Building Mobile and Resilient Containerized Applications in a Hybrid Multi-Cloud Environment

Using Hitachi Cloud Connect for Equinix

Hitachi Vantara January 2024

Table of **Contents**

Notices and Disclaimer	2
About This Guide	3
Intended Audience	3
Document Revisions	
References	3
Comments	3
Executive Summary	4
Introduction	5
Solution Overview	6
Benefits	6
Key Components	6
Validation	8
Validation Method	8
High Level Diagram	9
Hardware and Software	10
Test Scenarios	11
Guidelines and Recommendations	13
Validation Results	14
Test 1: Prepare the Environment	14
Test 2: Deploy a Stateful Application in Azure Red Hat OpenShift Cluster	
Test 3: Manually Migrate Stateful Applications Across OpenShift Clusters	29
Test 4: Migrate a Stateful Application Across OpenShift Clusters Using Kasten K10 Multi-Cluster	37
Test 5: Recover from a Ransomware Attack	

Notices and Disclaimer

© 2024 Hitachi Vantara LLC. All rights reserved.

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

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

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

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

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

1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and

2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

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

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

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

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

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

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

About This Guide

This reference architecture documents how to set up backup and restore operations between near-cloud Red Hat® OpenShift cluster and AWS and Azure clusters using Kasten K10 Multi-Cluster Manager and Hitachi Storage Plug-in for Containers (HSPC). Additionally, the document includes test procedures to validate the resiliency of the solution, which you can leverage for your own proof-of-concept before deploying the solution.

Intended Audience

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

Document Revisions

Revision Number Date		Author	Details
v1.0	January 2024	Hitachi Vantara LLC	Initial Release

References

- Azure Red Hat OpenShift v4.11
- Hitachi Storage Administration
- Hitachi Storage Plug-in for Containers Quick Reference Guide v3.12.0
- Red Hat OpenShift Container Platform installation on AWS v4.12
- Red Hat OpenShift Container Platform installation on vSphere v4.12
- Veeam Kasten K10 Guide

Comments

Send any comments on this document to <u>GPSE-Docs-Feedback@hitachivantara.com</u>. 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.

Executive Summary

This reference architecture documents the process of cloud-based migration of a containerized application with the Kubernetes volume snapshot function using Hitachi Storage Plug-in for Containers (HSPC) and Kasten K10 Multi-Cluster Manager by Veeam when a Hitachi Virtual Storage Platform 5200 (VSP 5200) storage system is used as the storage backend. HSPC leverages Thin Image (TI) point-in-time snapshots that are instantaneous and space efficient.

Using MySQL stateful application as an example, this document describes how to use HSPC for backup and restore, disaster recovery, and data mobility. In addition, it includes some real-world use cases. The environment used for this validation includes two Red Hat OpenShift clusters, one at the near-cloud VMware environment, one in Amazon Web Services (AWS), and one Azure-managed Red Hat OpenShift Cluster.

For all clusters, storage is provided from a VSP 5200 storage system located at the near-cloud data center. Keeping the application data in a centralized location has a number of benefits including costs, performance, and security. The near-cloud data center is a colocation operated by Equinix. This solution bridges the cloud divide and ensures availability of data across all clusters.

The Equinix colocation was selected because it offered high-speed and low latency connections to the major hyperscalers, such as AWS and Azure. Hitachi Vantara collaborated with Equinix to offer a near-cloud hybrid offering called **Hitachi Cloud Connect for Equinix**.

This offering allows clients to locate Hitachi products such as the VSP storage systems at Equinix International Business Exchange[™] (IBX) data centers worldwide. In addition, there is an option for clients 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, organizations can locate their data residing on VSP storage systems next to clouds to leverage hybrid- or multi-cloud capabilities while still maintaining physical control of the data.

If you want to discuss hosting these types of solutions at Equinix, contact your Hitachi Vantara sales team. For more information, visit the Hitachi Cloud Connect for Equinix webpage at: <u>https://hitachivantara.com/en-us/products/storage/flash-storage/cloud-connect-for-equinix.html</u>.

Introduction

Red Hat OpenShift is a hybrid-cloud application platform that leverages the power of Kubernetes and combines reliable and proven services to make the process of developing, modernizing, deploying, running, and managing applications more streamlined. OpenShift ensures a uniform user experience whether applications are deployed on public-cloud, on-premises, hybrid-cloud, or edge architecture. Azure Red Hat OpenShift (ARO) delivers on-demand, fully managed OpenShift clusters with high availability, co-managed and operated in partnership with Microsoft and Red Hat.

The installation program of OpenShift Container Platform offers flexibility to deploy on a wide range of platforms. You can deploy OpenShift Container Platform on bare metal, AWS, Azure, GCP, VMware vSphere, and so on.

You can install OpenShift Container Platform using either installer-provisioned (IPI) or user-provisioned infrastructure (UPI) methods. In this reference architecture, Red Hat OpenShift clusters in near-cloud VMware and on AWS were deployed using the IPI method.

Hitachi Storage Plug-in for Containers is a software solution comprising of libraries, settings, and commands that enable you to create a container for running stateful applications. The software enables stateful applications to persist and maintain data after the life cycle of the container has ended. HSPC provides persistent volumes (PV) backed by Hitachi storage systems.

Kasten K10 is an enterprise-grade robust data management platform by Veeam that helps organizations to back up and restore container-based applications on Kubernetes/OpenShift. The capabilities include automating and orchestrating data backup, recovery, disaster recovery, and application mobility across multiple Kubernetes clusters and cloud environments. Kasten K10 offers support for a variety of Kubernetes distributions, as well as public and private cloud providers and storage solutions.

The environment used for this validation includes a Red Hat OpenShift cluster at the near-cloud data center, a Red Hat OpenShift cluster in AWS, and an ARO cluster in Azure. All clusters share the same VSP 5200 storage system located in the near-cloud data center for persistent volume requirement for stateful applications. Keeping the data at the near-cloud location ensures data availability to any cloud vendor at close proximity and avoids cloud locking. The near-cloud data center is a colocation operated by Equinix.

To summarize, our hybrid cloud environment consists of the following three domains. The relationship across the domains is shown in *Figure 1*.

- A near-cloud Equinix colocation data center (named SV5), located in San Jose, California.
- A cloud hosted by AWS in Northern California.
- A cloud hosted by Azure in California.



Figure 1: Hybrid Cloud Environment

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

Solution Overview

HSPC integrates the OpenShift Container Platform with the Hitachi storage system by using the Container Storage Interface (CSI). Integrating backup software Kasten K10 with HSPC enables protection from data loss and on demand application mobility in the OpenShift Container Platform by using the Hitachi storage system functions (such as Thin Image snapshots and ShadowImage clones).

In addition, implementing HSPC enables the availability of high-performance and high-reliability persistent volumes.

Benefits

The following lists the benefits of the business continuity solution using Red Hat OpenShift Cluster, Hitachi storage system, and Kasten K10 Multi-Cluster:

- Allows business to resume operations quickly when a disaster brings down a cluster environment.
- On-demand application mobility: Provides the flexibility to quickly snap data copies in multiple environments for on-demand analytics, data mining, disaster recovery testing, development testing, and similar use cases.
- The backup and restore operations of Kubernetes clusters in a hybrid cloud environment can be centralized with a single pane of glass UI provided by Kasten K10 Multi-Cluster manager.
- Recover from ransomware attacks: Granular, schedule-based snapshots with immutability (using the Data Retention Utility) enables the administrator to recover from a point-in-time snapshot before the attack.
- A substantial reduction in cloud egress costs can be achieved by sharing the same near-cloud storage between AWS and Azure cluster.
- Closes the gap between cloud environments and guarantees data availability across different clusters.

Key Components

The following lists the major components of the solution. For specifications, see the Hardware and Software section.

- Red Hat OpenShift Container Platform: This solution involved two Red Hat OpenShift Clusters. The first cluster consisted of
 three Control Plane nodes and two Worker nodes that were configured in the VMware environment at the near-cloud Equinix
 data center. The second cluster consisted of three Control Plane nodes and two Worker nodes deployed in AWS. Some of
 the key components of Red Hat OpenShift Container Platform are:
 - OpenShift Control Plane node: Runs services required for controlling the OpenShift Container Platform cluster and manages node workloads.
 - OpenShift Worker node: Worker nodes are part of the Kubernetes clusters which are responsible for running the containers and applications. Worker nodes have two main components, the Kubelet Service and the Kube-proxy Service.
- Namespace: Provides the scope of namespace to divide cluster resources among users.
- Persistent Volume and Persistent Volume Claim (PVC): A part of the storage of the cluster that is statically provided by the cluster administrator or dynamically provided by using the "StorageClass" object.
- Azure Red Hat OpenShift: This solution involved one ARO cluster. This cluster consisted of three Control Plane nodes and three Worker nodes. Microsoft ARO is a fully managed Red Hat OpenShift cluster, jointly operated and monitored by Microsoft and Red Hat.
- HSPC: A CSI plugin from Hitachi used to provision persistent volume from Hitachi storage systems to Red Hat OpenShift or Kubernetes cluster to preserve and maintain data after the container life cycle ends.
 - CSI-controller: Mainly incorporates the CSI controller service for storage operation. This service is deployed as "Deployment" and is run only on the control plane.
 - CSI-node: Mainly incorporates the CSI node service that manages volumes in each node. This service is deployed as "DaemonSet". This component is required for all nodes.
- Veeam Kasten K10 Multi-Cluster Manager: Kasten K10 provides a user-friendly data management platform to perform backup or restore, disaster recovery, and mobility of containerized applications. The K10 Multi-Cluster manager provides a platform for K10 operations across multiple OpenShift clusters in a hybrid-cloud environment.
- VSP Storage Systems: A VSP 5200 storage system was used for persistent volume in Red Hat OpenShift clusters deployed in near-cloud, AWS, and Azure for stateful applications.

- Network Switches: Cisco Nexus 9000 Series switch was used to connect to AWS Direct Connect and Azure ExpressRoute. The following accessories are required for establishing a WAN between the near-cloud data center and the clouds:
 - 10/25Gbase-LR-S Optics: Long Range transceivers required to connect long distances.
 - Single-Mode Fiber Cables: Required for long-distance communications.
- Equinix Fabric: Connected equipment at the Equinix near-cloud data center to AWS cloud and Azure cloud.
- AWS Cloud: Equipment at Equinix was connected to AWS cloud using a 10 Gbps Direct Connect link. On AWS, a Virtual Private Cloud was created in the region us-west-1. Some of the key services used in AWS cloud are EC2, S3, Route53, Classic load balancer, and Network load balancer.
- Azure Cloud: Equipment at Equinix was connected to Azure cloud using a 10 Gbps ExpressRoute link. On Azure, a Virtual Network was created in the region West US. Some of the key services used in Azure cloud are ARO, virtual machines, and load balancer.

Validation

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

Validation Method

This solution consists of the following test cases.

Test case 1 involves setting up the environment, which includes three clusters - two Red Hat OpenShift clusters in the nearcloud center and AWS cloud and one Azure RedHat OpenShift cluster.

To validate test case 2, a persistent volume was allocated from the VSP 5200 storage system located in near-cloud to deploy a stateful MySQL application in the ARO cluster.

Test case 3 involves taking backup of the stateful application in the ARO cluster and restoring in the target cluster in AWS. Before performing the backup operation, fresh data was inserted into the MySQL application in the ARO cluster. A snapshot of the persistent volume was created with Kubernetes Volume Snapshot function with HSPC. In the target cluster, a PVC of the snapshot volume was created and used as a source to create a clone volume. The stateful MySQL application was restored using the clone volume in the target cluster in AWS.

To validate test case 4, fresh data was inserted into the MySQL application in ARO cluster, and after restoring the backup, the database records were verified at the AWS location to ensure data consistency. The Kasten K10 Multi-Cluster user interface was used to perform this test case. A Global Location Profile was created with AWS S3 bucket as the storage provider, followed by creating Global Policies to automate the workflows for managing data (such as snapshot and restore). The subsequent step was to add Distributions, which defines the clusters where K10 resources must be allocated. Finally, snapshot and restore operations were carried out using the Global Policies.

Test case 5 shows how business continuity can be performed if a ransomware attack corrupts the application data. To validate this test case, a stateful MySQL application was used and the Hitachi Data Retention Utility (DRU) feature was set on the snapshot volume to restrict read and write. If a ransomware attack corrupts the application data, the data can be restored from the snapshot. You can perform the recovery process in either of the Red Hat OpenShift Container Platform cluster in AWS or in ARO. The process involves creating a PVC of the DRU-enabled snapshot, creating a snap-on-snap copy of that PVC, and then restoring the stateful MySQL application using the cloned PVC in the target cluster.

High Level Diagram

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



Figure 2: Test Environment

Hardware and Software

Table 1 provides the hardware specifications for the equipment used in this validation.

	Item	Description	Version	Function
Equinix Near-Cloud Data Center	Hitachi VSP 5200	1 TB cache (2) 20-core MPUs (4) RAID6 6D+2P parity groups (1) 10 GbE iSCSI port	SVOS RF 9.8.6 90-09-01-00/01	Storage system used to store application data.
	Hitachi Advanced Server DS220	 (2) 18-core Intel Xeon Gold 6140 @ 2.3 GHz 128 GB cache (1) Intel Ethernet Network Adapter XXV710 	BMC 4.70.06 BIOS S5BH3B22.H00	4-node VMware vSphere cluster used to deploy 5-nodes near-cloud Red Hat OpenShift cluster.
	Cisco Nexus C93180YC-FX	(48) 1/10/25-Gbps fiber ports (6) 40/100-Gbps QSFP28 ports	NXOS 9.3(4)	Network switch at the near- cloud data center servicing AWS Direct Connect.
AWS	Amazon EC2	(4) Intel Xeon Platinum 8000 series processor, 16 GB RAM	Instance type: t3.xlarge AMI Name: rhcos- 412.86.202306132230- 0-x86_64 AMI ID: <u>ami-</u> 03260f4b6e0166045	5-nodes Red Hat OpenShift cluster in the cloud.
	Amazon S3	(1) S3 Standard bucket	N/A	Storing Kasten K10 Multi- Cluster Global Location profile.
Azure	Azure Virtual Machine	Master Nodes: (8) Intel Xeon CPU E5-2673 v3 @ 2.40GHz processor, 32 GB RAM Worker Nodes: (4) Intel Xeon CPU E5-2673 v3 @ 2.40GHz processor, 16 GB RAM	Master Nodes: Standard D8s v3 Worker Nodes: Standard D4s v3	6-nodes Azure Red Hat OpenShift cluster.

