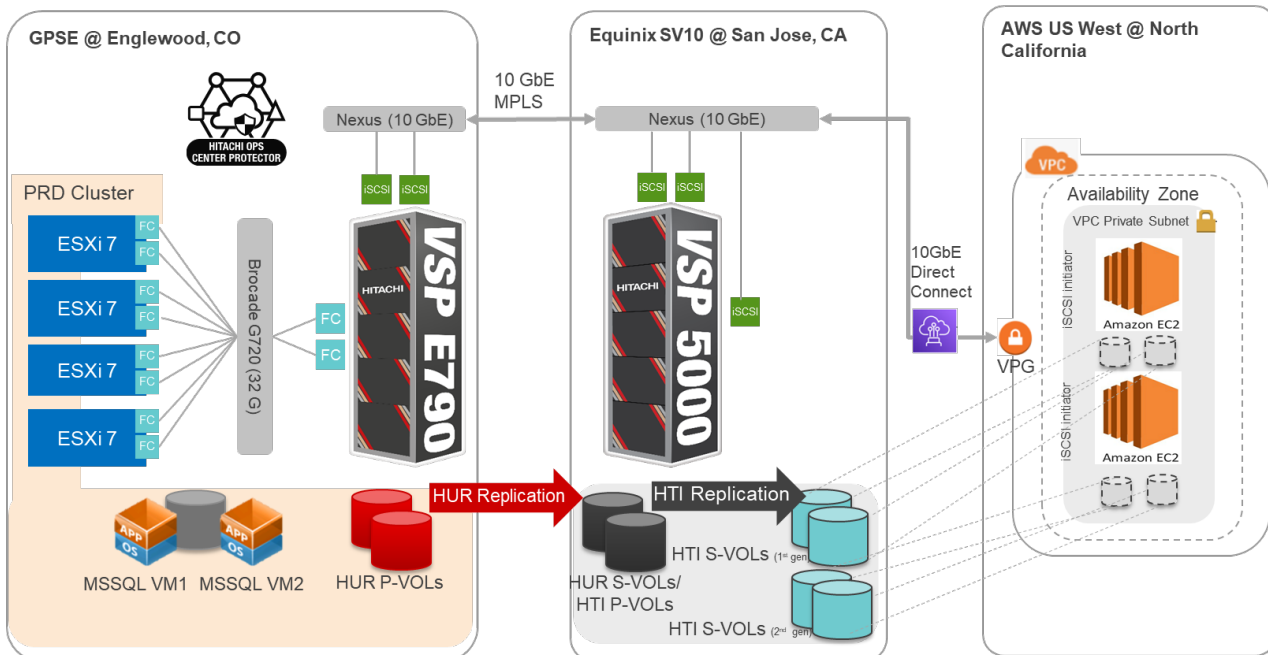


Figure 2. Test Environment Reference Architecture

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

Reference Architecture

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: <ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> 1. Create HUR pairs using Ops Center Protector. 2. Create HTI snapshots using Ops Center Protector. 	<ul style="list-style-type: none"> • HUR pairs are created. • HTI snapshots are created.
3	Access snapshots from AWS EC2: <ol style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> • Data created at primary site is present on snapshots. • EC2 virtual machine can write to snapshots without affecting primary site.
4	Multigeneration HTI snapshots: <ol style="list-style-type: none"> 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 2 nd generation snapshots contains records written by the primary site but not records from the 1 st generation snapshots.
5	Planned outage with data return: <ol style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> • Snapshots can still be taken while HUR roles are swapped. • Data created at DR site is

Reference Architecture

Test	Description	Success Criteria
	<ol style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> • returned to primary site.
6	<p>Unplanned outage with data return:</p> <ol style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> • Snapshots can still be taken while HUR roles are swapped. • Data created at DR site is returned to primary site.
7	<p>Restore from snapshots:</p> <ol style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> • 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.

- 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...
+												
<input type="radio"/>	45056 (00:B...	45056 (00:B...	611020	611020	D2_RDM_0	0	Compressio...	Normal	Deduplicati...	2.00 TIB	3%	Unmanaged
<input type="radio"/>	45057 (00:B...	45057 (00:B...	611020	611020	D2_RDM_1	0	Compressio...	Normal	Deduplicati...	2.00 TIB	1%	Unmanaged
<input type="radio"/>	45058 (00:B...	45058 (00:B...	611020	611020	D2_Win_Qu...	0	Thin	Normal	No	1.00 GIB	16%	Unmanaged
<input type="radio"/>	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.

- 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

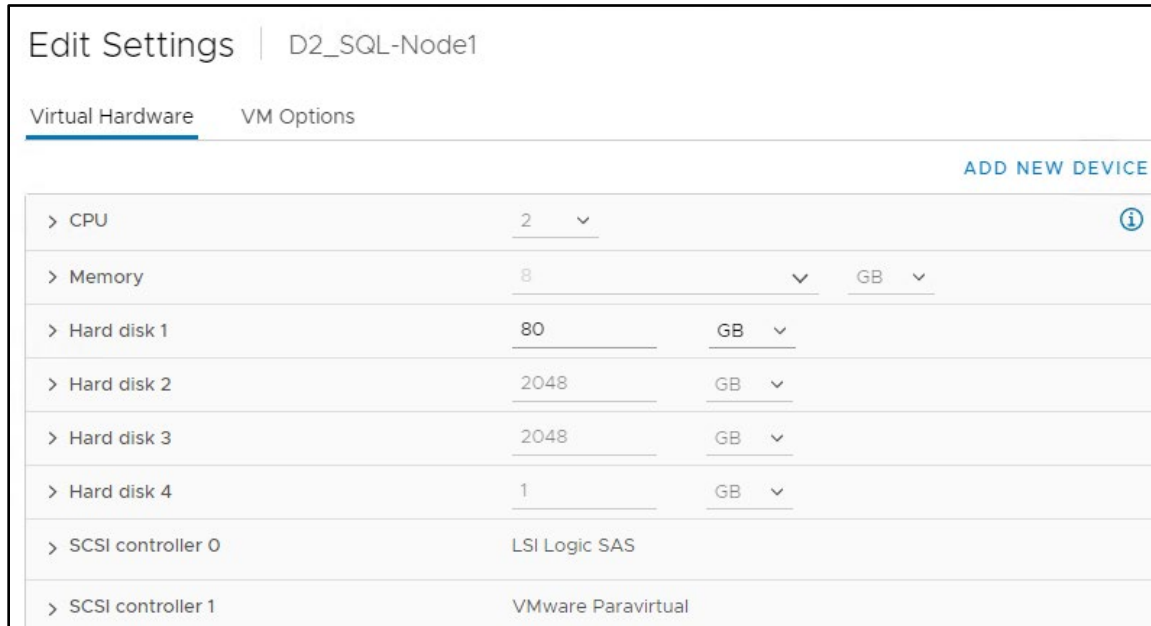
Remote Paths

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.


- 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.
- 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:



- 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.



- 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.

```

SQLQuery2.sql - D2S...administrator (53)*
USE master;
SELECT
    name 'Logical Name',
    physical_name 'File Location'
FROM sys.master_files;

```

	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_DATABASE...	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:


Object Explorer

Connect

- D2_SQL_EC2_1\EC2SQL1 (SQL Server 15.0.2000.5 - JUNO\Administrator)
 - Databases
 - Security
 - Server Objects
 - Replication
 - PolyBase
 - Always On High Availability
 - Management
 - Integration Services Catalogs
 - SQL Server Agent (Agent XPs disabled)
 - XEvent Profiler

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

 **Add Servers**

CSV Import

+

Fibre Servers

+

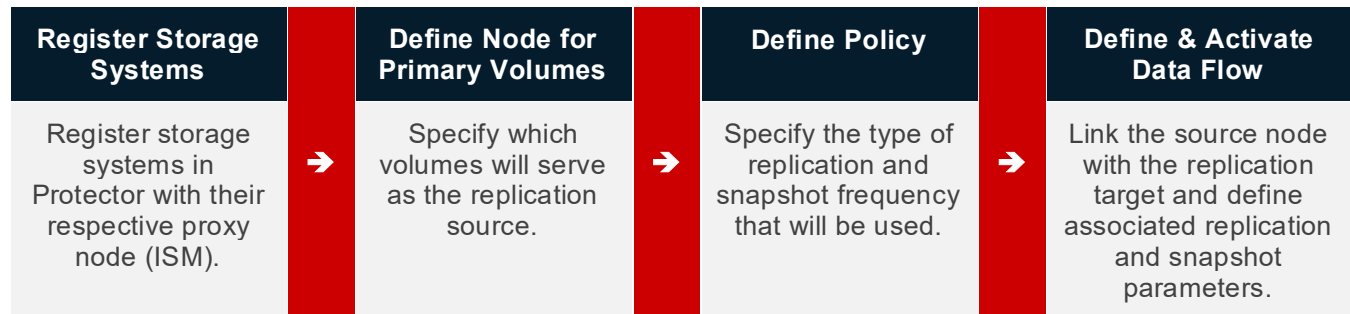
iSCSI Servers

+

SERVER NAME	DESCRIPTION	IP ADDRESS	OS TYPE
D2_SQL_EC2_1.juno.com	Description	172.77.25.233	WIN ▼
CHAP SECRET	iSCSI LIST		iSCSI NAMES
			iqn.1991-05.com.microsoft:d2-sql-ec2-1.juno.com

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.



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**.

Create Node - Hitachi Block Device

Specify Node name

Node Name

E790_92.71

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

Tags

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

Resources or replication relationships created or adopted by Protector must only be managed, modified and deleted via Protector. Failure to do so will cause unpredictable consequences and is not supported unless specifically advised to do so by the Protector documentation.

I confirm that I have read and understood this requirement

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

Create Node - Hitachi Block Device

Allocate node to Access Control Resource Group

This node will be added to the 'default' resource group. Select an additional resource group as required.

Name	Description
No available Access Control Resource Groups.	

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

Create Node - Hitachi Block Device

Select proxy node

Proxy Node

probehos

Proxy node requires a block command device configured and CCI / RAIDCOM management tools installed.

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

Create Node - Hitachi Block Device

Select Metadata directory

Common Metadata Directory

Note: The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.

- g. 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**.

Create Node - Hitachi Block Device

Specify Device

Select from detected storage devices

Recommended in environments where performance and reliability are critical.
If the desired serial does not appear in the above list then it may not have an available fiber command device on the selected proxy.

Specify by IP or Hostname with a port

To be used as a fall-back option, or if fiber command devices cannot be configured.

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

Create Node - Hitachi Block Device

Specify credentials for device

Storage Device Serial Number

Username

Password

Protector does not support using passwords which contain some special characters. See CCI / RAIDCOM documentation for further details.

The device account requires the following roles: Storage Administrator (Provisioning, Local Copy, Remote Copy), Security Administrator (View Only, View and Modify) and Support Personnel.

- i. In the Specify Configuration for Global Replication Reports screen, click **Next**.
- j. Specify **All** for LDEV Range and click **Next**.

Create Node - Hitachi Block Device

Specify LDEV Provisioning Range

LDEV Range

All

User defined

Start

End

- k. Continue to click **Next** until you reach the summary screen.
- l. Review the information in the summary screen and then click **Finish**.

Create Node - Hitachi Block Device

Summary of 'E790_92.71'

Proxy Node
probehost

Storage Device Serial Number
611020

Username
maintenance

LDEV Provisioning Range
All Available

Configured Command Devices

Type	LDEV ID	IP Address	Port
IP	-	172.23.67.92	31001

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**.

Create Node - Hitachi Block Host

Specify Node name

Node Name

Block_Host_E790_92_71_5P

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

Tags

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

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

Create Node - Hitachi Block Host

Select Hitachi Block Device

Hitachi Block Device

E790_92.71

- 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**.