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 vSphere cluster.
Red Hat OpenShift	4.11.25	Red Hat OpenShift cluster deployed in near-cloud.
	4.12.25	Red Hat OpenShift cluster deployed in AWS.
Azure Red Hat OpenShift	4.11.26	ARO Cluster deployed in Azure.
Hitachi Storage Plug-in for Containers	3.12	HSPC plugin integrates Kubernetes or OpenShift with Hitachi storage systems using Container Storage Interface.
Kasten K10 Multi-Cluster	6.0.2	Kasten K10 Multi-Cluster is a data management platform from Veeam which provides backup operation, disaster recovery, and application mobility for OpenShift applications across multiple clusters.

Item	Version	Function
MySQL	5.7.41	A stateful database application used to validate data consistency.
	T 11 0 0 0	

Table 2: Software Components

Test Scenarios

Table 3 lists the test scenarios performed in the validation.

#	Description	Success Criteria
1	 Prepare the environment: Deploy two Red Hat OpenShift clusters. One in a VMware environment in near- cloud and another in AWS. Deploy one Azure Red Hat OpenShift cluster in Azure. Define storage, network, and iSCSI connections. Use Dynamic Provisioning pool to provision persistent volume for stateful application in Azure and AWS. Deploy HSPC in both clusters. Deploy Kasten K10 and K10 Multi-Cluster in all clusters. Discover ARO and Red Hat OpenShift Cluster in AWS from Kasten K10 Multi- Cluster Manager deployed in near-cloud. 	Environment is set up as per specifications.
2	 Deploy a stateful application in the Red Hat OpenShift Container Platform clusters. This test case is performed in Azure. The persistent volume is provisioned in the ARO cluster from the Hitachi VSP 5200 storage system located in near-cloud. 1. Define the storage class for the VSP 5200 storage system with the required settings. 2. Deploy MySQL database as a stateful application on the ARO cluster with persistent volume claim. 3. Create a new table and ingest new records. 	Persistent volume from the VSP storage system can be provisioned to the ARO cluster. Stateful application can be deployed successfully.
3	 Migrate a stateful application across OpenShift clusters using HSPC (this test case is performed manually instead of Kasten K10): Ingest data into MySQL application in the ARO cluster. Create a Kubernetes volume snapshot. Create PV and PVC of the snapshot volume. Create a clone PVC using the PVC created in step 3 as the source PVC. Use the clone as a volume source to deploy MySQL stateful application in the Red Hat OpenShift Container Platform cluster on AWS. Verify whether the ingested data is visible to the target MySQL environment. 	Verify that the snapshot created in the ARO cluster can be manually restored in the Red Hat OpenShift Container Platform cluster in AWS.
4	 Migrate a stateful application across OpenShift clusters using Kasten K10 Multi-Cluster: Ingest data into MySQL application in Azure. Create an S3 bucket in AWS. Create a global location profile using this bucket. Create a global snapshot policy. Create a global distribution for snapshot policy and add the cluster. Run the snapshot policy for the MySQL application to take the backup. Create a global distribution for restore. Create a global distribution for import policy and add the cluster. Run the policy to restore the application in the target cluster. Verify whether the MySQL application is being restored and the ingested data is visible to the target MySQL environment. 	Verify that the backup taken in the ARO cluster can be restored in the Red Hat OpenShift Container Platform cluster in AWS using Kasten K10 Multi- Cluster.
5	 Recover from a ransomware attack: This test case is performed manually instead of Kasten K10. The Data Retention Utility feature is set on the snapshot volume to protect the backup from any write operations and define the data retention term for the protected volumes. Ingest data into MySQL application in Azure. Create a Kubernetes volume snapshot. 	Revert to clean stateful MySQL application from snapshot data with DRU.

#	Des	scription	Success Criteria
	3. ⊿	Set DRU attribute in the snapshot volume using Command Control Interface.	
	ч.	the data from the snapshot taken in step 2.	
	5.	Create a PVC using the snapshot volume created in step 2.	
	6.	Create a Kubernetes volume snapshot (snap-on-snap) of the PVC created in step 5.	
	This creates a cascaded snapshot volume.		
	7.	Create PVC of the cascaded snapshot (snap-on-snap) volume.	
	8.	Create a clone PVC using the PVC created in step 7 as the source PVC.	
	9.	Use the clone PVC as a volume source to deploy MySQL stateful application in the	
		Red Hat OpenShift Container Platform cluster in AWS.	
	10.	Verify whether the ingested data is visible to the target MySQL environment.	

Table 3: Test Scenarios

Guidelines and Recommendations

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

- While installing a Red Hat OpenShift cluster in a private environment (for example, in an existing Amazon Virtual Private Cloud with a specific AWS Identify and Access Management user), use "*CredentialMode*" to set as "Manual" in the install-config.yaml file. The default mode is "Mint", which assumes that you have administrative privileges.
- While running the OpenShift installation, install-config.yaml file is used by the installer. You must keep a backup of this file. If the installation fails and must be re-run, copy the OpenShift installer and install-config.yaml to a new directory and then run from there. You must not re-use the same directory, or else X.509 certificate error occurs.
- Prepare a separate node outside the cluster for cluster deployment and install OpenShift CLI (oc) command to interact with OpenShift Container Platform for administration.
- While migrating an application using Kasten K10 across clusters, a location profile is mandatory. Without the location profile, import policy would not generate, and restoration is not possible to other clusters. However, to restore an application in the same cluster, a location profile is not required.
- While building a POD with persistent volume, HSPC automatically performs a series of tasks such as provisioning the volume, creating an iSCSI target (or FC host group), attaching the volume to it, discovering the volume on the target node, and then attaching the volume as a block device or creating a file system on it.
- In Kubernetes environment, a "VolumeSnapshot" object cannot be attached to a POD because it is not a persistent volume. To access the snapshot data, create a clone volume and then attach the clone volume to a POD.
- Retention time cannot be reduced while DRU setting is active on a volume.
- ARO does not allow scaling the cluster workers to zero or attempt a cluster shutdown. Deallocating or powering down any virtual machine in the cluster resource group is not supported.

Validation Results

This section shows the steps and screenshots for each test scenario.

Test 1: Prepare the Environment

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

The test environment consists of three clusters: two multi-node Red Hat OpenShift clusters deployed using IPI method in nearcloud VMware environment and in AWS, and one Azure Red Hat cluster deployed in Azure. You must configure the following components for validation of test cases:

- Configure physical LAN and iSCSI connections for OpenShift clusters.
- Establish connections among three clusters.
- Provision DP pool to be used for persistent volume from the VSP 5200 storage system.
- Deploy two Red Hat OpenShift clusters, one in near-cloud VMware environment and another in AWS.
- Deploy one Azure Red Hat cluster in Azure.
- Install HSPC.
- Deploy Kasten K10 Multi-Cluster.

For steps to configure the following components, see <u>https://www.hitachivantara.com/en-us/pdf/architecture-guide/business-continuity-containerized-applications-in-hybrid-cloud-environment.pdf</u>.

- Deploy Red Hat OpenShift clusters in near-cloud VMware environment and AWS.
- Install HSPC.
- Deploy Kasten K10 and Kasten K10 Multi-Cluster.
- Access Kasten K10 dashboard.
- Discover AWS Red Hat OpenShift cluster and ARO cluster from Kasten K10 Multi-Cluster UI in near-cloud as a secondary cluster.

Deploy Azure Red Hat OpenShift Cluster

Azure Red Hat OpenShift is a fully managed Red Hat OpenShift service in Azure.

Prerequisites

Note that the following prerequisites are outside the scope of this document, so we do not describe them in detail. For more information, see <u>https://learn.microsoft.com/en-us/azure/openshift/quickstart-portal</u>.

- Access to the Azure portal.
- Create a service principal.
- Create a resource group.
- · Create a virtual network with two empty subnets.
- Obtain a Red Hat pull secret (optional).
- Install OpenShift CLI (oc) on the admin node to interact with OpenShift Container Platform from a command-line interface.

Create Azure Red Hat OpenShift Cluster

To deploy the ARO cluster from the Azure portal menu, complete the following steps:

1. From the Azure portal menu, search and select Azure Red Hat OpenShift clusters, and click Create.



2. In the **Basics** tab, specify the project and instance details such as region, cluster name, domain name, master VM size, worker VM size, and worker node count.

Home > Azure Red Hat OpenShi	ft clusters >	
Create Azure Red H	lat OpenShift cluster	
Select the subscription to manage d your resources.	eployed resources and costs. Use resource groups like folders to organize and n	nanage all
Subscription * 🕕	hv-gpse	~
Resource group * ①	hitachi-cloud-connec	\sim
Instance details	Create new	
Region * 🛈	(US) West US	~
OpenShift cluster name * 💿	azureopenshift	
Domain name * 🛈	ulloue2swb6226	
	The OpenShift console will be available at https://console-openshift- console.apps.ulloue2swb62	
OpenShift cluster version * 💿	4.11.26	~
Master VM size *	3x Standard_D8s_v3 8 vcpus. 32 GB memory Change size	
Worker VM size *	3x Standard_D4s_v3 4 vcpus. 16 GB memory Change size	
Worker node count * 🛈	0	3

3. In the Authentication tab, specify the service principal client ID, service principal client secret, and Red Hat pull secret.

Home > Azure Red Hat OpenShift clusters >					
Create Azure Red Hat	Create Azure Red Hat OpenShift cluster				
Basics Authentication Network	ing Tags Review + create				
Service principal information					
Service principal type	○ Create new				
	• Existing				
Service principal client ID * 🛈					
Service principal client secret * ①					
Pull secret					
Red Hat pull secret 🛈					

4. In the **Networking** tab, specify the virtual network name, two empty subnets (one for the control plane and one for worker nodes), and select the network settings.

\equiv Microsoft Azure	\wp Search resources, services, and docs (G+/)					
Home > Azure Red Hat OpenShift clusters > Create Azure Red Hat OpenShift cluster						
Basics Authentication Networ	Basics Authentication Networking Tags Review + create					
Cluster network						
Virtual network * ①	(New) aro-vnet-llovv32n					
Master subnet *	(New) master-subnet (10.0.0/27)					
Worker subnet *	(New) worker-subnet (10.0.0.128/25)					
Pod CIDR 10.128.0.0/14 10.128.0.0 - 10.131.255.255 (262144 addresses)						
Service CIDR 172.30.0.0/16 172.30.0.0 - 172.30.255.255 (65536 addresses)						
Network settings						
API server visibility * 🛈						
Private						
Ingress visibility * 🛈	Public					
O Private						

5. In the Tags tab, add tags to organize resources.

6. In the **Review + create** tab, click **Create** when the validation completes.

= M	licrosoft Azure				$ \mathcal{P} $ Search resources, services, and docs (G+/)		
Home >	Home > Azure Red Hat OpenShift clusters >						
Creat	e Azure Re	d Hat Ope	enShi	ft cluster			
		-					
🔗 Valio	lation passed						
Basics	Authentication	Networking	Tags	Review + create			

7. When the deployment completes, navigate to the cluster.

Home >			
G azureopenshift ☆ Azure Red Hat OpenShift cluster			×
₽ Search «	📋 Delete 🖒 Refresh 🖉 Connect		
Overview			JSON View
Activity log	Resource group : hitachi-cloud-connect	Provisioning state : Succeeded	
^A ဝ Access control (IAM)	Location : West US	OpenShift console : https://console-openshift-console.apps.qlkhtgj85f8cb5b24;	
🗳 Tags	Subscription : <u>hv-gpse</u>	OpenShift version : 4.11.26	
Settings	Subscription ID : 8f5b067e-c9da-42ea-b37	API server visibility : Private	
Settings		API server IIRI + https://api.olkhtpi85f8ch5h2/9+	

8. When the installation completes, you can access the console URL:

Overview	
Cluster	
Getting started resources ②	
Set up your cluster	Suild with guided documentation
Finish setting up your cluster with recommended configurations.	Follow guided documentation to build applications and familiarize yourself with key features.
Add identity providers →	Monitor your sample application →
Configure alert receivers →	Get started with Quarkus using a Helm Chart $ ightarrow$
View all steps in documentation p	View all quick starts
Details View settings	Status
Cluster API address https://api.qlkhtgj85f8cb5b249.westus.aroapp.io:6443	Cluster Control Plane Operators A Insights Disabled
Cluster ID 9322ed62-b1c3-4fld-9e31-50832b3f0e54	Aug 21, 2023, 4:12 PM
Infrastructure provider Azure	Insights operator is disabled. In order to enable Insights and benefit from recommendations specific to your clus documentation: https://docs.openshift.com/container-platform/latest/support/remote_health_monitoring/enable.
OpenShift version 4.11.26	Jul 25, 2023, 3:02 PM Cluster storage operator monitors all storage classes configured in the cluster

Status of the master and worker nodes of the cluster:

Nodes								
▼ Filter ▼ Name ▼	Search by name							
Name 1	Status 1	Role 1	Pods 1	Memory 1	CPU 1	Filesystem I	Created 1	Instance type
🚷 azureopenshift-xxzww- master-0	Ready	master	38	10.53 GIB / 31.39 GIB	1.082 cores / 8 cores	55 GIB / 1,023.8 GIB	Jul 25, 2023, 10:33 AM	Standard_D8s_v3
O azureopenshift-xxzww- master-1	Ready	master	65	13.38 GiB / 31.39 GiB	1.525 cores / 8 cores	23.17 GiB / 1,023.8 GiB	🚱 Jul 25, 2023, 10:32 AM	Standard_D8s_v3
azureopenshift-xxzww- master-2	Ready	master	49	10.92 GiB / 31.39 GiB	1.040 cores / 8 cores	114.7 GiB / 1,023.8 GiB	Jul 25, 2023, 10:33 AM	Standard_D8s_v3
& azureopenshift-xxzww- worker-westus-9q5rl	Ready	worker	36	5.96 GIB / 15.64 GIB	0.914 cores / 4 cores	38.53 GIB / 127.8 GIB	Jul 25, 2023, 10:48 AM	Standard_D4s_v3
azureopenshift-xxzww- worker-westus-65fg4	Ready	worker	28	5.34 GiB / 15.64 GiB	0.845 cores / 4 cores	20.76 GIB / 127.8 GIB	Jul 25, 2023, 10:50 AM	Standard_D4s_v3
& azureopenshift-xxzww- worker-westus-cg27c	Ready	worker	25	3 GiB / 15.64 GiB	0.442 cores / 4 cores	11.71 GiB / 127.8 GiB	Aug 21, 2023, 4:17 PM	Standard_D4s_v3

Install Hitachi Storage Plug-in for Containers

Installing HSPC in the Azure Red Hat OpenShift Cluster is similar to the section Install HSPC in Near-Cloud Cluster.

The following screenshot shows the status of the operator after a successful installation.

Project: k	ube-system 🔻					
Installed Operators						
Installed O	perators are represented by ClusterServiceVersions with	hin this Namespace. For more information, see the Understanding C	perators documentation g. Or create an Operator and ClusterServic	ceVersion using the Operator SDK 🗗.		
Name	Search by name.					
Name	1	Managed Namespaces	Status	Last updated	Provided APIs	
HINCH	Hitachi Storage Plug-in for Containers 112.0 provided by Hitachi	(KS) kube-system	Succeeded Up to date	Aug 22, 2023, 8:14 PM	HSPC	I

From the console, navigate to Workloads, click Pods, and ensure that the status of the operator pod is running.

Project kube-system 💌								
Pods								Create Pod
Filter Name Search by name								
Name 1	Status I	Ready I	Restarts I	Owner I	Memory I	CPU I	Created 1	
Phspc-operator-controller-manager-66f5f6858 fgc5g	- 3 Running	1/1	1	B hspc-operator-controller-manager-66f5f6858	37.0 MiB	0.002 cores	🕲 Jul 25, 2023, 12:31 PM	E

Verify that the status is Ready.

oc get hspc -n kube-system
NAME READY AGE
hspc true 6m02s

Create StorageClass and Volume SnapshotClass

After installing HSPC, you must create storage class to provision persistent volume from the VSP 5200 storage system. Additionally, a volume snapshot class is required to take point in time snapshot. The following screenshots show the status of storage class and volume snapshot class.

The YAML file used for storage class: StorageClasses > StorageClass details sc-vsp5200 sc-vsp5200 Details YAML veClass kind: Stora 1 apiVersion: storage.k8s.io/v1 2 metadata: name: sc-vsp5200 4 uid: 6f2e9fe5-5206-4662-a1de-7edd8320b737 resourceVersion: '133926' 6 creationTimestamp: '2023-07-25T09:28:52Z' 8 annotations: kubernetes.io/description: Hitachi Storage Plug-in for Containers 9 storageclass.kubernetes.io/is-default-class: 'true 10 11 > managedFields: provisioner: hspc.csi.hitachi.com 52 parameters: 53 csi.storage.k8s.io/fstype: ext4 54 csi.storage.k8s.io/provisioner-secret-namespace: default 55 56 csi.storage.k8s.io/provisioner-secret-name: secret-vs csi.storage.k8s.io/node-stage-secret-name: secret-vsp 57 csi.storage.k8s.io/controller-expand-secret-name: secret-vsp5200 58 csi.storage.k8s.io/node-publish-secret-namespace: default 59 csi.storage.k8s.io/controller-publish-secret-name: secret-vsp5200 60 csi.storage.k8s.io/controller-publish-secret-namespace: default 61 poolID: '0' 62 csi.storage.k8s.io/node-publish-secret-name: secret-vsp5200 63 connectionType: iscsi 64 65 csi.storage.k8s.io/controller-expand-secret-namespace: default portID: CL1-C 66 serialNumber: '40028' 67 csi.storage.k8s.io/node-stage-secret-namespace: default 68 reclaimPolicy: Delete 69 70 allowVolumeExpansion: true 71 volumeBindingMode: Immediate

Status of the StorageClasses:

StorageClasses		1	Create StorageClass
Name Search by name. 7			
Name 1	Provisioner 1	Reclaim policy 1	
SS sc-vsp5200 – Default	hspc.csi.hitachi.com	Delete	1

VolumeSnapshotClasses > VolumeSnapshotClass details vsc) snapshotclass-sample Events YAML Details apiVersion: snapshot.storage.k8s.io/v1 1 2 deletionPolicy: Delete 3 driver: hspc.csi.hitachi.com kind: VolumeSnapshotClass 4 metadata: annotations: 6 k10.kasten.io/is-snapshot-class: 'true' 7 8 creationTimestamp: '2023-07-25T09:30:04Z 9 generation: 1 10 > managedFields: -name: snapshotclass-sample 34 resourceVersion: '133372' 35 uid: 745d9cec-a779-48db-adb7-db49dc80266c 36 37 parameters: 38 csi.storage.k8s.io/snapshotter-secret-name: secret-vsp5200 39 csi.storage.k8s.io/snapshotter-secret-namespace: default 40 poolID: '0'

The YAML file for volume snapshot class:

Status of the VolumeSnapshotClasses:

VolumeSnapshotClasses			reate VolumeSnapshotClass
Name V Search by name. 7			
Name 👃	Driver 💲	Deletion policy 1	
VISS snapshotclass-sample	hspc.csi.hitachi.com	Delete	I

Install Kasten K10 in Azure Red Hat OpenShift Cluster

Deploying Kasten K10 in the ARO cluster is similar to deploying it in the near-cloud OpenShift cluster. The following screenshots show the status of Kasten K10.

Status of Kasten K10 Operator:

Project: ka	asten-io 🔻					
Installe	Installed Operators					
Installed Op	perators are represented by ClusterServiceVersions with	hin this Namespace. For more information, see the Understanding O	perators documentation 🖉. Or create an Operator and ClusterServi	ceVersion using the Operator SDK 🖉		
Name 🚽	Search by name/					
Name	1	Managed Namespaces	Status	Last updated	Provided APIs	
KIO	Kasten K10 (Enterprise – Term) 6.0.5 provided by Kasten by Veeam, kasten.io	NS kasten-lo	Succeeded Up to date	Aug 22, 2023, 10:20 PM	K10restore K10	

Status of Kasten K10:

Project: kasten-io 🔻					
Installed Operators > Operator details warren Kotten KIO (Enterprise - Term) 6.0.5 provided by Kasten by Veeam, kasten lo					Actions 👻
Details YAML Subscription Events All instance	nces K10restore K10				
K10s					Create K10
Name - Search by name /					
Name I Kin	ind I	Status I	Labels I	Last updated	
🔇 ki0 ki0	10	Conditions: Initialized, Deployed	No labels	3 Jul 25, 2023, 3:55 PM	I

Prepare Kasten K10 Multi-Cluster in Azure Red Hat OpenShift Cluster

Download and untar the Kasten K10 Multi-Cluster tool from the URL: https://github.com/kastenhq/external-tools/releases.

Discover Azure RedHat OpenShift Cluster as the secondary cluster from Kasten K10 Multi-Cluster UI in near-cloud.

Status of the discovered cluster "azurecluster" in Kasten K10 dashboard in near-cloud:

→ C ▲ Not secure k10route1-kasten-io.apps.ocpcluster.juno.com/k10/#/clusters	Q &
KASTEN By Vecam	🗓 Docs 🖄 kube:admin 🗸 🗘 2
K10 Multi-Cluster Dashboard	K10 Global Resources RBAC Entries Create and distribute policies and profiles to multiple clusters and manage multi-cluster permissions. Global Policies Global Policies Distributions
 3 Clusters 13 Policies 213 Applications O Non-Compliant Applications 	Data Usage Total Across Clusters Snapshot Object 883.0 GiB 883.0 GiB 30.7 MiB Recent Activity Image: Colspan="2">Actions Colspan="2">Failed Actions
Clusters	
L≅ A-Z Filter by Name 3 clusters	Completed with Errors 🚺 🕀 Add Clusters
CLUSTER	APPLICATIONS POLICIES ACTIONS + 1D
image: magnetic state and the last of the l	70 0 0 70 3 0 0 v
image: state with the state of the	
primury Ocpcluster dist: No Juster 1 Job Luster-Type primary	70 0 71 3 0 0 0 .

Test 2: Deploy a Stateful Application in Azure Red Hat OpenShift Cluster

This test case describes the process of deploying MySQL stateful application in Azure Red Hat OpenShift cluster using persistent volume from the near-cloud VSP 5200 storage system. HSPC enables the application to use a persistent volume from the VSP 5200 storage system.

- 1. Deploy a stateful MySQL application.
 - a. Create a project (Kubernetes namspace) for the MySQL application. From the Red Hat OpenShift console, navigate to Home, click Projects, and then click Create Project. In the Create Project menu, enter a project name and click Create.

Create Project
An OpenShift project is an alternative representation of a Kubernetes namespace.
Learn more about working with projects 🗗
Name * 💿
devapps
Display name
Description
Cancel Create

Status of the project.

Projects						Create Project
▼ Filter ▼ Name ▼ deva						
Name deva 🗙 Clear all filters						
Name 1 Display name 1	Status 💲	Requester 💲	Memory 1	CPU 1	Created 1	
PR devapps No display name	⊘ Active	kube:admin	-	-	Aug 16, 2023, 9:45 AM	:

b. Create a MySQL service. From the Red Hat OpenShift console, navigate to **Networking**, click **Services**, and then click **Create Service**. In the Create Service menu, populate the YAML file with the required information and click **Create**.

Creat	te Service				
Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.					
1	apiVersion: v1				
2	kind: Service				
3	metadata:				
4	namespace: devapps				
5	name: azapps2				
6	labels:				
7	app: azapps2				
8	spec:				
9	ports:				
10	- port: 3306				
11	name: azapps2				
12	clusterIP: None				
13	selector:				
14	app: azapps2				

c. Verify the status of the MySQL service:

Project: devapps 🔻				
Services				Create Service
Name Search by name				
Name 1	Labels 💲	Pod selector 1	Location 1	
S azapps2	app*azapps2	Q app=azapps2	None	:

d. Create a MySQL statefulset application. From the Red Hat OpenShift console, navigate to **Workloads**, click **StatefulSet**, and then click **Create StatefulSet**. In the Create StatefulSet menu, populate the YAML file with the required information and click **Create**.

Crea	te StatefulSet
Create b	y manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.
1	apiVersion: apps/v1
2	kind: StatefulSet
3	metadata:
4	namespace: devapps
5	name: azappsz
6	spec:
7	selector:
8	
10	app: d2dppsz
10	podWapagementBolicy, Bapallol
11	polyanagementrolity; Parallel
12	template.
14	metadata
15	labels
16	ann: azanns?
17	snec:
18	terminationGracePeriodSeconds: 30
19	containers:
20	- name: azaphs2
21	image: mysgl:5.7
22	args:
23	- "ienore-db-dir=lost+found"
24	env:
25	- name: MYSOL ROOT PASSWORD
26	value: pass123
27	- name: MYSQL DATABASE
28	value: devmysqldb1
29	- name: MYSQL_USER
30	value: admin
31	- name: MYSQL_PASSWORD
32	value: secret
33	ports:
34	- containerPort: 3306
35	name: mysql
36	volumeMounts:
37	- name: dev-vol1
38	mountPath: /var/lib/mysql
39	volumeClaimTemplates:
40	- metadata:
41	name: dev-vol1
42	spec:
43	storageClassName: sc-vsp5200
44	accessModes: ["ReadWriteOnce"]
45	resources:
46	
4/	Storage: 20061

e. Verify whether the StatefulSet is running.

StatefulSets > StatefulSet details								Actions •
Details Metrics YAML	Pods Environment Events							
▼ Filter ▼ Name ▼ Sean	ch by name	Ready	Restarts 1	Node I	Memory 1	CPU I	Created 1	
@ azapos2+0	C Bunning	1/1	0	azureopenshift-xxzww-worker-westus-9d5rl	202.2 MIB		3 Aug 16, 2023, 9:52 AM	1

f. Verify whether the PVC is created from the VSP 5200 storage system. Using storage class dynamically provisions a persistent volume in the VSP 5200 storage system. The following screenshots show the status of the PVC and PV created.