Create Node - Hitachi Block Host

Specify Logical Devices

Enter Logical Devices using any of the following formats:


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

Included Logical Devices

45056-45057

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

Create Node - Hitachi Block Host

 Summary of 'Block_Host_E790_92_71_SP'

Hitachi Block Device
E790_92.71

Logical Devices
45056-45057

Excluded Logical Devices

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

Create Policy

Specify name and description

Name
D2_HUR_HTI_SP

Description
This policy is to create cascaded HUR and HTI pairs with two 2TB volumes B0:00 and B0:01. HUR S-VOLS will be shared with HTI P-VOLS. HTI S-VOLS/ HUR -SVOLS will be mounted to EC2 Instance as iSCSI target.

Tags

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

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

Create Policy

Specify Hitachi Block Storage classification attributes

Use Hitachi Block Host selections
 Specify additional selections

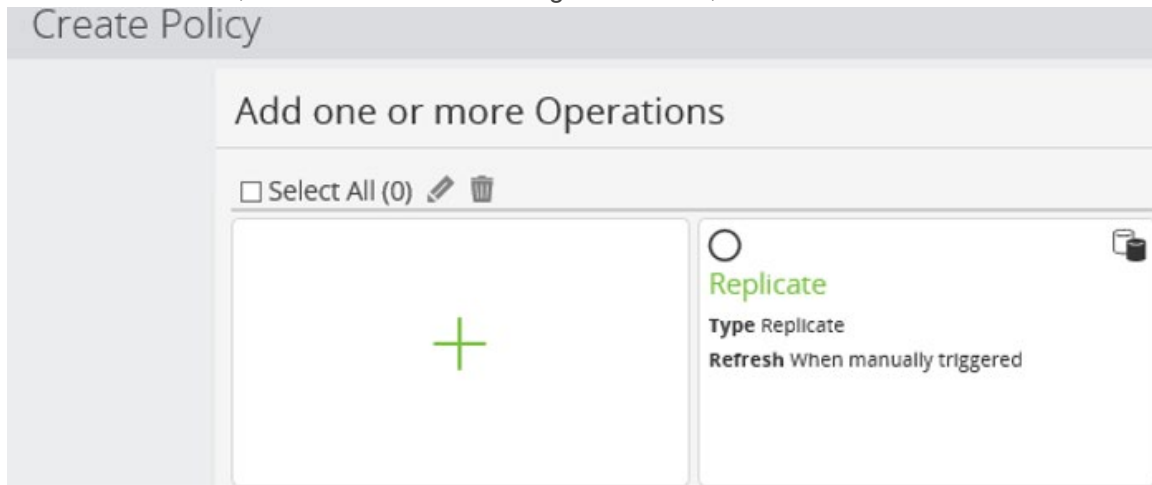
Enter values using any of the following formats:

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

- 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**.

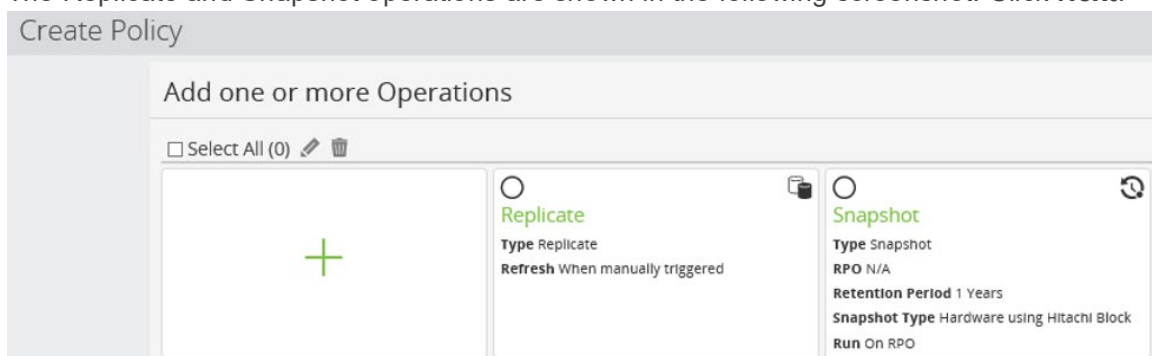


- k. In the Add One or More Operations screen, click the + (plus) symbol again to add a snapshot operation.
- l. Select **Snapshot** and click **Next**.
- m. In the **Specify snapshot operation attributes** screen, set **Recovery Point Objective=None** and **Retention=1 year**. Deselect **Quiesce...** 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**.

Create Data Flow

Specify name and description

Name

Description

Tags

Enter comma separated tags. Tags can include alphanumeric characters, underscore and hyphen. To

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

Create Data Flow

Nodes Node Groups

Node Name

Select a Node Type

- 5200_30.10
- Block_Host_E790_92_7...
- E790_92.71
- opscenter
- probehost
- sv10probehost

100%

Block_Host_E790_92_71_SP
Hitachi Block Host

Policies

- D2_HUR_HTI_SP
- Snapshot

- 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.

Create Data Flow

Nodes Node Groups

Node Name

Select a Node Type

- 5200_30.10
- Block_Host_E790_92_7...
- E790_92.71
- opscenter
- probehost
- sv10probehost

100%

Block_Host_E790_92_71_SP
Hitachi Block Host

5200_30.10
Hitachi Block Device

Routed Policies

No valid policies are being routed to this mover.

Mover Settings

Transfer Type

Batch

Label

Enable network data compression

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

The screenshot shows the 'Create Data Flow' interface. On the left, a list of nodes includes '5200_30.10', 'Block_Host_E790_92_71...', 'E790_92.71', 'opscnter', 'probehost', and 'sv10probehost'. The main canvas shows a data flow from 'Block_Host_E790_92_71_SP' (Hitachi Block Host) to '5200_30.10' (Hitachi Block Device). On the right, the 'Mover Settings' panel is visible, with 'Transfer Type' set to 'Continuous'. Other settings include 'Label', 'Enable network data compression', and 'Bandwidth Settings'.

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

The screenshot shows the 'Create Data Flow' interface. The source node 'Block_Host_E790_92_71_SP' is selected, indicated by a red triangle. On the right, the 'Policies' panel shows 'D2_HUR_HTI_SP' selected with a checkmark, and 'Snapshot' is unselected.

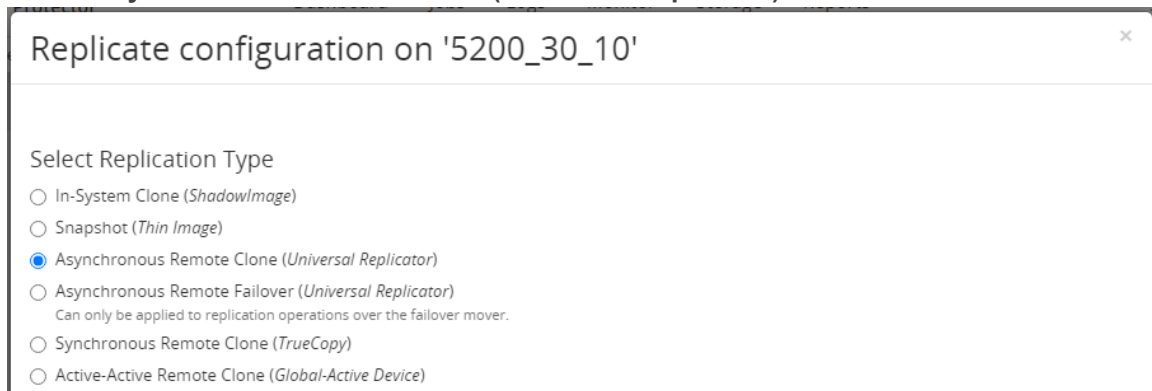
- g. Select the target node and then on the right, click **Replicate** first. This brings up a prompt for Replicate configuration options.

The screenshot shows the 'Create Data Flow' interface. The target node '5200_30.10' is selected, indicated by a blue square. On the right, the 'Policies' panel shows 'D2_HUR_HTI_SP' selected with a checkmark, and 'Replicate' is also selected with a checkmark.

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

The screenshot shows a dialog box titled 'Replicate configuration on '5200_30_10''. It contains the following text: 'Select Creation Mode' followed by two radio button options: 'Configure new replication' (which is selected) and 'Adopt an existing replication'.

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

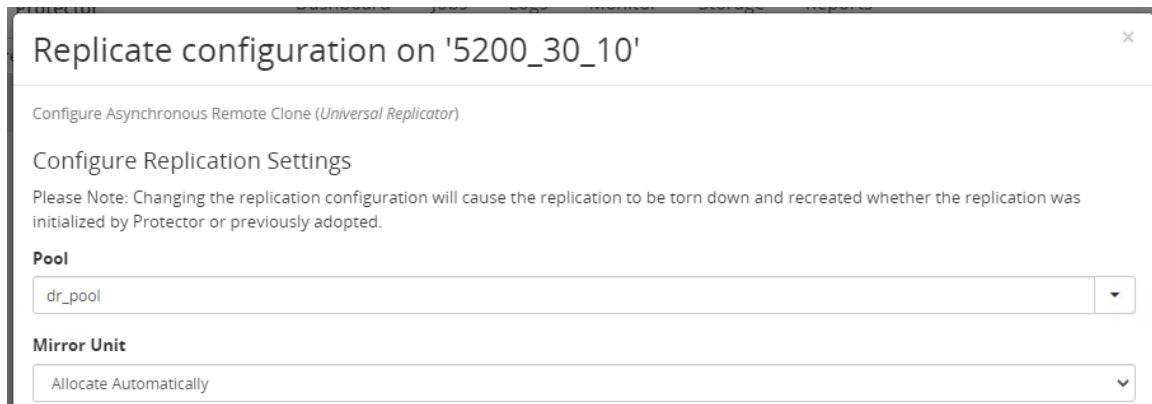


Replicate configuration on '5200_30_10'

Select Replication Type

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

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



Replicate configuration on '5200_30_10'

Configure Asynchronous Remote Clone (*Universal Replicator*)

Configure Replication Settings

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

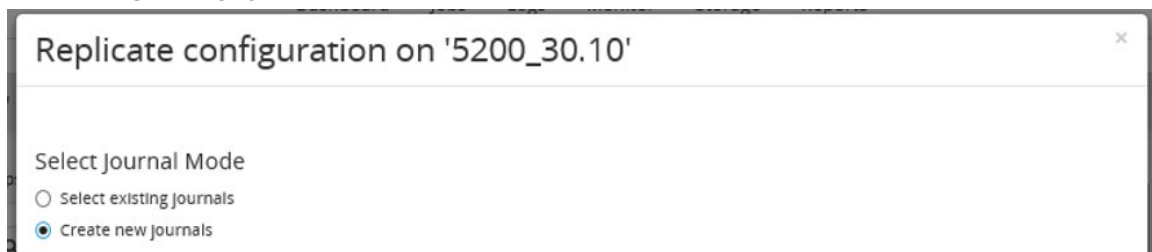
Pool

dr_pool

Mirror Unit

Allocate Automatically

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



Replicate configuration on '5200_30.10'

Select Journal Mode

- Select existing journals
- Create new journals

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

Replicate configuration on '5200_30.10'

Create Journals

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

Source Journal Pool

E790_92.71

prd_pool

Destination Journal Pool

dr_pool

Journal Sizes

400 GB

Journal Names

[JNL_%JOURNAL_ID%_%DATA_FLOW_NAME%

Journal names are limited to 32 characters, after variable resolution.
 Display variables which can be used for the journal name >



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**.