	Project devapor	•							
	. roject octopps								
<pre>view view view view view view view view</pre>	PersistentVo	olumeClaims							Create PersistentVolumeClaim
<pre>Net of the second second</pre>	T Cilere en	News - Creatible and							
Numeric Bate 1 Peterbeller Gate 2 Weij Result	T riter •	Name • Search by name /							
Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park Image: Constrained Park	Name 1	Sta	tus I	PersistentVolumes 1	Capacity		Used 1	StorageClass I	
The construction of the second	PVC dev-vol1-a	izapps2-0	Bound	pvc-0ba60c3e-0a8d-468d-bb64- 8a25e6dcbc72	200 GiB		-	SC sc-vsp5200	1
<pre>pro-0-basic - dots-dots-dots-dots-dots-dots-dots-dots-</pre>				002000000					
<pre>cp cc c</pre>	PersistentVolumes 1	Persistent\/okme.details							
Image: State Price: State Image: State	😰 pvc-Ob	a60c3e-0a8d-468d-b	b64-8a25e6dcbc72 💿 Boun	0					Actions
<pre>tem: tem</pre>	Deteile MAM	41							
Productival default in the second of the	Details YAM	1L							
me me <td>PersistentVol</td> <td>lume details</td> <td></td> <td></td> <td></td> <td></td> <td></td> <td></td> <td></td>	PersistentVol	lume details							
and control of the second of t	Name					Status Bound			
india tundia	pvc-ubabucse-uad	50-4060-0004-5825e0000072				Capacity			
And and a set of the s	Labels				Edit 🖉	200Gi			
And and a second	i to tabeta					Access modes ReadWriteOnce			
<pre>tetra control to the second of the seco</pre>	Annotations 3 annotations a					Velume mode			
Date Details Organization Details Optimization Details VAML Details Understandards Details VAML Details Image:	Reclaim policy					Filesystem			
Control Control Presentation Control <	Delete					StorageClass			
Very Lastes/during * Persistent/during status Protected/during * Persistent/during Protected/during * Per	Created at					SC-Vsp5200			
Presteriet Values } Presteriet Values details Presteriet Values Presteriet Values Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value Presteriet Value	Aug 16, 2023, 9:5	52 AM				PersistentVolumeClaim (EVO) dev-vol1-azapos2-0			
<pre>11 volume.touchmetes.io/provisioner-deftion-secret-namespace: default 12 finalizers: 13 - kubernetes.io/pro-protection 14 - external-articher/nspc-csi-hitachi-com 15 > managedFields:</pre>	2 a 3 4 5 6 7 8 9 10	<pre>pzVersion: v1 etadata: name: pvc-0ba60c3e-0a8c uid: a10b5736-bc7f-4867 resourceVersion: '14763 creationTimestamp: '202 annotations: pv.kubernetes.io/prov volume.kubernetes.io/</pre>	-468d-bb64-8a25e6dcbc72 -b99f-f85ea2ac873c 459' 3-08-16T04:22:092' isioned-by: hspc.csi.hitach provisioner-deletion-secret	i.com -name: secret-vsp5200					
<pre>14 - external-attacher/htpc-csi-hitachi-com 15 > manggedFields: 75 Spc: 76 capacity: 77 storage: 20061 77 storage: 20061 78 cSi: 79 driver: hspc.csi.hitachi.com 80 yolumetendie: 01scsiistachi.com 80 yolumetendie: 01scsiistachi.com 81 fsType: ext4 82 volumetendie: 01scsiistachi.ext 83 hostvodcoption: '' 84 size: 20061 85 portifs: '' 86 nickname: spc-bila3a0ed9 87 ports: cli.4 88 lde/TDex: '00181' 89 connectionType: lassi 80 yolumetendie: 129' 92 controllerPublishSecretEef: 93 name: secret-vsp5200 94 namespace: default 95 nodeStagEscretEef: 96 name: secret.vsp5200</pre>	11 12 13	finalizers: - kubernetes.io/pv-pr	provisioner-deletion-secret otection	-namespace: default					
<pre>75</pre>	14 15 >	 external-attacher/h managedFields: 	spc-csi-hitachi-com						
<pre>77 storage: 20061 78 csi: 79 driver: hspc.csi.hitachi.com 80 volumetandle: 01scsi-90000040028129spc-bl18300ad9 81 fsType: ext4 82 volumettributes: 83 hostWodeoption: '' 84 size: 2006i 85 portIFs: '' 86 nickname: spc-bl18300ad9 87 ports: CL-C 88 lidevTDMex: '00:81' 89 connectionType: iscsi 89 storage.twb:shcertef: 90 storage.twb:iskscertef: 91 name: spc-reteft: 93 name:spc: default 94 namespace: default 95 nodeStatestertef: 96 name: secret.vsp5200 97 name:spc: twsp5200 98 modePublishScertef: 99 modePublishScertef: 99 modePublishScretef: 90 modePublishScretef: 91 mame:spc: vsp5200 92 modePublishScretef: 93 modePublishScretef: 94 mamespace: vsp5200 95 modePublishScretef: 96 mame: secret.vsp5200 96 mamespace: vsp5200 97 mamespace: vsp5200 98 modePublishScretef: 99 modePublishScretef: 90 mamespace: vsp5200 90 mate: secret.vsp5200 91 mate: secret.vsp5200 92 modePublishScretef: 93 mate: secret.vsp5200 94 mate: secret.vsp5200 95 modePublishScretef: 95 modePublishScretef: 96 mate: secret.vsp5200 97 mate: secret.vsp5200 98 modePublishScretef: 99 modePublishScretef: 90 mate: secret.vsp5200 90 modePublishScretef: 90 mate: secret.vsp5200 91 mate: secret.vsp5200 92 modePublishScretef: 93 mate: secret.vsp5200 94 mate: secret.vsp5200 95 modePublishScretef: 96 mate: secret.vsp5200 96 mate: secret.vsp5200 97 mate: secret.vsp5200 98 modePublishScretef: 99 modePublishScretef: 90 modePubli</pre>	75 s	pec:							
<pre>78 cs: 79 driver:hspc.csi.hitachi.com 80 volumetandle: 01scsi-90000040028129spc-b118300ad9 81 ff5Type:extA 82 volumettriButes: 83 hostWodeOption: '' 84 size: 2006i 85 portIFs: '' 86 nickiname: spc-bila300ad9 87 ports: CL-C 88 ldevIDMex: '00:81' 89 connectionType: iscsi 89 storage.tubernets.io/csiProvisionerIdentity: 1690266643750-8081-hspc.csi.hitachi.com 91 ldevIDMex: '129' 92 controllerVulishsceretef: 93 name: secret.vsp5200 94 namespace: default 95 nodeStageStorettef: 96 name: secret.vsp5200 97 namespace: default 98 nodeStageStorettef: 99 nodeStageStorettef: 90 nodeStageStorettef: 90 nodeStageStorettef: 91 nodeStageStorettef: 92 nodeStageStorettef: 93 nodeStageStorettef: 94 nodeStageStorettef: 95 nodeStageStorettef: 95 nodeStageStorettef: 96 nodeStageStorettef: 97 nodeStageStorettef: 98 nodeStageStorettef: 99 nodeStageStorettef: 99 nodeStageStorettef: 99 nodeStageStorettef: 90 nodeStageStoret</pre>	77	storage: 200Gi							
80 volumetandLe: d1scji-90000040028-129spc-bl183004d9 81 f5Type: ext4 82 volumettributes: 83 hostVodedption: '' 84 size: 2004i 85 portIFs: '' 86 nickname: spc-bl18300ad9 87 ports: Cl-C 88 ldevIDMe: '0*81' 89 connectionType: iscsi 89 storage.twemetes.io/csiProvisionerIdentity: 1690266643750-8001-hppc.csi.hitachi.com 91 ldevIDMe: 'scsiProvisionerIdentity: 1690266643750-8001-hppc.csi.hitachi.com 92 controllerPublishsceretef: 93 name: secret-xsp5200 94 namespace: default 95 nadeStageScrettef: 96 name: secret-tesp5200 97 namespace: default 98 modePublishScrettef: 99 modePublishScrettef: 99 modeScrettef: 99 namespace: default 99 modeScrettef: 99 modeScrettef: 99 modeScrettef:	78 79	csi: driver: hspc.csi.hita	chi.com						
<pre>s2 volumestributes: s3 hostbodeOption: '' s4 size: 2006i s5 pottPs: '' s6 nickname: spc-bila300ad9 s7 pots: clC s8 ldevIDex: '0e:Bl' s0 connectionType: iscsi s0 storage.kubernetes.j0.cisProvisionerIdentity: 1090256643750-8001-hspc.csi.hitachi.com 91 ldevIDex: '129' s0 controllerublishscretatef: 93 name: secret-vsp5200 94 namespace: default 95 nodeStageScretatef: 96 name: secret-vsp5200 97 namespace: default 98 nodeStageScretatef: 99 nodeStageScretatef: 99 nodeStageScretatef: 99 nodeStageScretatef: 99 nodeStageScretatef: 99 nodeStageScretatef: 99 nodeStageScretatef: 99 nodeStageScretatef:</pre>	80 81	volumeHandle: 01scs	i900000040028129spc-b	11a3a0ad9					
<pre>83 hostVodeOption: '' 84 size:2006i 85 potTPs: '' 86 nickname: spc-blia308ad9 87 pots: cl C 88 ldevTDev: '0e181' 89 connectionType: iscsi 99 storage.kuewnetks.io/csiProvisionerIdentity: 1690266643750-8001-hspc.csi.hitachi.com 91 ldevTDev: '129' 92 controllerivblishSecretRef: 93 name: secret_vsp5200 94 namespace: default 96 name: secret.vsp5200 97 namespace: default 98 nodePublishSecretRef: 99 nodePublishSecretRef: 99 nodePublishSecretRef: 99 name: secret.vsp5200 99 name: secret.vsp5200 99 namespace: default</pre>	82	volumeAttributes:							
85 pottps:'' 86 nickname: spc-blis3a0ad9 87 pots: clC 88 ldevIDMey: 'a0eBi' 99 storage.kubernetes.io/csiProvisionerIdentity: 1690266643750-8001-hspc.csi.hitachi.com 91 ldevIDMey: 'lz.9' 92 controllerAublishGerettef: 93 name: secret.vsp5200 94 namessace: default 95 nodeStageScrettef: 96 name: secret.vsp5200 97 name: secret.vsp5200 98 modePublishSecrettef: 99 name: secret.vsp5200	83 84	hostModeOption: " size: 200Gi							
87 ports: cl.C 88 ldevIDHex: '00:81' 89 connectionType: iscsi 90 storage.Aubernetes.io/csiProvisionerIdentity: 1650266643750-8081-hspc.csi.hitachi.com 91 ldevIDDe: '123' 92 controllerMulishGeretRef: 93 name: secretNef: 94 namesspace: default 95 nodeStageScretRef: 96 name: secretNef: 97 name: secret.vsp5208 98 modePublishSecretRef: 99 modeStageScret.vsp5208	85 86	portIPs: ''	anada						
<pre>88 lde/tDHex: '00:81' 89 connection/pe: iscsi 90 storage.kukennetes.io/csiProvisionerIdentity: 1090256643750-8001-hspc.csi.hitachi.com 91 lde/tDDec: '129' 92 controller#ublishsecretRef: 93 name:sacret-vsp5200 94 name:sacret-vsp5200 95 name:sacret-vsp5200 96 name:sacret-vsp5200 97 name:sacret-vsp5200 98 nodePublishSecretRef: 99 name:sacret-vsp5200 99 name:sacret-vsp5200 90 storage.ideFault 90 storage.ideFault 91 storage.ideFault 92 storage.ideFault 93 storage.ideFault 94 storage.ideFault 95 storage.ideFault 95 storage.ideFault 96 storage.ideFault 97 storage.ideFault 98 storage.ideFault 99 storage.ideFault 90 storag</pre>	87	ports: CL1-C	00005						
90 storage.kubernets.io/csiProvisionerIdentity: 1690266643750-0001-hspc.csi.hitachi.com 91 ldevIDDec: '129' 92 controllerPublisheeretRef: 93 name:sacret.vsp5200 94 name:sacretsef: 95 modeStageScretRef: 96 name:sacret.vsp5200 97 namespace: default 98 nodePublishScretRef: 99 namespace: default 98 nodePublishScretRef: 99 name:sacret.vsp5200	88 89	ldevIDHex: '00:81' connectionType: isc	si						
91 Interface: 92 controller/unlishecretRef: 93 name:sace: 94 namespace: 95 nodeStageScretRef: 96 name:sace: 97 namespace: 98 nodeHullishScretRef: 99 name:sace: 99 name:sace:	90	storage.kubernetes.	io/csiProvisionerIdentity:	1690266643750-8081-hspc.csi.hitach	hi.com				
93 name: secret-vsp5200 94 namespace: default 95 nodestagescretRef: 96 name:sac: default 97 namespace: default 98 nodePublishEscretRef: 99 name:scret-vsp5200	91	controllerPublishSecr	etRef:						
<pre>95 nodeStageScrettef: 96 name:secret-vsp5200 97 namespace: default 98 nodePublishScrettef: 99 name:secret-vsp5200</pre>	93 94	name: secret-vsp520	0						
96 name:space: 97 namespace: 98 node+ublishSecretBef: 99 name:space:	95	nodeStageSecretRef:							
98 nodePublishSecretRef: 99 name: secret-vsp5200	96 97	name: secret-vsp520 namespace: default	0						
37 Hallic. Sculet-YSpS200	98	nodePublishSecretRef:	0						
100 namespace: default	100	namespace: default							

2. Access the MySQL application.

102name: secret-vsp5200103namespace: default

a. Log in to MySQL pod and verify whether the 200 GB persistent volume is created and mounted in /var/lib/mysql.

e azapps2	P azapps2-0 € Running												
Details Metr	rics	YAML	. Ε	nviro	nment	Logs	Events	Terminal					
Connecting to 🧿	azapps2	2											
sh-4.2% df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/sda4 /dev/sdf tmpfs tmpfs tmpfs tmpfs tmpfs	Size 128G 64M 7.9G 64M 7.9G 128G 196G 14G 7.9G 7.9G 7.9G	Used 34G 0 0 51M 34G 271M 24K 0 0	Avai1 956 64M 7.96 64M 7.86 956 1866 1866 146 7.96 7.96 7.96	Use% 27% 0% 0% 1% 27% 1% 1% 0% 0%	Mounted / /dev /sys/fs, /dev/shu /etc/pa: /etc/hos /var/lil /run/sed /proc/ai /proc/si /sys/fin	on (cgroup sswd sts)/mysql crets/kube cpi csi mware	rnetes.io/:	serviceaccount					

- b. Log in to MySQL database using mysql -u root -p.
- c. Verify whether the "devmysqldb1" database is created.
- d. Select the "devmysqldb1" database.

Project: de	vapps 🝷					
Pods > Po	d details					
🕑 azap	ops2-0	2 Running				
Details	Metrics	YAML	Environment	Logs	Events	Terminal
Connecting	to 🙆 azan	nc?				
Connecting	to 👅 azap	ps2				
mysql>						
mysql> s +	how databas	es;				
Databa	se					
inform	ation_schem	ia				
devmys	qldb1					
perfor	nance scher	la I				
sys						
+		·+				
5 POWS 1	n set (0.00	sec)				
mysql> u	se devmysql	.db1;				
Database	changed					

e. Create a table "employee" and ingest some records to the table.

Project: devapps	-					
Pods > Pod details P azapps2-	0 🛛 Running					
Details Metric	s YAML	Environment	Logs	Events	Terminal	
mysql> select * + id name	from employee +	; # 				
1 Jishan	Jishan@abc.	com				
2 Amrit 3 Adip	Amrit@abc.co	n l				
4 KSing	KSing@abc.c	m				
++	+	+				
5 rows in set (6	.00 sec)	+				
nysql>						

f. HSPC automatically creates an iSCSI target on port CL1-C of the storage system. Verify whether the 200 GB volume was created in the VSP 5200 storage system from Storage Navigator.

spc-0057925ee62726518aeb7340a491 (13)										Last Up	dated : 2023/08/1	.6 04:27 🐧
VSP-5200-SV10(S/N:40028) > Ports/Host Groups/iSCS	<u>I Targets</u> > <u>CL1-C</u> > spc-0057925ee627	26518aeb7340a4	9									
Volume Migration 🔻												
iSCSI Target Alias	spc-0057925ee62726518aeb7340a491	(13)		÷	lost Mode				00 [Standard]			
iSCSI Target Name	iqn.1994-04.jp.co.hitachi:rsd.r90.t.400	28.1c013		F	Port Security				Enabled			
Port ID	CL1-C			A	uthentication		Method		Comply with Host	Setting		
Virtual Storage Machine	VSP 5200, 5600 / 40028						Mutual CH	AP	Disabled			
							User Nam	•				
Hosts LUNS Host Mode Options CHAP Us	sers			^								
Add LUN Paths Copy LUN Paths Edit Comman	nd Devices More Actions										Selected	l: 0 of 4
Select All Pages Column Select	ttings									Options 👻	€€1	1 ⇒ →
			Capacity				Used Capac	ty				
Port ID LUN ID LDEV ID LDEV	V Name (ID)	Emulation Type	Total 1▼	Reserved	Used	Used (%)	Tier 1	Tier 2	Capacity Saving	Capacity Saving Status	++ Provisioning Type	CLPR
CL1-C 277 00:00:81 spc-	-b11a3a0ad9 dr_pool(0)	OPEN-V CVS	200.00 GB	0.00 GB	7.46 GB	3	-	-	Disabled	Disabled	DP	0:CLPR0

Test 3: Manually Migrate Stateful Applications Across OpenShift Clusters

This test case describes the process of migrating a stateful application from OpenShift cluster in Azure to AWS using Kubernetes commands and HSPC. The VSP 5200 storage system provides the persistent volume required for stateful MySQL applications in both clusters.

Snapshot Operation

To perform snapshot operation of an application in Azure Red Hat OpenShift Cluster, complete the following steps:

To create a MySQL application with a persistent volume of 200 GB from the VSP 5200 storage system and ingest data to the database, see the section <u>Test 2: Deploy a Stateful Application in Azure Red Hat OpenShift Cluster</u>. The following screenshot shows that the data is available in the MySQL application.

Project: d	evapps •	•					
Pods > P	od details						
🕑 aza	pps2-(C Running					
Details	Metric	s YAML	Environment	Logs	Events	Terminal	
Connectin	gto Ċ az	apps2					
mysql> +	select * :	from employee;	+				
id	name	email					
1	Jishan	Jishan@abc.c	com				
2	Amrit	Amrit@abc.co	m				
4	Adip KSing	KSing@abc.com	1 MTT				
5	JackM	JackM@abc.co	100				
+	+	+	+				
5 rows	in set (0	.00 sec)					
mysql>							

- 1. Create a snapshot.
 - a. Create a snapshot of the persistent volume created for the application. From the Red Hat OpenShift console, navigate to Storage, click VolumeSnapshots, and then click Create VolumeSnapshot. In the Create VolumeSnapshot menu, enter the required information such as PVC, snapshot name, snapshot class, and click Create.

Project: devapps 🔹		
		PersistentVolumeClaim details
Create VolumeSnapshot	Edit YAML	Name
·		PVC dev-vol1-azapps2-0
PersistentVolumeClaim *		
PVC dev=voll=azapps2=0		Namespace
det ten drapper o		NS devapps
Name *		
dev-voll-azanns2-0-snanshot		Status
		O Bound
Snapshot Class *		StorageClass
VSC snanshotclass_cample	-	SC sc-vsp5200
and and an array of the second s	•	
		Requested capacity
Create Cancel		200 GiB
		Access mode
		Single user (RWO)
		Volume mode
		Filesenters

b. Verify whether the snapshot is created.

Project: devapps 🔹							
VolumeSnapshots						Create Vol	umeSnapshot
▼ Filter ▼ Name ▼	Search by name 7						
Name 1	Status 💲	Size 1	Source 1	Snapshot content 📫	VolumeSnapshotClass 1	Created at 🕴	
VS dev-vol1-azapps2-0- snapshot	🖉 Ready	200 GiB	PVC dev-vol1-azapps2-0	VSC snapcontent-fd711d01- bd5d-4621-a04d- 6971145f2c65	VSC snapshotclass-sample	🚱 Aug 16, 2023, 10:47 AM	:

c. In Storage Navigator, verify whether the snapshot volume 00:00:89 is created successfully.

istory											
Copy Type: TI	•										l€ € 1 / 1 € €
TI History (Page.1)											
\$Filter ON OFF											
Date and Time	Primary Voli	ime	Secondary V	olume	Mirror	Daol ID	Diff Compare Volume	Description	Description		î
	LDEV ID	Provisioning Type	LDEV ID	Provisioning Type	Unit	P00110	bir compare volume	Code	Charle Provin		
2023/08/16 05:18:06	00:00:81	DP	00:00:89	DP	3	0		2011	PSUS		
2023/08/16 05:18:03	00:00:81	DP	00:00:89	DP	3	0	·	2001	PAIR		
	Copy Type: TI TT History (Page.1) RFilter SFilter ON OFF Date and Time 2023/08/16 05:18:06 2023/08/16 05:18:03 2023/08/16 05:18:03	Primary (Page.1) A Filter ON OFF Data and Time 2222/06/16 05:18:00 2022/06/16 05:18:00	Date and Time Primary Volume 2022/06/16 0518/06 000081 De 2022/06/16 0518/06 000081 DP	TI Image: Constraint of the second of the sec	Date and Time Prime volume Secondary Volume 2023/06/16 05:18:0 Option:10 Provisioning Type DBV /D Provisioning Type 2023/06/16 05:18:00 Option:10 Option:10 Option:10 DP 2023/06/16 05:18:00 Option:10 DP Option:10 DP	Dep Types: TI TI TI History (Depc.1) Arise: 0 original Date and Time Primary Volume Secondary Volume Minor Date and Time EDP 10 Rowsioning Type LDP/ 10 Porvisioning Type Unit 2023/db(16 05:t8)66 Oxide(16 DP) 00:00:82 DP 3 3	Determine Primary Volume Secondary Volume Mirror Pool ID Date and Time EDD /D Provisioning Type LDR /D Pool ID Pool ID 2022/db(16 05x18)60 Op/db(16) DP 00/db(16) DP 3 0	Provide TIT TIT USONY (Project.) TIT USONY (Project.) TIT USONY (Project.) Date and Time Mercory Volume Mercory Volume Date and Time Provisioning Type DDV 10 Provisioning Type Unit. Peol 10 Diff Compare Volume 2023/06/16 65181.06 600.068.1 DP 000/0018 0 - 2023/06/16 65181.06 000.068.1 DP 000/0018 0 -	Between Volume Marrier Original Colspan="2">Def Compare Volume Def Compare Volume	Prime TI I I I I I I I I I I I I I I I I I I	Prime Trime Prime P

Restore Operation

To restore an application in an OpenShift cluster in AWS, complete the following steps:

- 1. Identify the volume handle string for the snapshot secondary volume 00:00:89. Volume handle string for this volume is "60060e80089c5c0000509c5c0000**0089--spc-1d1bccb46c**".
- 2. Create a PV and PVC using the volume 00:00:89 with the pre-defined volume handle string.
 - a. Create a project called "devapps" for the MySQL application.

Projects						С	reate Project
▼ Filter ▼ Name ▼	deva						
Name deva 🗙 Clear all fi	iters						
Name 💲	Display name 🌐	Status 1	Requester 1	Memory 1	CPU 1	Created 1	
PR devapps	No display name	Active	kube:admin	-	-	🚱 Aug 16, 2023, 5:43 AM	:

b. Create a manifest file for PV using the volume handle string. This way, storage class does not dynamically create a new volume. Instead, it uses the existing volume to preserve the snapshot data. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumes**, and then click **Create PersistentVolume**. In the Create PersistantVolume menu, populate the YAML file with the required information and click **Create**.

Create b	te PersistentVolume by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.
1	apiVersion: v1
2	kind: PersistentVolume
3	metadata:
4	name: devappspv
5	namespace: devapps
6	spec:
7	capacity:
8	storage: 200Gi
9	accessModes:
10	- ReadWriteOnce
11	persistentVolumeReclaimPolicy: Retain
12	storageClassName: sc-vsp5200
13	csi:
14	driver: hspc.csi.hitachi.com
15	volumeHandle: 60060e80089c5c0000509c5c00000089spc-1d1bccb46c
16	claimRef:
17	name: devappspvc
18	namespace: devapps
19	

c. Verify whether the PV is created.

PersistentVolumes					Create Persistent	Volume
Name 👻 Search by name						
Name 🗍	Status 🗍	Claim 1	Capacity 🗍	Labels 🔱	Created 1	
PV devappspv	Available	(PVC) devappspvc	200Gi	No labels	Aug 16, 2023, 6:39 AM	:

d. Create a manifest file for PVC using the PV created in step 2b. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumeClaims**, and then click **Create PersistentVolumeClaim**. In the Create PersistantVolume menu, populate the YAML file with the required information and click **Create**.



e. Verify whether the PVC is created.

Project: devapps 🔹						
PersistentVolu	meClaims				Create Persiste	ntVolumeClaim
▼ Filter ▼ Name	e Visit Search by name 7					
Name 🗍	Status 💲	PersistentVolumes 1	Capacity 💲	Used 1	StorageClass 🔱	
PVC devappspvc	Bound	PV devappspv	200 GiB	-	SC sc-vsp5200	:

- 3. Create a clone PVC using the "devappspvc" PVC as data source.
 - a. Create a manifest file. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumeClaims**, and then click **Create PersistentVolumeClaim**. In the Create PersistantVolume menu, populate the YAML file with the required information and click **Create**.



b. HSPC dynamically provisions a PV from the VSP 5200 storage system. Verify whether the PVC and PV are created.

Project: devapps 🔹						
ersistentVolume	eClaims				Create Pers	istentVolume
Filter 🝷 Name	Search by name					
Name 1	Status 1	PersistentVolumes 1	Capacity 1	Used 1	StorageClass	
PVC devappscione	🕏 Bound	PV pvc-99561b45-0f6f-498c- bd98-c09053516fff	200 GiB	-	SC sc-vsp5200	:
sistentVolumes > Pers	stentVolume details					
pvc-99561t	045-0†6†-498c-b	od98-c09053516ttt ⊚ №	und			
tails YAML						
1 kind: Ponsist	ent\/olume					
2 apiVersion: v	1					
3 metadata:						
4 name: pvc-9	9561b45-0f6f-498c-bd98-c	09053516fff				
5 uid: 9a2276	c1-0706-42d3-80a3-e9722e	fd5c72				
6 resourceVer	sion: '17201901'					
7 creationTim	estamp: '2023-08-16T06:4	6:48Z'				
8 annotations						
9 pv.kubern	etes.io/provisioned-by:	hspc.csi.hitachi.com				
10 volume.ku	bernetes.io/provisioner-	deletion-secret-name: secret-vsp52	00			
11 volume.ku	bernetes.io/provisioner-	deletion-secret-namespace: default				
12 finalizers:						
13 - kuberne	tes.io/pv-protection					
14 > managedFiel	ds:…					
72 spec:						
73 capacity:						
74 storage:	200Gi					
75 csi:						
76 driver: h	spc.csi.hitachi.com					
77 volumeHan	dle: 01scsi900000040	028140spc-9feee3c777				
78 fsType: e	xt4					
79 volumeAtt	ributes:					
80 hostMod	eoption:					
81 51Ze: 2	0001					
82 portIPs	: 					
83 nicknam	e: spc-9teee3c777					
oq ports:	CLI-C					
as IdevIDH						
connect	In the second					

The dynamically created PV is the Thin Image clone volume. In the following screenshot, volume 00:00:8C is the designated clone volume.

н	isto ry									
	Сору Туре: ТІ 🛛									
	TI History (Page.1)			_		_	_		_	
	☆ Filter ON OFF									
	B	Primary Volu	ime	Secondary V	olume	Mirror	0.170	D:// 0	Description	
	Date and Time	LDEV ID	Provisioning Type	LDEV ID	Provisioning Type	Unit	POOLID	Dirr Compare volume	Code	Description
	2023/08/16 06:50:39	00:00:89	DP	00:00:8C	DP	3	0	-	2092	CLONE END
	2023/08/16 06:46:52	00:00:89	DP	00:00:8C	DP	3	0	·	2091	CLONE START
	2023/08/16 06:46:51	00:00:89	DP	00:00:8C	DP	3	0	-	2001	PAIR

- 4. Restore the MySQL application in the AWS cluster.
 - a. Create a MySQL service. From the Red Hat OpenShift console, navigate to **Networking**, click **Services**, and then click **Create Service**. In the Create Service menu, populate the YAML file with the required information and then click **Create**.

Create Service Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.	
1 apiVersion: v1 2 kind: Service	
3 metadata:	
4 namespace: devapps	
5 name: azappsz	
9 norts	
10 - port: 3306	
11 name: azapos2	
12 clusterIP: None	
13 selector:	
14 app: azapps2	

b. Verify whether the MySQL service is created.

Project: devapps 🔻				
Services				Create Service
Name Search by name 7				
Name 🗘	Labels 💲	Pod selector 1	Location 1	
S azapps2	app=azapps2	Q app=azapps2	None	:

c. Create a MySQL statefulset application. From the Red Hat OpenShift console, navigate to Workloads, click StatefulSet, and then click Create StatefulSet. In the Create StatefulSet menu, populate the YAML file with the required information and click Create. In the volume section, use the "devappsclone" claim created in step 3. This ensures that the MySQL application uses the clone PVC for persistent data.

Project:	devapps 👻
Creat Create by	The StatefulSet manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.
1	apiVersion: apps/v1
2	kind: StatefulSet
4	namesnace: devanns
5	name: azapps2
6	spec:
7	selector:
8	matchLabels:
9	app: azapps2
10	serviceName: "azapps2"
11	podManagementPolicy: Parallel
12	replicas: 1
13	template:
14	metadata:
15	labels:
16	app: azappsz
10	Spel:
10	containers:
20	- name: azanns2
21	image: mysol:5.7
22	args:
23	- "ignore-db-dir=lost+found"
24	env:
25	- name: MYSQL_ROOT_PASSWORD
26	value: pass123
27	- name: MYSQL_DATABASE
28	value: devmysqldb1
29	- name: MYSQL_USER
30	
32	- Hame: MTSQL_FASSWORD
33	ports:
34	- containerPort: 3306
35	name: mysal
36	volumeMounts
37	- name: dev-vol1
38	mountPath: /var/lib/mysql
39	volumes:
40	- name: dev-vol1
41	persistentVolumeClaim:
42	

d. Verify whether the StatefulSet is running.