Replicate configuration on '5200_30_10'

Select Remote Path Group

Automatically Selected
 User Selected

Select Source Node

Select a Source Remote Path Group

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

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

Replicate configuration on '5200_30_10'

Configure Asynchronous Remote Clone (Universal Replicator)

Configure Resource Group

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

Select a Resource Group

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

Replicate configuration on '5200_30_10'

Configure Asynchronous Remote Clone (Universal Replicator)

Secondary Volume Host Groups

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

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

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

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

D2_SPAL (CL3-A-0)
▼ ✕

If replication SVOLs are exposed to a host, users must ensure that they are not in use during replication resynchronization, otherwise critical system failure may occur on the host machine.



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**.

Replicate configuration on '5200_30_10'

Configure Asynchronous Remote Clone (Universal Replicator)

Specify Naming Options

Secondary Logical Device Name

Match Origin
 Custom

Protector_LDEV_%SECONDARY_LDEV_ID%_%CREATION_DATE%%CREATION_TIME%

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

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

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

Replicate configuration on '5200_30_10'

Configure Asynchronous Remote Clone (Universal Replicator)

Configuration Summary

Type
Asynchronous Remote Clone (Universal Replicator)

Pool
dr_pool on 5200_30_10

Resource Group
Automatically Selected

Mirror Unit
0

On Destination Write Failure
Ignore

Source Journal
Automatically created in pool prd_pool

Destination Journal
Automatically created in pool dr_pool

Secondary Host Groups
D2_SPAL

Volume Naming
Automatically Generated

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

Snapshot configuration on '5200_30.10'

Configure Snapshot Settings

Storage Node
5200_30.10

Snapshot Pool
dr_pool (Dynamic Pool)

Advanced Configuration

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

Snapshot configuration on '5200_30.10'

Configure Resource Group

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

User Selected

Select a Resource Group

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

The screenshot shows a dialog box titled "Snapshot configuration on '5200_30.10'". Under the "Configure Provisioning Options" section, three checkboxes are checked: "Consistency group", "Fully provisioned" (with a sub-note: "Creates secondary devices for this snapshot. If not enabled floating devices will be used."), and "Cascade mode" (with a sub-note: "Allows the creation of snapshots / clones of this snapshot."). Below this is the "Cascade Pool" section with a text input field containing "dr_pool". A note at the bottom states: "When using fully provisioning devices with cascade mode, a cascade pool may be required."



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

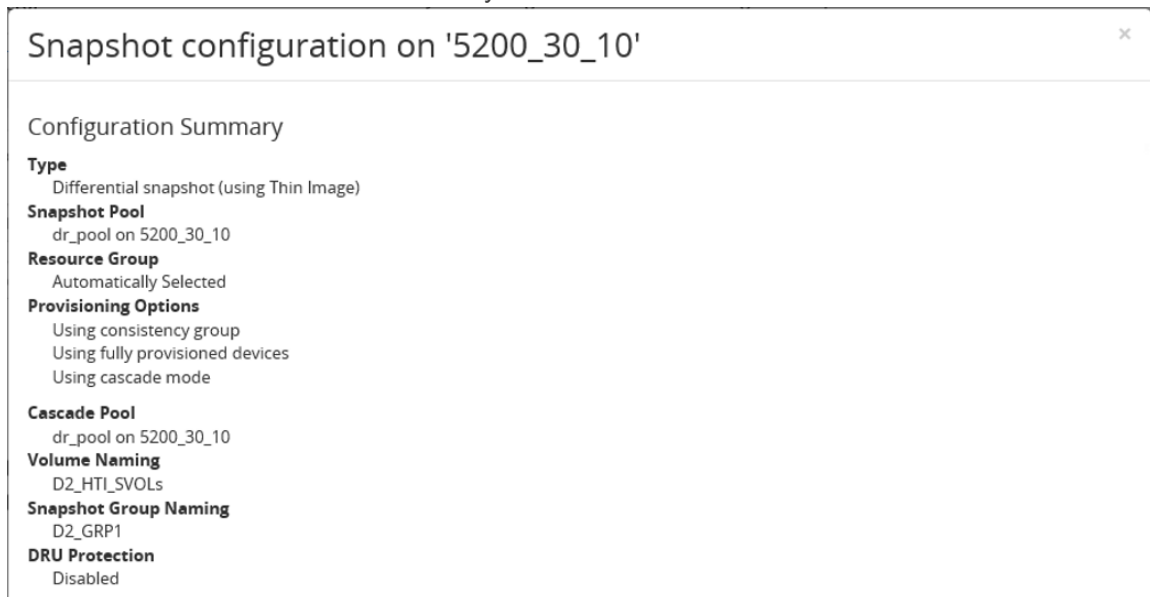
- u. In the Specify Naming Options screen, you can choose the naming option for the individual HTI secondary volumes as well as the naming option for the whole snapshot group. Click **Next**.

The screenshot shows a dialog box titled "Snapshot configuration on '5200_30_10'". Under the "Specify Naming Options" section, "Secondary Logical Device Name" has "Custom" selected (radio button), with a text input field containing "D2_HTI_SVOLs". A note below says: "Logical device names are limited to 32 characters, after variable resolution." Below that, "Display variables which can be used for the secondary LDEVs' name" has a right-pointing arrow. "Snapshot Group Name" has "Custom" selected (radio button), with a text input field containing "D2_GRP1". A note below says: "Snapshot group name is limited to 28 characters."

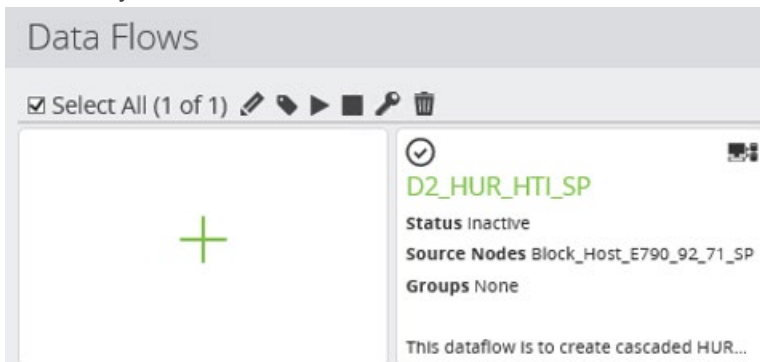
- v. For **Protection Type**, select **None**. Click **Next**.

The screenshot shows a dialog box titled "Snapshot configuration on '5200_30.10'". Under the "Configure DRU Options" section, a note says: "DRU protection is only available when using fully provisioned devices." The "Protection Type" dropdown menu is set to "None". Below it, "Duration of Settings Lock (Days)" is set to "0". A warning icon and note at the bottom state: "DRU settings cannot be removed until this duration has expired."

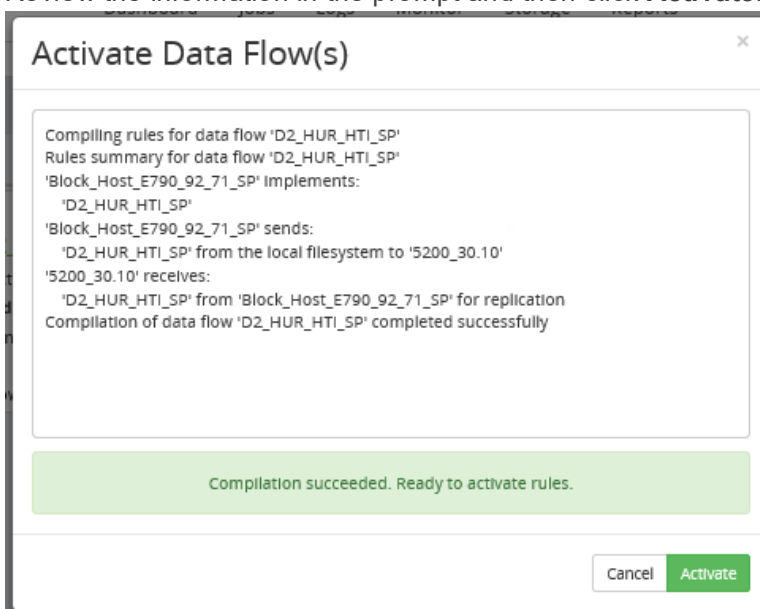
- w. Review the information on the summary screen and click **Finish**.



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**.

Origin Node	Source Node	Policy	Operation	Destination Node
Block_Host_E790_92_71_SP	5200_30_10	D2_HUR_HTI_SP	Snapshot	5200_30_10

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.

HITACHI | Ops Center Administrator Dashboard Jobs Monitoring

Dashboard > Storage Systems > 40028 > Volumes > 6

Storage volume 6 (00:00:06)

0

VOLUME LABEL D2_HTI_SVOLS	POOL ID 0	CAPACITY SAVING No	PROVISIONING STATUS UNMANAGED	ASSIGNED TO MIGRATION No
VIRTUAL VOLUME ID -	POOL NAME dr_pool	COMPRESSION ACCELERATI... -	ATTACHED SERVERS -	ALUA ENABLED No
VIRTUAL STORAGE SYSTEM... 40028	TYPE Thin	TOTAL 2.00 TiB	PORT ID CL4-B	T10PI Disabled
VIRTUAL MODEL VSP 5200, 5600	POOL TIER Platinum	BLOCK SIZE 4294967296 Blocks	HOST GROUP NAME/ISCSI T... HDIDProvisione	
	TIERING POLICY 0: All	THIN FREE 2.00 TiB	LUN 2047 (07FF)	
	ATTRIBUTE Thin	THIN USED 0 GiB	NUMBER OF LUN PATHS 0	
	STATUS NORMAL	USAGE 0%		
	PROTECTED			

HITACHI | Ops Center Administrator Dashboard Jobs Monitoring

Dashboard > Storage Systems > 40028 > Volumes > 7

Storage volume 7 (00:00:07)

0

VOLUME LABEL D2_HTI_SVOLS	POOL ID 0	CAPACITY SAVING No	PROVISIONING STATUS UNMANAGED	ASSIGNED TO MIGRATION No
VIRTUAL VOLUME ID -	POOL NAME dr_pool	COMPRESSION ACCELERATI... -	ATTACHED SERVERS -	ALUA ENABLED No
VIRTUAL STORAGE SYSTEM... 40028	TYPE Thin	TOTAL 2.00 TiB	PORT ID CL5-A	T10PI Disabled
VIRTUAL MODEL VSP 5200, 5600	POOL TIER Platinum	BLOCK SIZE 4294967296 Blocks	HOST GROUP NAME/ISCSI T... HDIDProvisione	
	TIERING POLICY 0: All	THIN FREE 2.00 TiB	LUN 2047 (07FF)	
	ATTRIBUTE Thin	THIN USED 0 GiB	NUMBER OF LUN PATHS 0	
	STATUS NORMAL	USAGE 0%		
	PROTECTED			

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.

	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**.