Project: devapps 👻							
StatefulSets > StatefulSet details							
Details Metrics YAML Po	ds Environment Events						
▼ Filter ▼ Name ▼ Search by	name						
Name †	Status 👔	Ready 1	Restarts 1	Node 1	Memory I	CPU I	Created 1
e azapps2-0	2 Running	1/1	0	🚯 ip-10-77-28-190.us-west-1.compute.internal	10	-	🚱 Aug 16, 2023, 6:59 AM

e. Log in to pod azapps2-0 and verify whether the 200 GB persistent volume is mounted on /var/lib/mysql.

ods > Pod details									
azapps2-0) CR	unning							
Details Metric	s Ya	AML	Envi	ronment	Logs	Events	Terminal		
Connecting to 🔗 a	zapps?								
Johnecung to 🕒 a	zuppsz.								
connecting to 🤤 a	zabhaz								
sh-4.2\$ df -h	20pp32								
sh-4.2\$ df -h Filesystem	Size	Used	Avail	Use% Moun	nted on				
sh-4.2\$ df -h Filesystem overlay	Size 100G	Used 28G	Avail 73G	Use% Moun 28% /	nted on				
sh-4.2\$ df -h Filesystem overlay tmpfs	Size 100G 64M	Used 28G Ø	Avail 73G 64M	Use% Moun 28% / 0% /dev	nted on				
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs	Size 100G 64M 7.8G	Used 28G Ø	Avail 73G 64M 7.8G	Use% Moun 28% / 0% /dev 0% /sys	nted on / s/fs/cgrou	up			
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm	Size 100G 64M 7.8G 64M	Used 28G 0 0	Avail 73G 64M 7.8G 64M	Use% Moun 28% / 0% /dev 0% /sys 0% /dev	nted on / s/fs/cgrou //shm	up			
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs	Size 100G 64M 7.8G 64M 7.8G	Used 28G 0 0 54M	Avail 73G 64M 7.8G 64M 7.7G	Use% Moun 28% / 0% /dev 0% /sys 0% /dev 1% /etc	nted on / //fs/cgrou //shm :/passwd	цр			
sh-4.2\$ df -h Filesystem overlay tmpfs shm tmpfs /dev/nvme0n1p4	Size 100G 64M 7.8G 64M 7.8G 100G	Used 28G 0 0 54M 28G	Avail 73G 64M 7.8G 64M 7.7G 73G	Use% Moun 28% / 0% /dev 0% /sys 0% /dev 1% /etc 28% /etc	nted on ; ;/fs/cgrou //shm :/passwd :/hosts	up			
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde	Size 100G 64M 7.8G 64M 7.8G 100G 196G	Used 28G 0 0 54M 28G 271M	Avail 73G 64M 7.8G 64M 7.7G 73G 186G	Use% Moun 28% / 0% /dev 0% /sys 0% /dev 1% /etc 28% /etc 1% /var	oted on ; ;/fs/cgrou //shm :/passwd :/hosts /lib/mysd	up			
sh-4.2\$ df -h Filesystem overlay tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde tmpfs	Size 100G 64M 7.8G 64M 7.8G 100G 196G 15G	Used 28G 0 0 54M 28G 271M 20K	Avail 73G 64M 7.8G 64M 7.7G 73G 186G 15G	Use% Moun 28% / 0% /dev 0% /dev 1% /etc 28% /etc 1% /var 1% /run	ited on / //fs/cgrou //shm :/passwd :/hosts //lib/mysu //secrets,	up ql /kubernete	s.io/servicea	ccount	
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde tmpfs tmpfs	5ize 100G 64M 7.8G 64M 7.8G 100G 196G 15G 7.8G	Used 28G 0 54M 28G 271M 20K 0	Avail 73G 64M 7.8G 64M 7.7G 73G 186G 15G 7.8G	Use% Moun 28% / 0% /dev 0% /sys 0% /dev 1% /etc 28% /etc 1% /var 1% /run 0% /pro	ited on //s/fs/cgrou //shm :/passwd :/hosts //lib/mysc //secrets, oc/acpi	up ql /kubernete	s.io/servicea	ccount	
sh-4.2\$ df -h filesystem overlay tmpfs shm tmpfs /dev/vsme0n1p4 /dev/vsde tmpfs tmpfs tmpfs	Size 100G 64M 7.8G 64M 7.8G 100G 196G 15G 7.8G 7.8G	Used 28G 0 54M 28G 271M 20K 0 0	Avail 736 64M 7.86 64M 7.76 736 1866 156 7.86 7.86	Use% Moun 28% / 0% /dev 0% /dev 0% /dev 1% /dev 1% /etc 1% /var 1% /run 0% /pro	<pre>//inted on //ishm //ishm //ishm //ib/mysd //iscrets/ //iscrets/ //isc/acpi //isc/acpi //isc/sci</pre>	up ql /kubernete	s.io/servicea	ccount	

f. Log in to MySQL and verify whether the "devmysqldb1" database is available.

Project: devapps 🔹						
Pods > Pod details Pod azapps2-0	C Running					
Details Metrics	YAML	Environment	Logs	Events	Terminal	
Connecting to 🬀 azap	ops2					
<pre>> Joatabase > information_sch devmysqldb1 mysql performance_sch sys </pre>	<pre></pre>	for completion c ire to get a qui	of table icker sta	and column rtup with	names -A	

g. Verify whether the ingested data from Azure RedHat OpenShift cluster is available here.

Pods > Po	od details pps2-0	C Running				
Details	Metrics	YAML	Environment	Logs	Events	Terminal
Connecting mysql> mysql>	g to ⓒ azag	pps2 from employe	-e:			
+						
id	name	email	+			
id + 1 2 3 4 5	name Jishan Amrit Adip KSing JackM	email Jishan@abc Amrit@abc. Adip@abc.c KSing@abc. JackM@abc.				

HSPC automatically creates an iSCSI target on port CL1-C of the storage system and assigns the volume to the appropriate worker node.

spc-190210a78dfc87a6f821889c2552 (0D)										_
VSP-5200-SV10(S/N:40028) > Ports/Host Groups/iSCS	<u>I Targets</u> > <u>CL1-C</u> > spc-190210	a78dfc87a6f8218	89c2552							
Volume Migration 🔻										
iSCSI Target Alias	spc-190210a78dfc87a6f821889	c2552 (0D)			Host Mode					00 [Standard]
iSCSI Target Name	iqn.1994-04.jp.co.hitachi:rsd.r9	0.t.40028.1c00d			Port Securi	ty				Enabled
Port ID	CL1-C				Authenticat	tion	Met	hod		Comply with Host Setting
Virtual Storage Machine	VSP 5200, 5600 / 40028						Mut	ual CHAP		Disabled
							Use	r Name		
Hosts LUNs Host Mode Options CHAP Us	sers			·	`					
Add LUN Paths Copy LUN Paths Edit Comman	nd Devices More Actions 🔻									
Select All Pages Column Set	ttings									Op
	Emulation	Pool Name	Capacity				Used Capaci	ty		Capacity Saving
	Туре	(ID)	Total	Reserved	Used 1 V	Used (%)	Tier 1	Tier 2	Tier 3	Capacity Saving
CL1-C 2135 00:00:8C spc-	9feee3c777 OPEN-V CVS	dr_pool(0)	200.00 GB	0.00 GB	6.31 GB	3	-	-	-	Disabled

Test 4: Migrate a Stateful Application Across OpenShift Clusters Using Kasten K10 Multi-Cluster

This test case describes the process of migrating a stateful application by performing backup and restore operations between two OpenShift clusters using Kasten K10 Multi-Cluster Global policy and HSPC. In this environment, backup and restore operations can be performed among three clusters residing on near-cloud, AWS, and Azure. To demonstrate this, we captured the snapshot of a stateful MySQL application running on Azure Red Hat OpenShift cluster and then restored it on a Red Hat OpenShift cluster in AWS. The entire process was performed from the Kasten K10 Multi-Cluster UI. The VSP 5200 storage system serves the persistent volumes required for stateful MySQL application in both clusters.

Snapshot Operation

Before performing snapshot operation of an application using Kasten K10, create the following:

- Global Location Profile: Profiles define credentials and locations required to move the data in and out of the cluster. In this scenario, an Amazon S3 bucket is used.
- Global Policy: Policies are used to automate your data management workflows. To achieve this, they combine actions you want to take (such as snapshot), a frequency or schedule for how often you want to take that action, and a label-based selection criteria for the resources you want to manage.
- Distribution: Distributions define which K10 resources belong to which clusters.

To perform snapshot operation of an application using Kasten K10, complete the following steps:

- 1. Create a new project "devmysql" and deploy a stateful MySQL application with a 100 GB persistent volume from the VSP 5200 storage system, as shown in the section <u>Test 2: Deploy a Stateful Application in Azure Red Hat OpenShift Cluster</u>.
- 2. Access the stateful MySQL application.
 - a. Log in to pod MySQL and verify whether the 100 GB persistent volume is mounted in "/var/lib/mysql", as per the manifest file.

ade b Rod dat	nile							
azapps	s1-0 c	Running						
•								
Details M	etrics	YAML	E	nvironment	Logs	Events	Terminal	
JUNNECCING TO	azabbs							
sh-4.2\$ df - Filesystem	h Size	Used	Avail	Use% Mounted	on			
sh-4.2\$ df - Filesystem overlay	h Size 1286	Used 16G	Avail 113G	Use% Mounted	on			
sh-4.2\$ df - Filesystem overlay tmpfs	h Size 128G 64M	Used 16G 0	Avail 113G 64M	Use% Mounted 12% / 0% /dev	on			
sh-4.2\$ df - Filesystem overlay tmpfs tmpfs	h Size 128G 64M 7.9G 64M	Used 16G 0 0	Avail 113G 64M 7.9G 64M	Use% Mounted 12% / 0% /dev 0% /sys/fs/ 0% /dev/shu	on /cgroup			
sh-4.2\$ df - Filesystem overlay tmpfs tmpfs shm tmpfs	h Size 128G 64M 7.9G 64M 7.9G	Used 16G 0 0 49M	Avail 113G 64M 7.9G 64M 7.8G	Use% Mounted 12% / 0% /dev 0% /sys/fs/ 0% /dev/shm 1% /etc/pa	on /cgroup 11 sswd			
sh-4.2\$ df - Filesystem overlay tmpfs tmpfs shm tmpfs /dev/sda4	h Size 1286 64M 7.96 64M 7.96 1286	Used 16G 0 0 49M 16G	Avail 113G 64M 7.9G 64M 7.8G 113G	Use% Mounted 12% / 0% /dev 0% /sys/fs/ 0% /dev/shm 1% /etc/pas 12% /etc/hos	on /cgroup m sswd			
sh-4.2\$ df - Filesystem overlay tmpfs tmpfs shm tmpfs /dev/sda4 /dev/sde	h Size 128G 64M 7.9G 64M 7.9G 128G 98G	Used 16G 0 0 49M 16G 271M	Avail 113G 64M 7.9G 64M 7.8G 113G 93G	Use% Mounted 12% / 0% /dev 0% /sys/fs, 0% /dev/sm 1% /etc/pas 12% /etc/hos 1% /var/lit	on /cgroup m sswd sts p/mysql			
sh-4.2\$ df - Filesystem overlay tmpfs tmpfs shm tmpfs /dev/sda4 /dev/sde tmpfs	h Size 128G 64M 7.9G 64M 7.9G 128G 98G 14G	Used 16G 0 0 49M 16G 271M 24K	Avail 1136 64M 7.96 64M 7.86 1136 936 146	Use% Mounted 12% / 0% /dev 0% /sys/fs, 0% /dev/shu 1% /etc/pas 1% /var/lit 1% /run/sec	on /cgroup m sswd sts p/mysql crets/kuber	metes.io/s	serviceaccount	
sh-4.2\$ df - Filesystem overlay tmpfs shm tmpfs /dev/sda4 /dev/sda4 tmpfs tmpfs	h Size 1286 64M 7.9G 1286 986 14G 7.9G	Used 16G 0 0 49M 16G 271M 24K 0	Avail 113G 64M 7.9G 64M 7.8G 113G 93G 14G 7.9G	Use% Mounted 12% / 0% /dev 0% /sys/fs, 0% /dev/shm 1% /etc/pas 12% /etc/hos 1% /var/lit 1% /run/sec 0% /proc/ad	on /cgroup m sswd sts b/mysql crets/kuber cpi	metes.io/s	serviceaccount	
sh-4.2\$ df - Filesystem overlay tmpfs tmpfs tmpfs /dev/sda4 /dev/sda tmpfs tmpfs tmpfs	h Size 1286 64M 7.96 64M 7.96 1286 986 1286 986 146 7.96 7.96	Used 16G 0 0 49M 16G 271M 24K 0 0	Avail 113G 64M 7.9G 64M 7.8G 113G 93G 14G 7.9G 7.9G	Use% Mounted 12% / 0% /dev 0% /sys/fs/ 0% /dev/shm 1% /dev/shm 1% /dev/shm 1% /dev/shm 1% /dev/shm 1% /var/lit 1% /var/lit 1% /var/sec 0% /proc/sc	on (cgroup a sswd sts o/mysql crets/kuber cpi csi	metes.io/s	serviceaccount	

b. Log in to MySQL database using "mysql -u root -p" and verify whether the "devmysqldb" database is created as per the manifest file.

mysql> show databases	;
Database	•
information_schema devmysqldb	
mysq1 performance_schema sys	
+	+ ec)

c. Create a table "employee" and ingest some new records to the table.



- 3. Verify that the application created in <u>step 1</u> is registered in the Kasten K10 UI.
- 4. Navigate to the Kasten K10 Multi-Cluster Dashboard and verify the available clusters and registered applications.



5. Navigate to **Clusters**, select **azurecluster**, and view the registered applications. Kasten K10 registers detected namespaces as an application. The following screenshot shows that the namespace "devmysql" is detected as an application.

KASTEN by vecam	🗓 Docs 🛛 k10-multi-cluster-user v 🗘 2
$<$ Clusters $<$ azurecluster \preceq	
Applications View details or perform actions on applications.	Cluster-Scoped Resources M Latest snapshot was jul 31, 638am • Create a Policy >
1 application 😸 🗄 Page 1 🕢 🕥	
devmysql	
Not Protected by Policies f Create a Policy >	
1000 GB 🔒 1 -€1 🎯 1 🗇 8	
anapahot restore export details	

6. Create a Global Location Profile.

a. In the K10 Global Resources section of the K10 Multi-Cluster Dashboard, click Global Profile.

Create and distribute policies and profiles to multiple clusters and	1 Global Profile
manage multi-cluster permissions.	2 Global Policies
	2 Distributions

Global K10 Reso	DUTCES
User Permissions Manage permissions for multi-cluster users	Location Profiles Create profiles that define credentials and locations needed to move data in and out of the cluster. You'll
Location Profiles Manage global cloud location profiles	 ⊕ New Profile
Infrastructure Profiles Manage global infrastructure profiles	No Profiles
Policies Manage global policies	No profiles have been created yet. Create a profile.
Distributions	

b. In this scenario, an Amazon S3 bucket named bucketk10b is created and designated as the destination for Global Location Profile. For instructions to create an S3 bucket, see the <u>Amazon S3 User Guide</u>.

Buckets (13) Info Buckets are containers for data stored in S3. Learn more		C	Copy ARN Empty Delete Create bucket
Q 10b >	X 1 match		< 1 > ©
Name	▼ AWS Region		
O bucketk10b	US West (N. California) us-west-1	Objects can be public	July 31, 2023, 15:00:43 (UTC+05:30)

c. In the Global K10 Resources window, click New Profile.

d. Enter the required information (such as Profile Name, Storage Provider, AWS region, Bucket Name, AWS Access Key, Secret Key, and so on) and click **Save Profile**.



e. Verify that the profile is created.

	BAL PROFILE Belong	gs to the distributions dist-aws-restore, dist-	demo10-restore, dist-demo10-snapshot, dist-onprem-snapshot			
Ø	LOCATION PROFILE	k10b 🛛		>yaml	edit	iii delete
	CLOUD PROVIDER AWS S3	_{REGION} US West (N. California) • us-west-1	BUCKET NAME bucketk10b			

- 7. Create a Global Snapshot Policy.
 - a. From the K10 Global Resources page, click Global Polices and then click New Policy.
 - b. Enter the snapshot related information (such as Policy Name, Backup Frequency, target application, application resources, and so on). Select **Enable Backup via Snapshot Exports**, select the location profile that you created, and click **Create Policy**. This is required to generate an import policy while restoring the application.

	New Policy	
Name		
he display name for this policy		
az-snapshot-policy		
Comments		
Action		
he action that should be taken w	hen this policy is executed	
Snapshot		 Import
Backup Frequency		
 Hourly 	O Daily	O Weekly
 Monthly 	O Yearly	On Demand
Export Location Profile The profile that restore points globalkastenk10b	will be exported to	-
Export Location Profile The profile that restore points $\overrightarrow{l_{O}}$ globalkastenk10b Storage class exceptions	will be exported to	· · ·
Export Location Profile The profile that restore points and profile that restore points globalkastenk10b Storage class exceptions Advanced Export Set	will be exported to	
Export Location Profile The profile that restore points image: start and star	will be exported to ttings aces this policy should target. S	elect applications by name or by
Export Location Profile The profile that restore points globalkastenk10b Storage class exceptions Advanced Export Set Select Applications Droose which application namesp abel. By Name	will be exported to ttings aces this policy should target. S By Labels	elect applications by name or by
Export Location Profile The profile that restore points globalkastenk10b Storage class exceptions Advanced Export Set Select Applications Choose which application namesp abel. By Name Choose one or more applications	will be exported to ttings aces this policy should target. S By Labels to target with this policy.	elect applications by name or by
Export Location Profile The profile that restore points globalkastenk10b Storage class exceptions Advanced Export Set Select Applications Choose one or more applications to Choose one or more applica	will be exported to ttings aces this policy should target. S By Labels to target with this policy.	elect applications by name or by
Export Location Profile The profile that restore points globalkastenk10b Storage class exceptions Advanced Export Set Select Applications hoose which application namesp abel. By Name Choose one or more applications to devmysql x	will be exported to ttings aces this policy should target. S By Labels to target with this policy.	elect applications by name or by
Export Location Profile The profile that restore points	will be exported to ttings aces this policy should target. S By Labels to target with this policy.	elect applications by name or by
Export Location Profile The profile that restore points	will be exported to ttings aces this policy should target. S By Labels to target with this policy. fexclude specified application re	elect applications by name or by None esources.
Export Location Profile The profile that restore points	will be exported to ttings aces this policy should target. S By Labels to target with this policy. rexclude specified application re	elect applications by name or by None esources. Filter Resources
Export Location Profile The profile that restore points	will be exported to ttings aces this policy should target. S By Labels to target with this policy. fexclude specified application re cal Resources	elect applications by name or by None sources. Filter Resources
Export Location Profile The profile that restore points	will be exported to ttings aces this policy should target. S By Labels to target with this policy. fexclude specified application re ced resources that are not capture. CusterRoles, and ClusterRoleBi	elect applications by name or by elect applications by name or by None sources. Filter Resources ed in application snapshots, such as ndings.
Export Location Profile The profile that restore points globalkastenk10b Storage class exceptions Advanced Export Set Select Applications Choose one or more applications of devrmysql × Select Application Resources Sptionally create filters to include All Resource Snapshot Cluster-Scope These include non-namespa Custom Resource Definition	will be exported to ttings aces this policy should target. S By Labels to target with this policy. fexclude specified application re ted Resources that are not captur . ClusterRoles, and ClusterRoleBi Resources Fitz	elect applications by name or by None sources. Filter Resources ed in application snapshots, such as ndings. er Cluster-Scoped Resources

c. Verify that the policy is created.

B GLOBAL POLICY	Not yet added to distribution.	Create a Distribution
az-snapshot-policy		2 edit
devmysql @ cluster-scoped resources		<i>دل></i> yami
Snapshot on-demand for exporting data.		10 delete
 Export onDemand snapshots using the export profile globalkastenk10b Export volume data for durable backups 		

- 8. Create a distribution.
 - a. From the K10 Global Resources page, click **Distributions** and then click **New Distribution**.
 - b. In the Add Distribution window, enter the required information (such as Distribution Name), specify the Azure cluster, specify the two resources created (Global Location Profile and Global Snapshot Policy), and then click **Add Distribution**.

	Ad	dd Distribution	
Nam The di	e isplay name for the distribution. Mu	st be Kubernetes-compatible (lowercase, dots, dasł	hes)
dist	t-az-snapshot-policy		
Clust Using Multip	ters labels, specify the clusters to which ole labels will be unioned (OR). Any o luster - azurecluster ×	you want to distribute resources. luster that matches any label will be targeted.	
Reso Select Availa	urces the global K10 resources to distribution ble Options (5) Sector	ute to clusters. elect All Selected (2)	Deselect All
Reso Select Availa	urces the global K10 resources to distribu- ible Options (5) Se aws-restore Import + restore policy depends on profile globalkastenk10b.	elect All Selected (2)	Deselect All
Reso Select Availa	urces the global K10 resources to distribu- tible Options (5) Se aws-restore Import + restore policy depends on profile globalkastenk10b. demo10-restore Import + restore policy depends on profile globalkastenk10b.	vite to clusters. Selected (2) Image: Selected (2) Image: Snapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Snapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Snapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Snapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Snapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Snapshot + export policy depends profile globalkastenk10b.	Deselect All
Reso Select Availa	urces the global K10 resources to distribu- ble Options (5) se aws-restore Import + restore policy depends on profile globalkastenk10b. demo10-restore Import + restore policy depends on profile globalkastenk10b. demo10-snapshot Snapshot + export policy depends on profile globalkastenk10b.	atte to clusters. Selected (2) Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Selected (2) Image: Sapshot + export policy depends profile globalkastenk10b. Image: Sapshot + export policy depends policy depends policy depends profile globalkastenk10b. Image: Sapshot + export policy depends policy depend	Deselect All

c. Verify that the distribution is added.

dist-az-snapshot-policy	C (♪ 回 道 force sync yami edit delete
CLUSTERS RESOURCES	

- 9. Collect a snapshot of the registered application using the Global Snapshot Policy.
 - a. From the K10 Multi-Cluster Dashboard, click Cluster "azurecluster" and then click Policies.

Verify that the Global snapshot on-demand policy created in the Global Policy section is available under Polices.



b. Click run once, which opens a Run Once window. To start the snapshot, click Yes, continue.



c. Open the Kasten K10 Multi-Cluster dashboard and check the status of the policy in the **Actions** window. To check the phase in progress, click the related action.

total actions	completed actions	failed actions	skipped actions O	avg duration 36 sec	live artifacts 3,965	retired artifacts
Actions (4)			۱۵۶ Filter	v		Page 1 🔇 🕥
40% Policy Run policy-run-t7xds	POLICY 2 az-snapshot-policy	ACTIONS (4) (2) (2)				START Today, 6:59am

d. Verify that the phase has changed to Completed Successfully. Click the relevant action to confirm that no error is present.

Clusters < azu	irecluster v		
\bigcirc	completed successfully az-snapshot-policy policy-run-t7xds Show Details	START END DUBATION Today, 6:59am Today, 7:02am 2 mins, 46 secs APPLICATIONS (All) (devmyse)	
Actions 🗿)		않 Filter Actions
COMPLETED Export policy-run-t7xds	PHASES C Exporting Metadata Monitoring Actions All phases completed successfully.	ROTECTE OBJECT ARTIFACTS none none roucy az-snapshot-policy	START Today, 6:59am DURATION 2 mins, 31 secs
COMPLETED Export scheduled-I5mv	PHASES O Exporting RestorePoint All phases completed successfully.	Mortecteb object Antractis none 647 @ spec roucr az-snapshot-policy	START Today, 6:59am DURATION 1 Sec
COMPLETED Export scheduled-w2tjv	PHASES © Exporting RestorePoint © All phases completed successfully.	HOTECTE DOBECT ARTEACTS devmysql 1 @ kanister POLICY 20 @ spec	START Today, 7:00am DURATION 1 min, 15 secs
COMPLETED Backup scheduled-w2tjv	Snapshotting Application Components Snapshotting Application configuration Snapshotting Workdoa dapps1 All phases completed successfully.	reorder:solution AnterActs dewmysql 1 Toucr 20 az-snapshot-policy	START Today, 6:59am Duktrion 41 secs

e. Integrating Kasten K10 with HSPC creates a Thin Image snapshot and splits the pairs. A clone volume 00:02:78 was created from snapshot volume 00:02:77. In Storage Navigator, confirm the pair status.

Copy Type: TI V										
TI History (Page.1)										
A Filter ON OFF										
Data and Time		Primary Volume Secondary Volume		olume	Mirror	De el ID		Description	Description	
Date and Time		LDEV ID	Provisioning Type	LDEV ID	Provisioning Type	Unit	POOLID	bin compare volume	Code	beschpton
2023/08/14 07:0	02:56	00:02:77	DP	00:02:78	DP	3	0	-	2092	CLONE END
2023/08/14 07:0	01:00	00:02:77	DP	00:02:78	DP	3	0	-	2091	CLONE START
2023/08/14 07:0	00:59	00:02:77	DP	00:02:78	DP	3	0	-	2001	PAIR
2023/08/14 07:0	00:14	00:02:44	DP	00:02:77	DP	3	0	-	2011	PSUS
2023/08/14 07:0	00:10	00:02:44	DP	00:02:77	DP	3	0	-	2001	PAIR

Restore Operation

You can restore an application from snapshot across clusters from the Kasten K10 Multi-Cluster. Restore operation consists of the following high-level steps:

- Copying the Import data.
- Creating a Restore policy.
- Implementing the Restore Policy.
- 1. To copy the Import data, from the K10 Multi-Cluster Dashboard, click the production Cluster "azurecluster" and then click **Policies**.
- 2. In the Polices window, select the Global Snapshot Policy created earlier, click **Show Import data**, and then click **Copy to clipboard**.

Policies		
olicies are used to automate your data management workflows. To ant to take (e.g., snapshot), a frequency or schedule for how often ibel-based selection criteria for the resources you want to manage	achieve this, they combine actions you you want to take that action, and a	
Create New Policy Filter by Name X	Importing Data	×
GLOBAL POLICY az-snapshot-policy Valid	The encoded text below contains import data needed by the receiving cluster. You'll be asked to paste this text when you create an import policy on the receiving cluster. Visit the Policies Page at any time to see this information.	revalidate 45 yemi 11 once
Snapshot on-demand for exporting data.	Copy to Clipboard	
B* Export onDemand snapshots using the export profile globalkastenk10b		

- 3. Create a restore policy.
 - a. From the K10 Global Resources page, click Polices and then click New Policy.
 - b. In the New Policy window, enter a Policy Name and select Import Frequency as On Demand.
 - c. In Config Data for Import section, paste the import policy copied in step 2.
 - d. Select Restore after Import and select the Global Location in Profile for Import.

e. Click Create Policy.