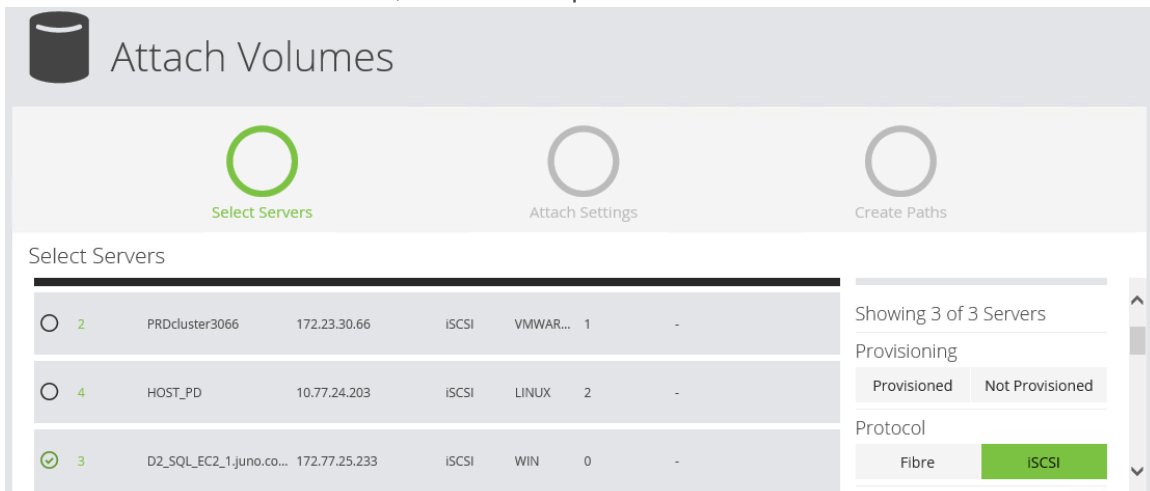
ashboard > Storage Systems > 40028 > Volumes

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...
<input type="radio"/>	0 (00:00:00)	0 (00:00:00)	40028	40028	backup_ta...	0	Thin	Normal	No	5.00 TIB	0%	Attached
<input type="radio"/>	3 (00:00:03)	3 (00:00:03)	40028	40028	JNL_0_D2_...	0	Thin	Normal	No	400.00 GIB	0%	Unattached
<input checked="" type="radio"/>	6 (00:00:06)	6 (00:00:06)	40028	40028	D2_HTI_S...	0	Thin	Normal	No	2.00 TIB	0%	Unmanag...
<input checked="" type="radio"/>	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**.



Attach Volumes

Select Servers

ID	Name	IP	Protocol	Host Mode	Paths
2	PRDcluster3066	172.23.30.66	iSCSI	VMWAR...	1
4	HOST_PD	10.77.24.203	iSCSI	LINUX	2
3	D2_SQL_EC2_1_juno.co...	172.77.25.233	iSCSI	WIN	0

Showing 3 of 3 Servers

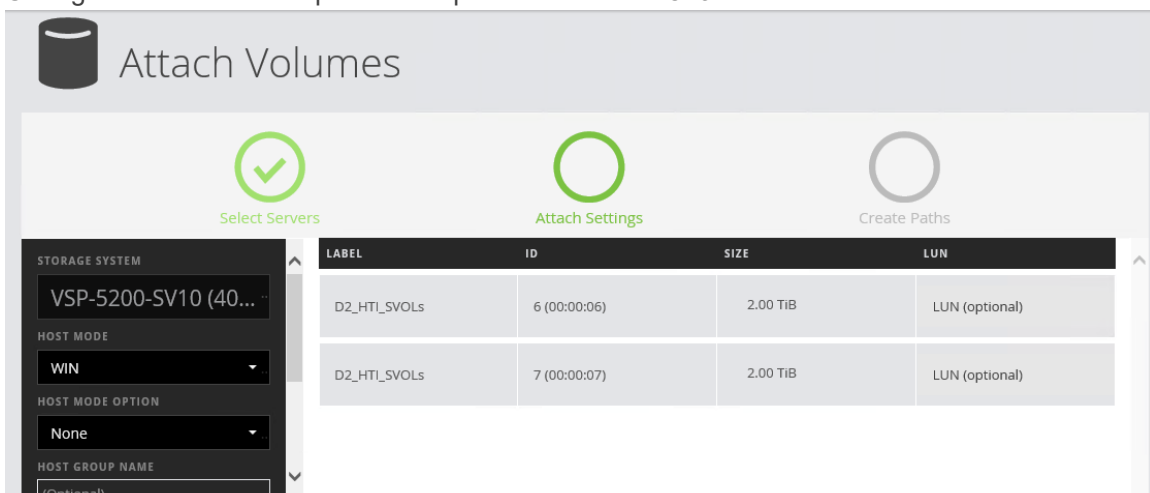
Provisioning

Provisioned Not Provisioned

Protocol

Fibre ISCSI

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



Attach Volumes

Select Servers Attach Settings Create Paths

STORAGE SYSTEM

VSP-5200-SV10 (40...

HOST MODE

WIN

HOST MODE OPTION

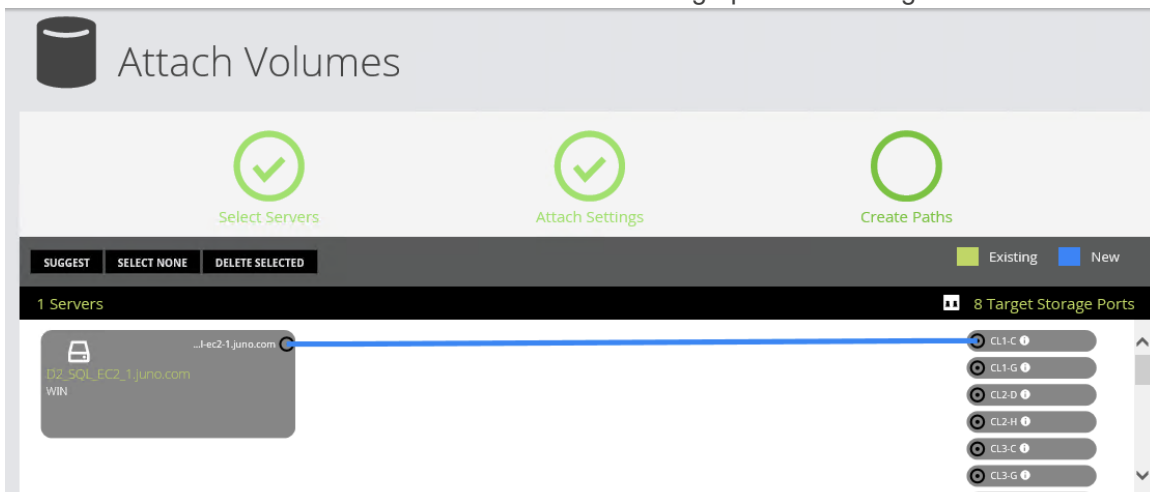
None

HOST GROUP NAME

(Optional)

LABEL	ID	SIZE	LUN
D2_HTI_SVOLs	6 (00:00:06)	2.00 TiB	LUN (optional)
D2_HTI_SVOLs	7 (00:00:07)	2.00 TiB	LUN (optional)

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



Attach Volumes

Select Servers Attach Settings Create Paths

SUGGEST SELECT NONE DELETE SELECTED

Existing New

1 Servers 8 Target Storage Ports

D2_SQL_EC2_1_juno.com WIN

CL1-C

CL1-G

CL2-D

CL2-H

CL3-C

CL3-G

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.

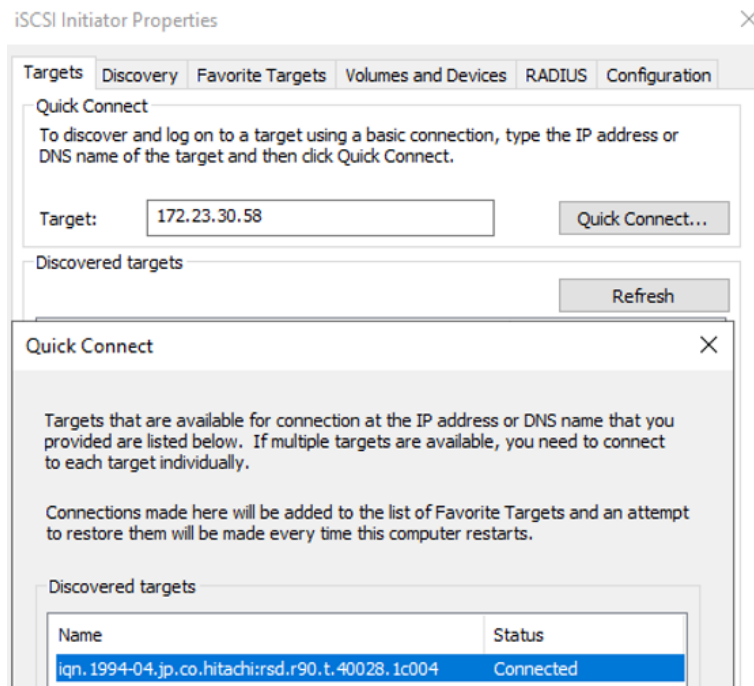
4. Discover the iSCSI target and volumes on the EC2 virtual machine.

- a. Identify the IP address of the storage ports mapped to the EC2 virtual machine.

The screenshot shows the AWS Management Console 'Ports' page for an EC2 instance. A table lists iSCSI targets. The target 'CL1-C' is highlighted, showing its IP address as 172.23.30.58.

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

- 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**.

The screenshot shows the Windows Disk Management console. Two disks are listed, both online and healthy.

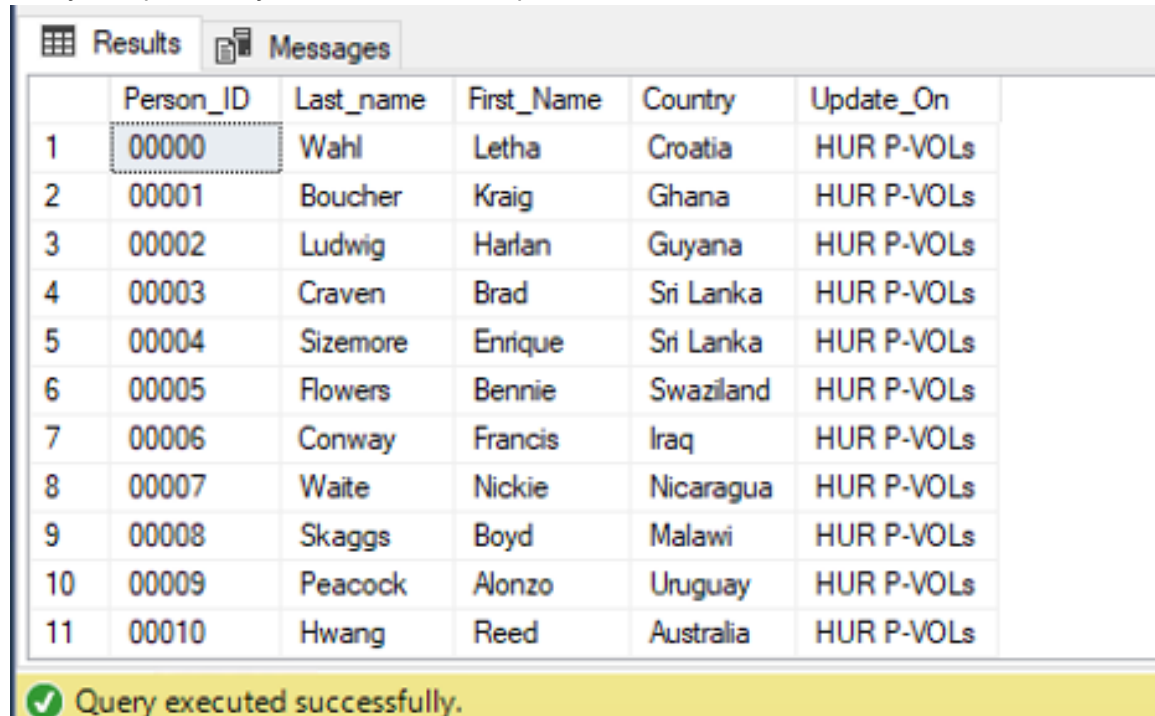
Disk	Volume	Capacity	File System	Health
Disk 1	SQLDBDATA (D:)	2047.98 GB	NTFS	Healthy (Primary Partition)
Disk 2	SQLDBLog (E:)	2047.98 GB	NTFS	Healthy (Primary Partition)

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

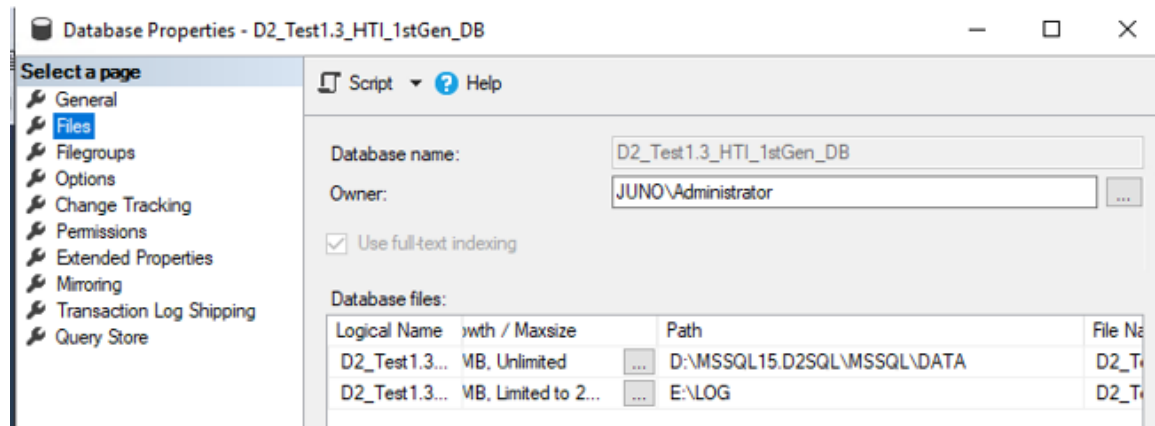
- c. Verify that previously written records are present.



	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.



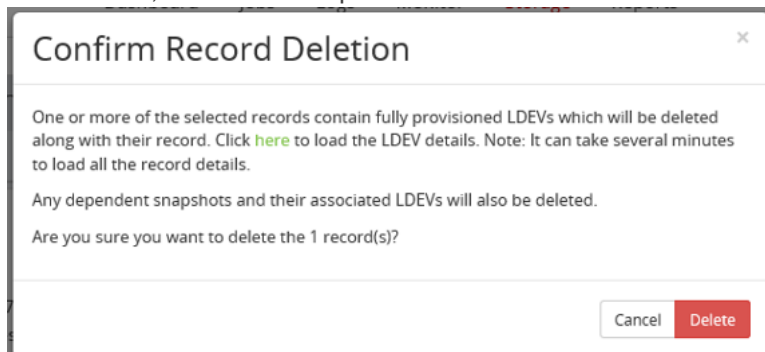
Data is written to the new database.

	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.

	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	FelipaAbemathy@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

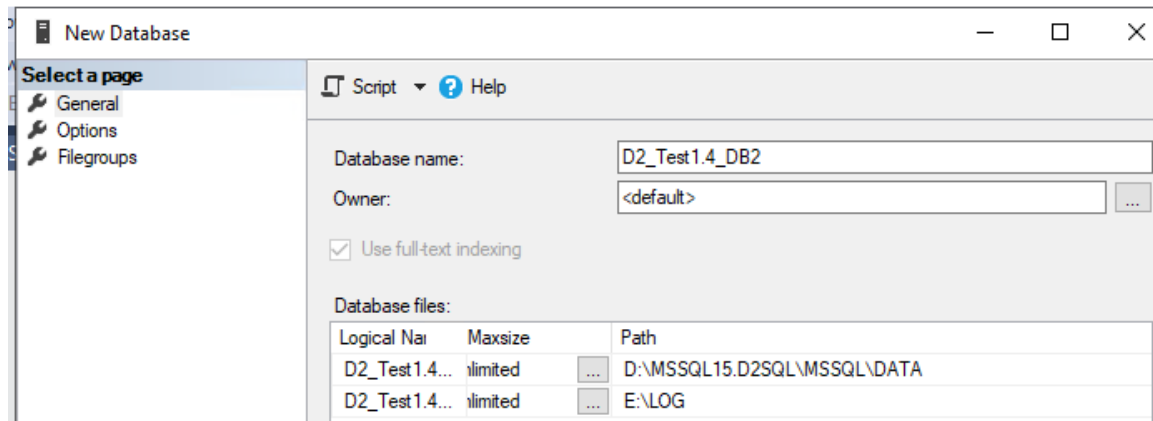
Query executed successfully. D2SQLDB\D2SC

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.

	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	FelipaAbemathy@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

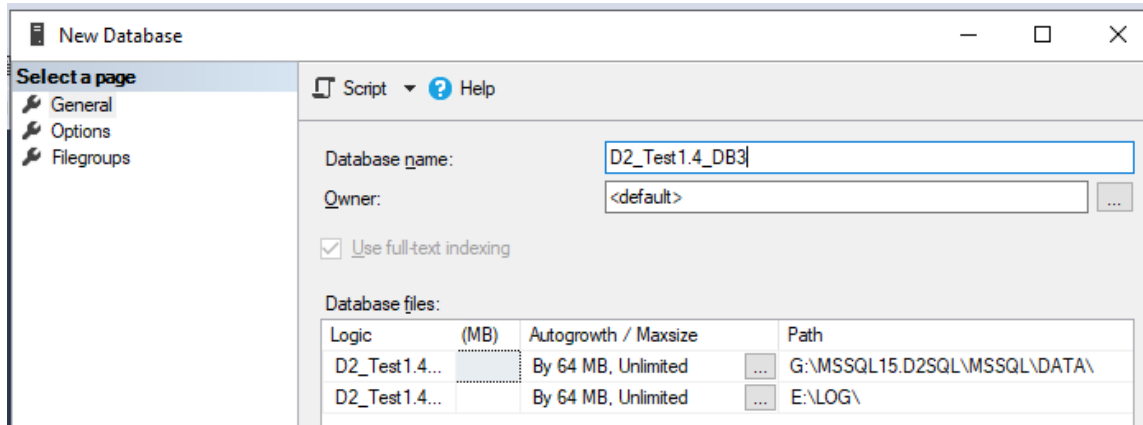
Query executed successfully. D2_SQL_EC2_1\EC2S

- 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.

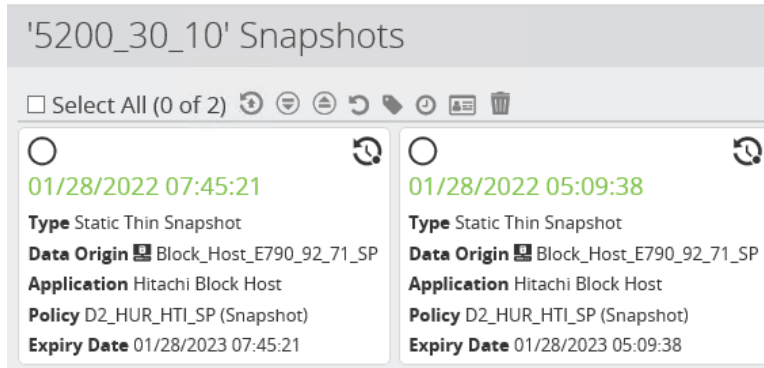


	timestamp	Product_ID	Product_Name	Product_Price	Update_On
1	2022-01-28 06:26:40.0000000	00000	061A27D2Z8K3RE48VT07TROS1	4931	2_HUR-P-VOLs
2	2022-01-28 06:26:40.0000000	00001	S78RQ696155N6LK0G0T46FF0U	83382	2_HUR-P-VOLs
3	2022-01-28 06:26:40.0000000	00002	IQF24115077U8S6Z2F06K39UA	93423	2_HUR-P-VOLs
4	2022-01-28 06:26:40.0000000	00003	C08CC2G28922THYZKMG3EKP1P	3163	2_HUR-P-VOLs
5	2022-01-28 06:26:40.0000000	00004	3BCU5G22NQ6GQBPR16AY8734D	4240	2_HUR-P-VOLs
6	2022-01-28 06:26:40.0000000	00005	U75F2XG65028G5513TIS2C634	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	4PWA659A4EGORIFOTHRI28E3	4709	2_HUR-P-VOLs
11	2022-01-28 06:26:40.0000000	00010	7ZEHFV9S9Y197Y02XRJ3418U2	3961	2_HUR-P-VOLs

Query executed successfully. D2SQLDB\D

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.

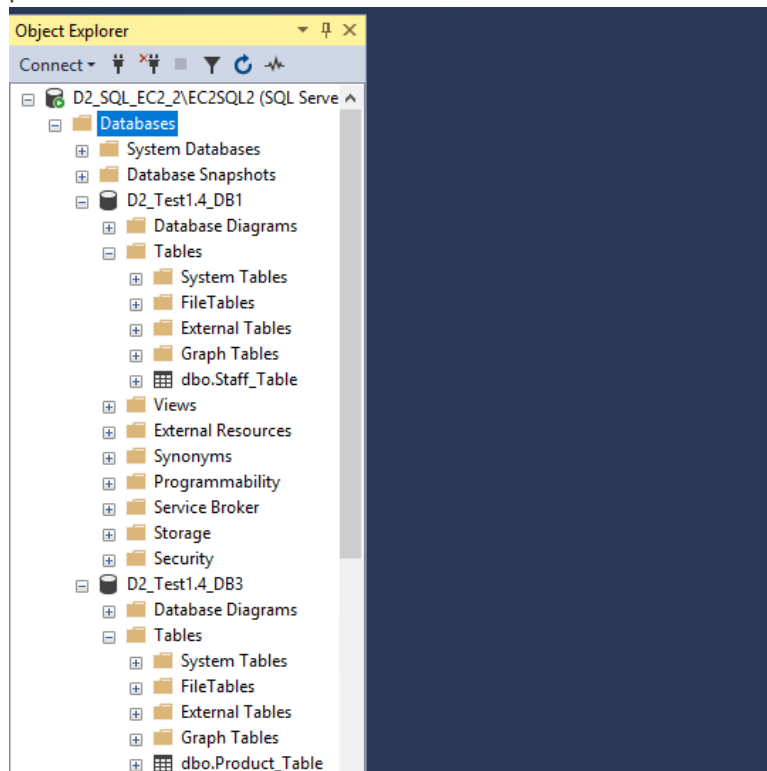


- b. In the second snapshot generation, the HTI primary volumes remain the same, but the HTI secondary volumes are different (from the first snapshot generation HTI secondary volumes).

The screenshot shows a table titled "'5200_30_10' Snapshot '01/28/2022 07:45:21' Pairs". It compares 'Original Primaries' and 'Original Secondaries' across various attributes.

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	-	-	-/-

- Map the second generation snapshots, as described in section [Test 3](#), to a second EC2 virtual machine.
- 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.

	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.Finehart@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	FelipaAbemathy@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

	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	3BCU5G22NQ6GQBPR6AY8734D	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	4PWA659A4EG0RIF0THRGI28E3	4709	2_HUR-P-VOLs
11	2022-01-28 06:26:40.0000000	00010	7ZEHFV9S9Y197Y02XRJ3418U2	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

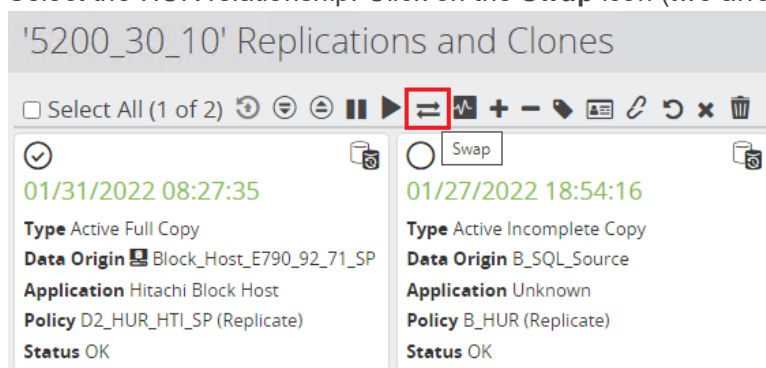
This test case demonstrates how HUR can be suspended in a planned outage; for example, you can perform maintenance in the primary data center and business can continue operations on the secondary storage system. It also shows how HTI snapshots can still be created during the outage. After the outage,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

	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	FelipaAbemathy@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**.

Swap '01/31/2022 08:27:35'

Confirm Swap

Confirm Swap

SWAP

Swapping a replication can potentially cause the **loss or destruction of data**. If you are certain you want to perform this operation enter 'SWAP' in the field above.

Direction

5200_30_10 -> E790_92.71

The final direction of the replication once the swap is complete

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

'5200_30_10' Replications and Clones

Select All (0 of 2) [Icons]

<p>01/31/2022 08:27:35</p> <p>Type Active Full Copy</p> <p>Data Origin Block_Host_E790_92_71_SP</p> <p>Application Hitachi Block Host</p> <p>Policy D2_HUR_HTI_SP (Replicate)</p> <p>Status Swapped</p>	<p>01/27/2022 18:54:16</p> <p>Type Active Incomplete Copy</p> <p>Data Origin B_SQL_Source</p> <p>Application Unknown</p> <p>Policy B_HUR (Replicate)</p> <p>Status OK</p>
--	--

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

'5200_30_10' Replication '01/31/2022 08:27:35' Pairs

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

	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	FelipaAbemathy@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\EC25

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.

New Database

Select a page: General, Options, Filegroups

Script Help

Database name: D2_TEST1.5_DB2

Owner: <default>

Use full-text indexing

Database files:

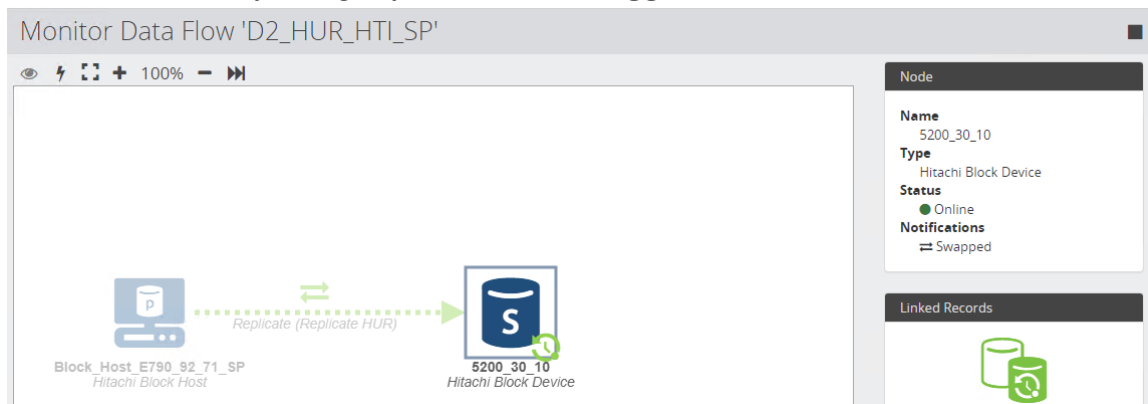
	Path
D2_TEST1....	E:\MSSQL15.D2SQL\MSSQL\DATA
D2_TEST1....	D:\LOG

- 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**.



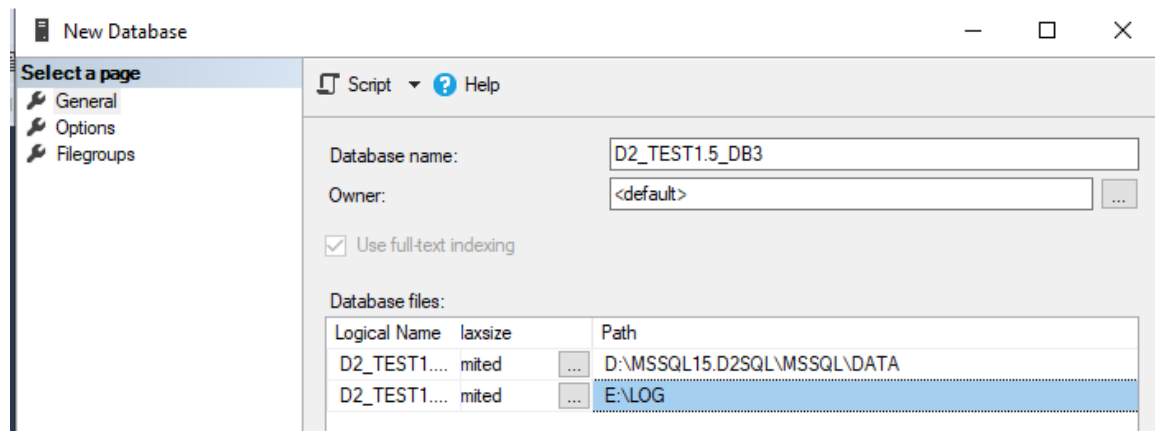
- c. After the snapshots are created, get the HTI snapshot details.

'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	-	-/-	

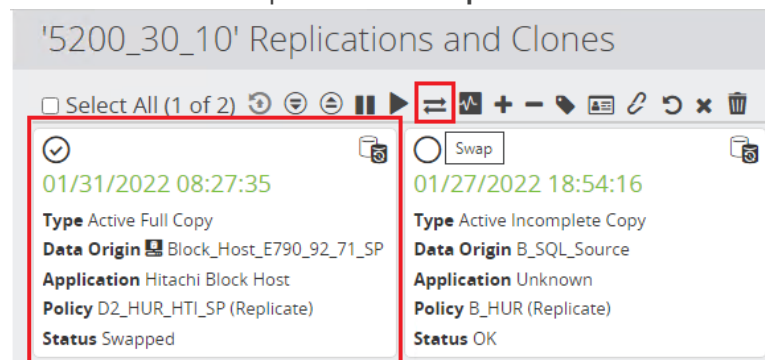
- d. In Administrator, **Refresh** the secondary storage system.
- e. Map the snapshot volumes, as described in section [Test 3](#), to a second EC2 virtual machine.
- f. In the EC2 virtual machine, bring the drives online and attach the databases: D2_Test1.5_DB1 (created while running on HUR primary 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.

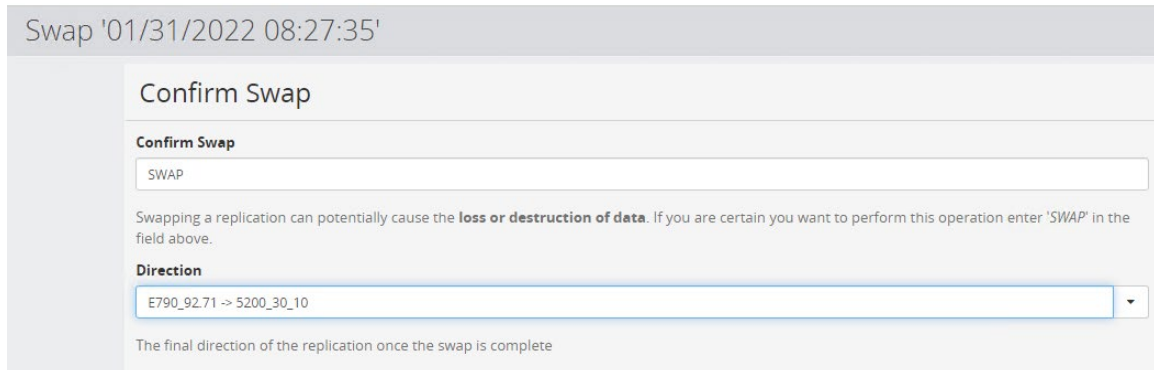


	timestamp	Employee_ID	Employee_Name	Country	Update_On
1	2022-01-31 14:15:44.0000000	80460	Letha Bolt	Croatia	HTI_S-VOLs_1stGen
2	2022-01-31 14:15:44.0000000	19222	Vanita Kelleher	Ghana	HTI_S-VOLs_1stGen
3	2022-01-31 14:15:44.0000000	86849	Marcos Abraham	Guyana	HTI_S-VOLs_1stGen
4	2022-01-31 14:15:44.0000000	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**.

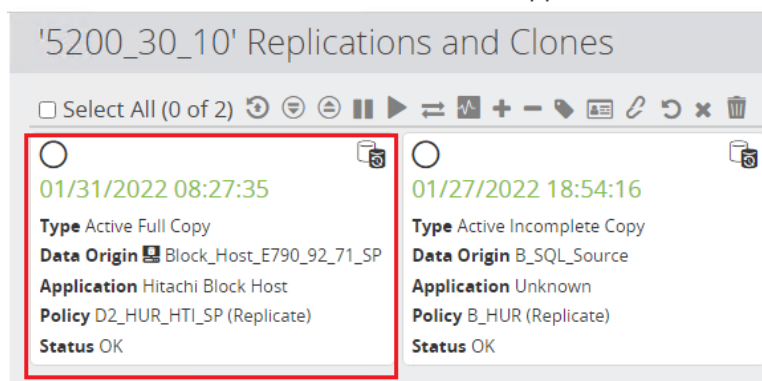


- f. After the swap is done, the original primary volumes are now P-VOLs again, and original secondary volumes are S-VOLs again.

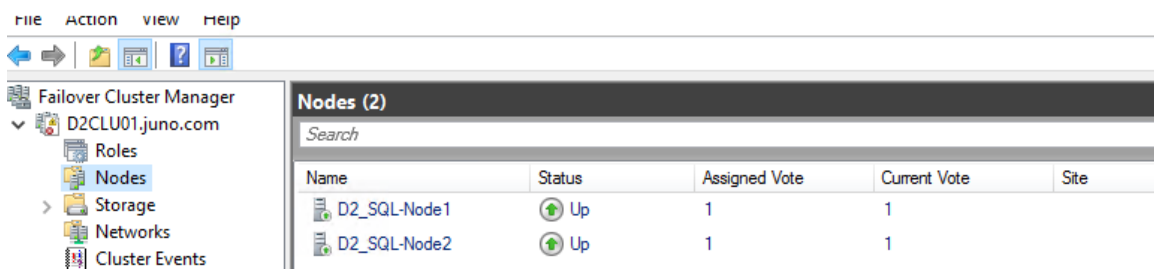
'5200_30_10' Replication '01/31/2022 08:27:35' Pairs

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%	-	-/-


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



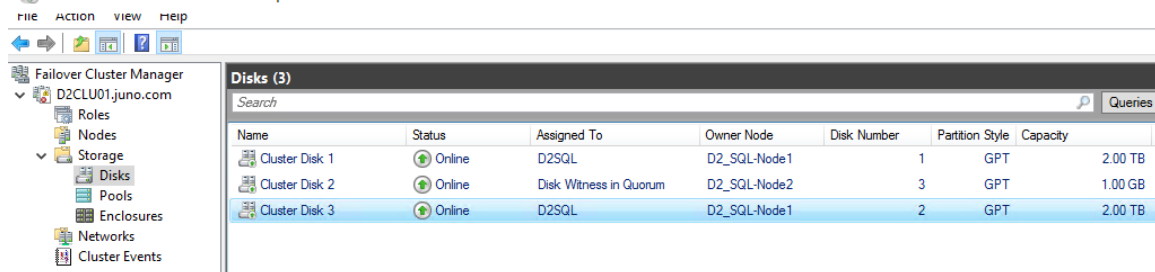
- 8. Power on the SQL virtual machines at the primary site.
 - a. After several minutes, check Windows Failover Cluster Manager to ensure the nodes are online.



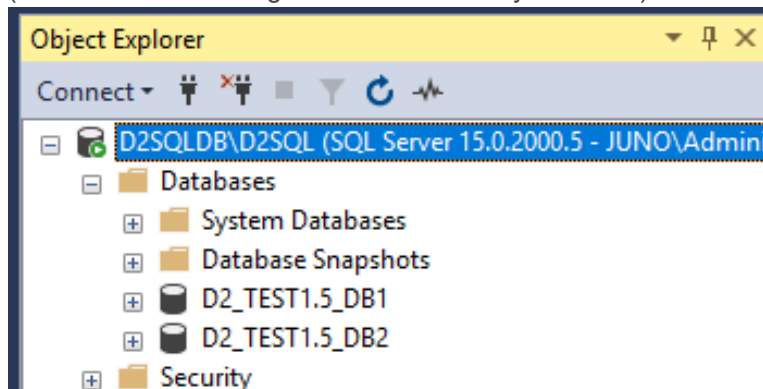
- b. Go to **Roles** and start SQL.



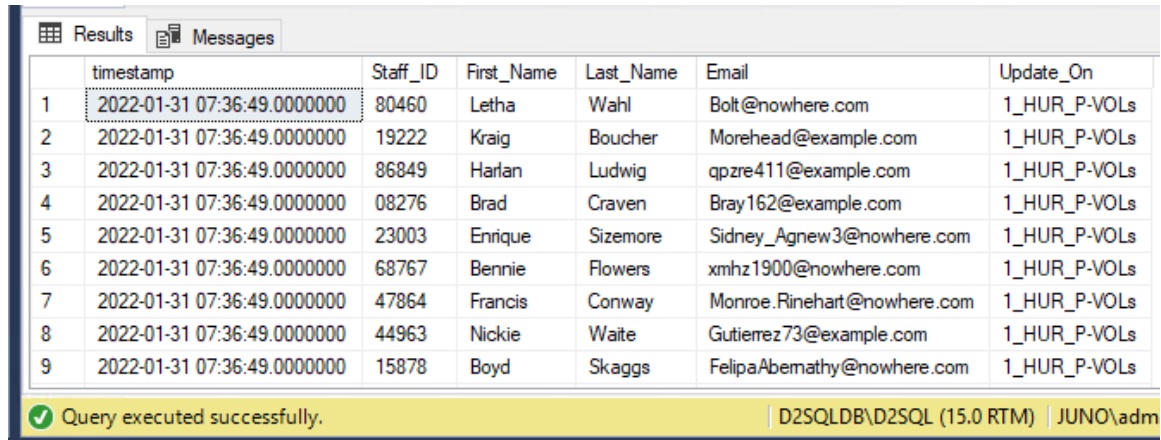
- c. Go to **Storage > Disks** and bring the disks online.



9. Verify that the new records written to SQL at the DR site are present.
- 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).



- Check the content of the databases.



	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.

	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	FelipaAbemathy@nowhere.com	HUR_P-VOLs

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

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.

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

Reference Architecture

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

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore	Operational State
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000b001)	5	disk	2.00 TB	Not Consumed	Dead or Error
<input type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000be0...	3	disk	10.00 GB	Not Consumed	Dead or Error
<input checked="" type="checkbox"/>	HITACHI Fibre Channel Disk (naa.60060e80212b0c0050602b0c0000b00...	10	disk	2.00 TB	D_VMFS	Dead or Error
<input type="checkbox"/>	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:
- Click **Storage**, and then click on the replication target storage system.
 - Click **Replication and Clones**.
 - Click the HUR relationship. The following screen should display:

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

Summary	State
<p>Capture Date 02/06/2022 11:26:56</p> <p>Application Hitachi Block Host</p> <p>Data Flow D2_HUR_HTI</p> <p>Policy D2_HUR_HTI_SP</p> <p>Operation Replicate</p> <p>Mover Type Continuous</p> <p>Type Active Full Copy</p> <p>Pool</p>	<p>Condition OK</p> <p>Paused No</p> <p>Mounted No</p> <p>Exposed to Host No</p> <p>Swapped No</p> <p>Suspend for Swap No</p> <p>Eligible for Teardown No</p>

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

Swap '02/06/2022 11:26:56'

Confirm Swap

Confirm Swap

SWAP

Swapping a replication can potentially cause the **loss or destruction of data**. If you are certain you want to perform this operation enter 'SWAP' in the field above.

Direction

5200_30_10 -> E790_92.71

The final direction of the replication once the swap is complete

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

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

Summary	State
<p>Capture Date 02/06/2022 11:26:56</p> <p>Application Hitachi Block Host</p> <p>Data Flow D2_HUR_HTI</p> <p>Policy D2_HUR_HTI_SP</p> <p>Operation Replicate</p> <p>Mover Type Continuous</p> <p>Type Active Full Copy</p> <p>Pool</p>	<p>Condition OK</p> <p>Paused No</p> <p>Mounted No</p> <p>Exposed to Host No</p> <p>Swapped No</p> <p>Suspend for Swap Yes</p> <p>Eligible for Teardown No</p>

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

02/06/2022 11:26:56

Type Active Full Copy

Data Origin Block_Host_E790_92_71_SP

Application Hitachi Block Host

Policy D2_HUR_HTI_SP (Replicate)

Status OK

'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	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.

	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	FelipaAbemathy@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.

The screenshot shows the 'New Database' dialog box with the following configuration:

- Database name: D2_TEST1.6_DB2
- Owner: <default>
- Use full-text indexing:
- Database files:

File Name	Size	Path
D2_TEST1...	MB, Unlimited	E:\MSSQL15.D2SQL\MSSQL\DATA
D2_TEST1...	MB, Unlimited	D:\LOG

- 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	Conlifedgor	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	Ampputply	Monroe	HUR_S-VOLs
11	2022-02-06 12:45:36.0000000	75351	Amptecticlet	Surfside	HUR_S-VOLs

Query executed successfully. | D2_SQL_EC2_1\

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



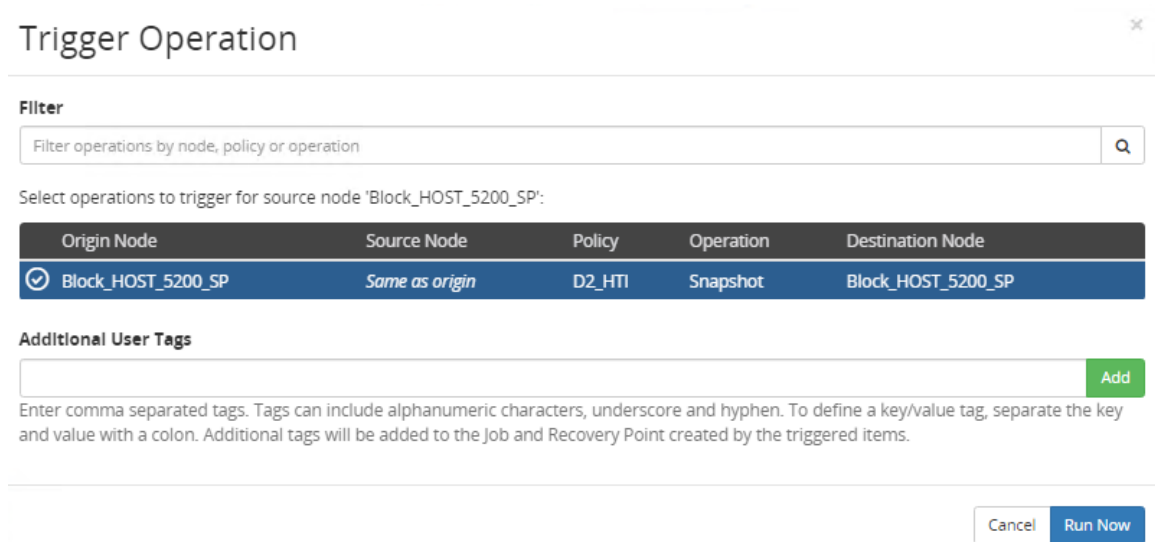
Note: Snapshots cannot be triggered in the original data flow while the primary ISM is unavailable. A new data flow that uses the secondary ISM is needed.

- Create a Block Host specifying the HUR secondary volumes from the VSP 5200 storage system.
- 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**.

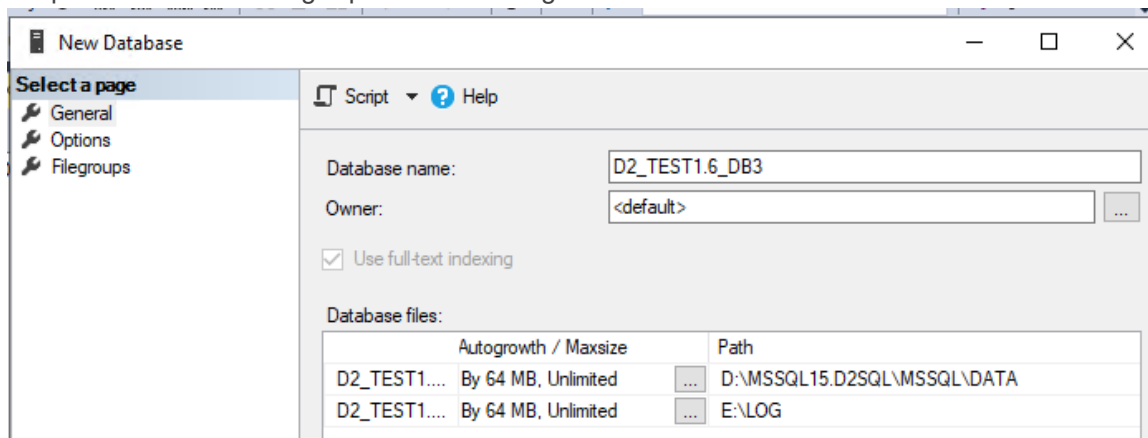


- f. After the snapshots are created, get the HTI snapshot details.



- 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.

	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

- 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	-/-

- a. 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.

Swap '02/06/2022 11:26:56'

Confirm Swap

Confirm Swap

Swapping a replication can potentially cause the **loss or destruction of data**. If you are certain you want to perform this operation enter 'SWAP' in the field above.

Direction

The final direction of the replication once the swap is complete

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

'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	PAIR	S-VOL	0%	-	-/-	← 0xb000	540028	PAIR	P-VOL	0%	-	-/-
✓ 0xb001	611020	PAIR	S-VOL	0%	-	-/-	← 0xb001	540028	PAIR	P-VOL	0%	-	-/-

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**.

'5200_30_10' Replications and Clones

✓ Select All (1 of 1) [Icons]

✓ 02/06/2022 13:35:15

Type Active Full Copy

Data Origin Block_Host_E790_92_71_SP

Application Hitachi Block Host

Policy D2_HUR_HTI_SP (Replicate)

Status Swapped

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

Swap '02/06/2022 13:35:15'

Confirm Swap

Confirm Swap

Swapping a replication can potentially cause the **loss or destruction of data**. If you are certain you want to perform this operation enter 'SWAP' in the field above.

Direction

The final direction of the replication once the swap is complete

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

'5200_30_10' Replication '02/06/2022 13:35:15' 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%	-	-/-

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

'5200_30_10' Replications and Clones

Select All (0 of 1)

02/06/2022 13:35:15

Type Active Full Copy

Data Origin Block_Host_E790_92_71_SP

Application Hitachi Block Host

Policy D2_HUR_HTI_SP (Replicate)

Status OK

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.

Failover Cluster Manager

D2CLU01.juno.com

Roles

Nodes

Storage

Networks

Cluster Events

Nodes (2)

Name	Status	Assigned Vote	Current Vote	Site
D2_SQL-Node1	Up	1	1	
D2_SQL-Node2	Up	1	1	

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

Failover Cluster Manager

D2CLU01.juno.com

Roles

Nodes

Storage

Networks

Cluster Events

Roles (1)

Name	Status	Type	Owner Node	Priority	Information
D2SQL	Running	Other	D2_SQL-Node1	Medium	

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

Failover Cluster Manager

D2CLU01.juno.com

Roles

Nodes

Storage

Disks

Pools

Enclosures

Networks

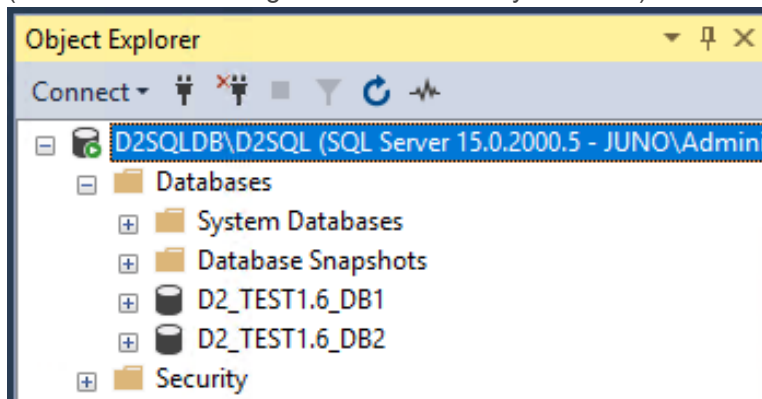
Cluster Events

Disks (3)

Name	Status	Assigned To	Owner Node	Disk Number	Partition Style	Capacity	Rep
Cluster Disk 1	Online	D2SQL	D2_SQL-Node1	1	GPT	2.00 TB	
Cluster Disk 2	Online	Disk Witness in Quorum	D2_SQL-Node2	3	GPT	1.00 GB	
Cluster Disk 3	Online	D2SQL	D2_SQL-Node1	2	GPT	2.00 TB	

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.

	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	FelipaAbemathy@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

	timestamp	Store_ID	Store_Name	City	Update_On
1	2022-02-06 12:45:36.0000000	80460	Conlifedgor	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	Ancycliphone	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	Ampputply	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	Computer	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	Transliector	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	Computer	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	Amlifiry	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	Comcyclor	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 Hemann	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:
- In SSMS, delete the database instance (D2_Test1.7_DB1).
 - In Windows Failover Cluster Manager, stop SQL and shut down the cluster.
 - Power off the virtual machines.
7. Now we will use Protector to restore the primary database with the HTI snapshots.
- The first step is to suspend the HUR pairs and swap their roles.
 - In Protector, click **Storage**, and then click the secondary storage system.
 - Click **Replication and Clones**.
 - Select the HUR relationship and click **Swap**.
 - Type **SWAP** to confirm the task. In the Direction field, select the direction that is opposite of the current direction and click **Finish**.

Swap '02/03/2022 11:07:42'

Confirm Swap

Confirm Swap

SWAP

Swapping a replication can potentially cause the **loss or destruction of data**. If you are certain you want to perform this operation enter 'SWAP' in the field above.

Direction

S200_30_10 -> E790_92.7

The final direction of the replication once the swap is complete

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



02/03/2022 11:07:42

Type Active Full Copy

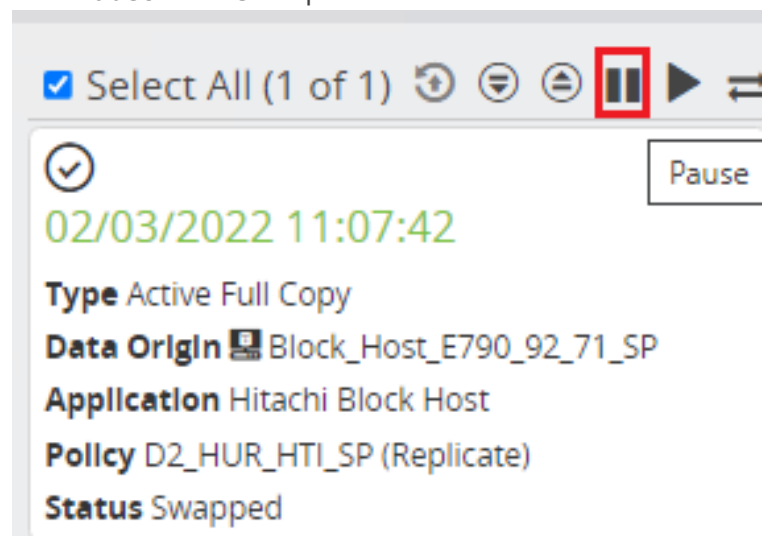
Data Origin Block_Host_E790_92_71_SP

Application Hitachi Block Host

Policy D2_HUR_HTI_SP (Replicate)

Status Swapped

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



Select All (1 of 1) [Refresh] [Down Arrow] [Up Arrow] [Pause] [Play] [Refresh]

02/03/2022 11:07:42

Type Active Full Copy

Data Origin Block_Host_E790_92_71_SP

Application Hitachi Block Host

Policy D2_HUR_HTI_SP (Replicate)

Status Swapped

Pause

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



02/03/2022 11:07:42

Type Active Full Copy

Data Origin Block_Host_E790_92_71_SP

Application Hitachi Block Host

Policy D2_HUR_HTI_SP (Replicate)

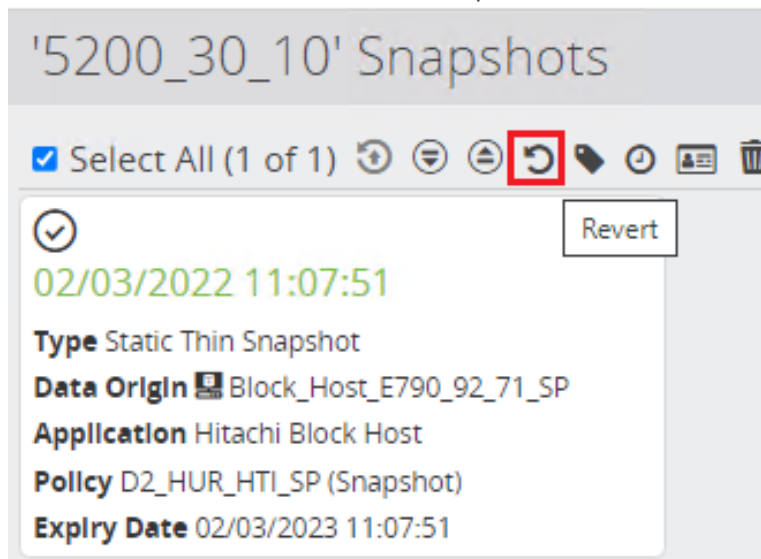
Status Paused, Swapped

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

'5200_30_10' Replication '02/03/2022 11:07:42' Pairs

Original Primaries							Original Secondaries						
ID	Storage	Status	Attribute	%	M	I/O Mode	ID	Storage	Status	Attribute	%	M	I/O Mode
✓ 0xb000	611020	SSUS	S-VOL	100%	-	-/-	← 0xb000	540028	PSUS	P-VOL	100%	W	-/-
✓ 0xb001	611020	SSUS	S-VOL	100%	-	-/-	← 0xb001	540028	PSUS	P-VOL	100%	W	-/-

- 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.

Revert '02/03/2022 11:07:51'

Confirm Revert

Confirm Revert

REVERT

Reverting can potentially cause the **loss or destruction of data**. If you are really certain you want to perform this operation enter 'REVERT' in the field above.

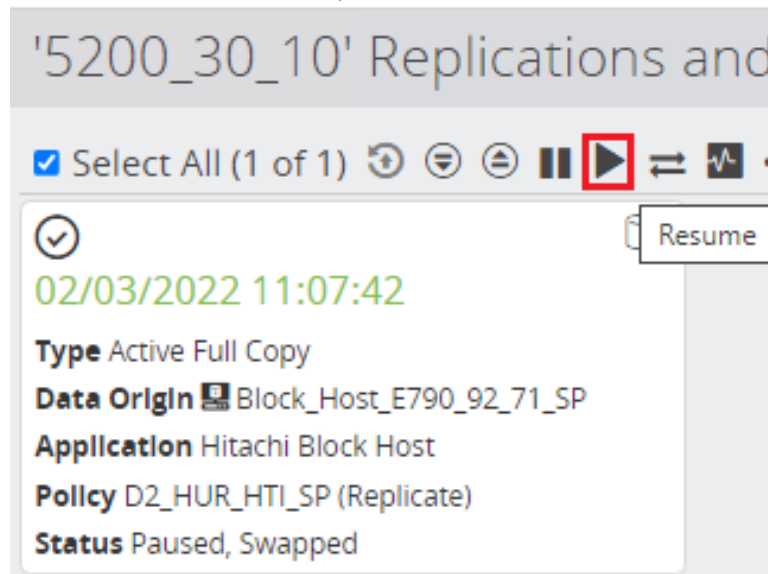
- iii. Confirm the operation completed by verifying that Status shows Completed Successfully.

Job Details

Details	Parameters
<p>Type Other</p> <p>Status ✓ Completed Successfully</p> <p>Node 5200_30_10</p> <p>Operation Revert Block Snapshot</p> <p>Subsystem Block</p> <p>Description Reverting Snapshot</p> <p>Started 02/03/2022 12:28:31</p> <p>Completed</p>	<p>Data Flow D2_HUR_HTI</p> <p>Operation Snapshot</p> <p>Policy D2_HUR_HTI_SP</p> <p>Recovery Point Capture Date 02/03/2022 11:07:51</p> <p>Source 5200_30_10</p> <p>Destination 5200_30_10</p> <p>User sysadmin</p>

- 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**.

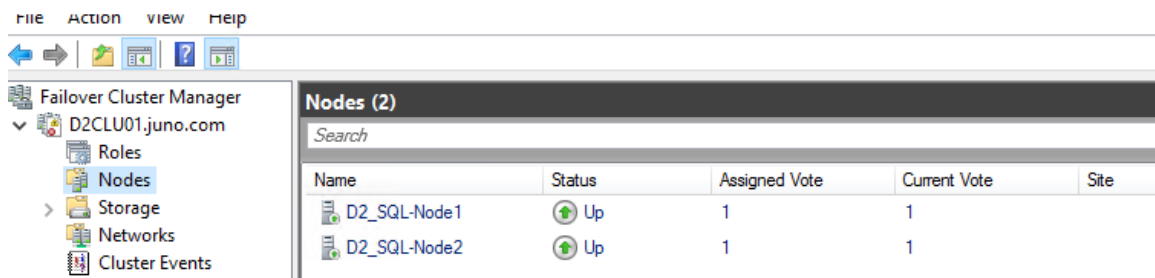


- ii. After the resume operation is done, **Swap** the HUR pairs to the original direction, and after the swap is done, the original primary volumes are now P-VOLs again, and original secondary volumes are S-VOLs again.


'5200_30_10' Replication '02/03/2022 11:07:42' 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%	-	-/-

- 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.

	timestamp	Product_ID	Product_Name	Product_Price	Update_On
1	2022-02-03 06:10:03.0000000	00000	Compuator	859179	HUR_P-VOLs
2	2022-02-03 06:10:03.0000000	00001	Charwoofaquiry	869131	HUR_P-VOLs
3	2022-02-03 06:10:03.0000000	00002	Transliector	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	Compuator	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.

	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

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

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