	New Policy	
Name		
The display name for this policy		
aws-restore-policy		
Comments		
		/
Action		
The action that should be taken whe	n this policy is executed	
Snapshot		Import
Restore After Import		
Automatically restore after imp	orting	
Data-Only Restore		
Restore only the volume data	and exclude other artifacts	auch as config files.
Don't wait for workloads	s to be ready	
Specifies whether the restore	action should skip waiting fo	er all workloads (Deployments,
statenuisets or Deploymentu	ontigs) to be ready before co	mpieting.
Restore cluster-scoped r	esources	
If the restore point contains or restored unless you select t	iuster-scoped (non-namespa this option. This helps preve	ced) resources, they will not be nt against unintended overwriting
of this cluster's resources.		
Apply transforms to re	stored resources	
On restore, change the cont	ents of spec resources. This m	ay be useful when migrating
names.	example, you can change stor	ige classes or edit container image
Select Application Resource	S	
optionary create inters to include	erencique specifieu applicatio	n resolutors.
All Resources		Filter Resources
Optional blueprint actions to be	run before or after restores c	omplete
Before		
After - On Success		
After - On Failure		
After - On Failure		
After - On Failure		
After - On Failure mport Frequency Hourly	Daily	O Weekly
Meter - On Failure mport Frequency Hourly Monthly	Daily	Weekly On Demand
Meter - On Failure mport Frequency Hourly Monthly	O Daily Yearly	Weekly On Demand
Meter - On Failure	Oaily Yearly	Weekly On Demand
After - On Failure mport Frequency Hourly Monthly Config Data for Import Patte the text that was presented to	Daily Vearly	Weekly On Demand
After - On Failure mport Frequency Hourly Monthly Config Data for Import Patte the text that was presented to future. Policy runs will synchronize th in text more	Daily Yearly	Weekly On Demand
After - On Failure mport Frequency Hourly Monthly Config Data for Import Patte the text that was presented to ducter. Policy runs will synchronize the last export.	O Daily Vearly	Weekly On Demand
After - On Failure Import Frequency Hourly Hourly Config Data for Import Tatte the text that was presented to ducter. Policy runs will synchronize the last export. bt:APpearticsychilexr12rigNeT1 GateAppeart	Daily Yearly you when the restore point whe restore points present in 1 FulleBlags I.C. 1211-134-91/1097-egat/C	Weekly On Demand source the source cluster at the time of Ppc/stukEU+e7st1h1+94-gbl0t-ks2ck
After - On Failure Import Frequency Hourly Hourly Config Data for Import State the text that was presented to Suster. Policy runs will synchronize the last export. Intel@poentics.chile.or	Daily Yearly you when the restore point whe restore points present in t FulleBlogs IO: 121111491/097-ggstX	Weekly On Demand or Demand was exported from the source the source cluster at the time of Pgc/stukEU+u7st1h7r98v@bUT-Ms2Dx
After - On Failure Import Frequency Hourly Hourly Config Data for Import State the text that was presented to Suster, Policy runs will synchronize the Inst export. IntroPoentCay XPREwrIDYBUFTI Galidat Frofile for Import Refer the profile that defines the loc	Daily Vearly vois when the restore point whe restore points present in t fulteBlock 10 to	Weekly On Demand was exported from the source the source cluster at the time of Pgc/stukEl+c7rt1hfr98v@k00T-Mc2Ce
After - On Failure Import Frequency Hourly Hourly Config Data for Import Tatte the text that was presented to Suster. Policy runs will synchronize the last export. IntroPresentary Officer Tright FTL Galaxie Frofile for Import iselect the profile that defines the loc D global kastenk10b	Daily Vearly vois when the restore point whe restore points present in t futBlogs IOL 131k1x+91/1097-egat/C ation for Importing data.	Weekly On Demand Ass exported from the source the source cluster at the time of rgc/stukEl+c7zE1hTr98xQkUUT-A22cc
After - On Failure Import Frequency Hourly Hourly Config Data for Import Tatte the text that was presented to fuctor. Policy runs will synchronize of the last export. IntroPresentary/OfficerID/gill/Ti Calkde Frofile for Import ielect the profile that defines the loc Dig globalkastenk10b	Daily Vearly you when the restore point whe restore points present in t fullettige i.o. 121kise91/097egetK attorn for Importing data.	Weekly On Demand
After - On Failure Import Frequency Hourly Hourly Config Data for Import Tate the text that was presented to Juder. Policy runs will synchronize the last export. InterPresentary:Officer-321(9017) Calibre Profile for Import Inter the profile that defines the loc Dig globalkastenk10b	Daily Vearly you when the restore point whe restore points present in t FutBlogs 10: 131 kike 9 LT d97 egot 0 ation for Importing data.	Weekly On Demand
After - On Failure Import Frequency Hourly Hourly Config Data for Import Patter the text that was presented to Juder. Policy runs will synchronize of the last export. InterPresentary:Officer=321(90171 Calibre Profile for Import Inter the profile that defines the loc College globalkastenk10b Advanced Settings	Daily Vearly you when the restore point whe restore points present in t heteBass Ide 131k3k=91/1097egat/0 ation for Importing data.	On Demand On Demand Ass exported from the source the source cluster at the time of App://stukEli-utr/shugeout-Aczox
After - On Failure Import Frequency Hourly Hourly Config Data for Import Patter the text that was presented to functer. Policy runs will synchronize of the last export. InterPresentage/Plicer/32/gBUT1 Calibre Profile for Import InterConfig Data technes the loc Config Data technes the loc Advanced Settings gnore Exceptions and Continue profiles barras retrineer	Daily Vearly Vearly voiu when the restore point whe restore points present in the restore points present in the restore points present in the restore points of the restore points attain for Importing data. e if Possible falling) is useful in environme	Weekly On Demand or Domand or operated from the source he source cluster at the time of or operative:
After - On Failure Import Frequency Hourly Hourly Config Data for Import Patter the text that was presented to functer. Policy runs will synchronize of the last export. InterPresentage/Placer321(90171 Calibre Profile for Import Inter the profile that defines the loc Complete Statement of the state of the last export. Advanced Settings gnore Exceptions and Continue proving exceptions for stars retrying/ reven state but the policy actions si	Daily Vearly Vearly Vearly vois when the restore points restore points present in t reterm porting data. e if Possible failing) is useful in environme tould continue best-effort.	Weekly On Demand or Demand or operated from the source he source cluster at the time of or operative stratistic strate operations or operative strate opera
After - On Failure Import Frequency Hourly Hourly Config Data for Import Tate the text that was presented to Suder. Policy runs will synchronize of the last export. InterPresentary:27/68/97164848 Frofile for Import Inter the profile that defines the loc D Advanced Settings spore Exceptions and Continue profile scattering for the policy actions of D Mobile scattering substances	Daily Vearly Vearly Vearly voiu when the restore points restore points present in t restore points present in t restore points present in t restore points re	Weekly On Demand or Demand

f. Verify that the policy is created.

🌐 GI	LOBAL POLICY	Not yet added to distribution.	① Create a Distribution
R	aws-restore-policy		[교] edit
	Import on-demand and restore after import us for exporting data.	sing the import profile globalkastenk10b.	>yami
			delete

- 4. Create a Distribution.
 - a. From the K10 Global Resources page, click Distributions and then click New Distribution.
 - b. In the Add Distribution window, enter the name, specify the AWS cluster, select the restore policy and location profile in Resources, and then click **Add Distribution**.

			ci ilo di		
Name The dis	e splay name for the distribution. M	lust be Kube	rnetes-co	mpatible (lowercase, dots, da	ishes)
dist-	-aws-restore-policy				
Clusto Using I Multip	ers labels, specify the clusters to whic le labels will be unioned (OR). Any	th you want t v cluster that	to distribi matches	Ite resources. any label will be targeted.	
Cli	uster - awscluster 🗙				
Resou	JICES	huto to cluste	0.55		
Select	the global KTO resources to distri	Salast All	ers.	ad (3)	Develop All
Availat		Select All	Select	ed (2)	Deselect All
M	aws-restore Import + restore policy depends on profile globalkastenk10b.	٥	W	aws-restore-policy Import + restore policy depend profile globalkastenk10b.	s on 🛛 🚫
R	az-snapshot-policy Snapshot + export policy depends or profile globalkastenk10b.	n 🕑	6	globalkastenk10b S3, us-west-1, "bucketk10b"	8
Q	demo10-restore Import + restore policy depends on profile globalkastenk10b.	٥			
R	demo10-snapshot Snapshot + export policy depends or profile globalkastenk10b.	n 📀			

c. Verify that the distribution is added.

ľ\$	distribution dist-aws-restore-policy	C Force sync yaml edit delete
	CLUSTERS RESOURCES	
	STATUS ✓ Synced 6 minutes ago	

- 5. Run the restore operation.
 - a. From the K10 Multi-Cluster Dashboard, click the secondary cluster "awscluster" and then click Policies.
 - b. Verify whether the Global restore on-demand policy created in the <u>Create Restore Policy</u> section is available under **Polices**.

R	GLOBAL POLICY	✓ revalidate
	Valid	 yaml
	Import on-demand and restore after import using the import profile globalkastenk10b. for exporting data.	्रे run once

c. Click run once, which opens the Run Once window. To start the restore, click Yes, continue.



d. Open the K10 Multi-Cluster Dashboard and check the status of the policy in the **Actions** window. To check the phase in progress, click the related action.



e. Verify that the phase has changed to Completed Successfully. To confirm that no error is present, click the relevant action.

KASTEN by Veeam			🗓 Docs 🛛 k10-multi-cluster-user 🗸 🗘 🗗
Clusters < aw	vscluster \vee		
\bigcirc	COMPLETED SUCCESSFULLY aws-restore-policy policy-run-srprs Show Details	START END DURATION Today, 7:22am Today, 7:25am 2 mins, 27 secs APPLICATIONS	
ctions 2			ំរុំ៖ Filter Actions 🗸 🗸
COMPLETED Restore scheduled-g6ljc	PHASES Restoring Application Components All phases completed successfully. 	TANDETHANESPACE ARTIFACTS devmysql none POLICY aws-restore-policy	START Today, 7:23am URATION 1 min, 3 Secs
COMPLETED	PHASES	PROTECTED OBJECT ARTIFACTS	START

f. From the K10 Multi-Cluster Dashboard, navigate to the cluster "awscluster" and verify that application "devmysql" is restored.

< Clusters < awscluster >	
Applications View details or perform actions on applications.	Cluster-Scoped Resources M Latest snapshot was Today, 6:59am Create a Policy >
1% Filter by Status devmysq(1 application Image: 1 to 1 application	
devmysql	
Not Protected by Policies	
Latest snapshot was Today. 7:00am	
∳ Create a Policy >	
100.0 GIB 🗍 - E1 🛞 t 🖏 8	
Snapshot restore export details	

g. From the OpenShift UI, verify that statefulset application azapps1 is created.

Project: devmysql 🔻								
StatefulSets > StatefulSet d	etails						1	Actions 🔻
Details Metrics	YAML Pods Environment	Events						
▼ Filter ▼ Name Name ↑	Search by name7 Status 1	Ready 1	Restarts 🗍	Node I	Memory 1	CPU 1	Created 1	
P azapps1-0	C Running	1/1	0	N ip-10-77-28-154.us-west- 1.compute.internal	187.7 MiB	0.000 cores	Aug 14, 2023, 7:24 AM	:

6. Verify that the data is available.

a. Log in to the pod azapps1-0 in devmysql namespace in the awscluster cluster and verify whether the 100 GB persistent volume is mounted.

ods > Pod details											
azapps1-	0 c	Running									
etails Metr	ics	YAML	Er	nvironmei	nt Log	js	Events	Termina	al		
onnecting to 🛛 🙆	azappsl										
sh-4.2\$ df -h Filesystem	Size	Used	Avail	Use% Moun	ted on						
sh-4.2\$ df -h Filesystem overlay	Size 100G	Used 30G	Avail 71G	Use% Moun 30% /	ted on						
sh-4.2\$ df -h Filesystem overlay tmpfs	Size 100G 64M	Used 30G 0	Avail 71G 64M	Use% Moun 30% / 0% /dev	ted on						
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs	Size 100G 64M 7.8G	Used 30G 0	Avail 71G 64M 7.8G	Use% Moun 30% / 0% /dev 0% /sys	ted on /fs/cgrou	p					
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm	Size 100G 64M 7.8G 64M	Used 30G 0 0	Avail 71G 64M 7.8G 64M	Use% Moun 30% / 0% /dev 0% /sys 0% /dev	ted on /fs/cgrou /shm	p					
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs	Size 100G 64M 7.8G 64M 7.8G	Used 30G 0 0 56M	Avail 71G 64M 7.8G 64M 7.7G	Use% Moun 30% / 0% /dev 0% /sys 0% /dev 1% /etc	ted on /fs/cgrou /shm /passwd	p					
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/nvme0n1p4	Size 100G 64M 7.8G 64M 7.8G 100G	Used 30G 0 0 56M 30G	Avail 71G 64M 7.8G 64M 7.7G 71G	Use% Moun 30% / 0% /dev 0% /sys 0% /dev 1% /etc 30% /etc	ted on /fs/cgrou /shm /passwd /hosts	p					
sh-4.2\$ df -h Filesystem overlay tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde	Size 100G 64M 7.8G 64M 7.8G 100G 98G	Used 30G 0 0 56M 30G 271M	Avail 71G 64M 7.8G 64M 7.7G 71G 93G	Use% Moun 30% / 0% /dev 0% /sys 0% /dev 1% /etc 30% /etc 1% /var	fs/cgrou /fs/cgrou /shm /passwd /hosts /lib/mysq	p 1					
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde tmpfs	Size 100G 64M 7.8G 64M 7.8G 100G 98G 15G	Used 30G 0 0 56M 30G 271M 20K	Avail 71G 64M 7.8G 64M 7.7G 71G 93G 15G	Use% Moun 30% / 0% /dev 0% /sys 0% /dev 1% /etc 30% /etc 1% /van 1% /run	fed on /fs/cgrou /shm /passwd /hosts /lib/mysq /secrets/	p 1 kuberr	netes.io/:	serviceacc	ount		
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde tmpfs tmpfs	Size 100G 64M 7.8G 64M 7.8G 100G 98G 15G 7.8G	Used 30G 0 0 56M 30G 271M 20K 0	Avail 71G 64M 7.8G 64M 7.7G 71G 93G 15G 7.8G	Use% Moun 30% / 0% /dev 0% /sys 0% /dev 1% /etc 30% /etc 1% /var 1% /var 1% /pro	/fs/cgrou /shm /passwd /hosts /lib/mysq /secrets/ z/acpi	p 1 kuberr	netes.io/	serviceacc	ount		
sh-4.2\$ df -h Filesystem overlay tmpfs tmpfs shm tmpfs /dev/nvme0n1p4 /dev/sde tmpfs tmpfs tmpfs	Size 100G 64M 7.8G 64M 7.8G 100G 98G 15G 7.8G 7.8G	Used 30G 0 6 56M 30G 271M 20K 0 0	Avail 71G 64M 7.8G 64M 7.7G 71G 93G 15G 7.8G 7.8G	Use% Moun 30% / 0% /dev 0% /sys 0% /dev 1% /etc 30% /etc 1% /var 1% /run 0% /pro 0% /pro	fs/cgrou /shm /passwd /hosts /lib/mysq /secrets/ z/acpi z/scsi	p 1 kuberr	netes.io/	serviceacc	ount		

b. Log in to MySQL and verify whether database devmysqldb and employee table is available.

Project: devm	ysql 🔻						
Pods > Pod det	tails						
🕑 azapp:	s1-0 ເວ	Running					
Details M	letrics	YAML	Environment	Logs	Events	Terminal	
Connecting to	C azapps	51					
mysql≻ mysql≻ show	database	s; -+					
Database		Ì					
informatic	on_schema						
mysql	co cchoma						
sys	ce_scheila						
5 rows in se	et (0.0 0	sec)					
mysql> use o Reading tab You can turn	devmysqld le inform n off thi	b; ation for s feature	completion of ta to get a quicker	ble and o startup	column names with -A		
Database cha mvsal≻ show	anged tables:						
+	devmvsal	+ db					
+		+					
+	+ (0.00 =	+ ec)					
mysql>	- (0100 3						

c. Verify whether the ingested data in Azure cluster "azurecluster" (as shown in <u>ARO cluster</u> section) is available in "awscluster" cluster in AWS.

ecting to C azapps1	
ql>	
<pre>ql> select * from employee;</pre>	
d name email	
+	
1 ankit ankit@rxy.com	
2 Raj raj@rxy.com	

The following screenshots show the status of PVC and PV created during restoration.

Persistent volume claim:				
Project: devmysql 🔻				
PersistentVolumeClaims				Create PersistentVolumeCla
Filter Name Search by name 7				
Name 1 Status 1	PersistentVolumes 1	Capacity 1	Used 1	StorageClass 1
PVC dev-vol-azapps1-0 Sound	PV pvc-986b06e0-4599-4169- bf89-ef6fff7534d5	100 GiB	-	SC sc-vsp5200

Persistent volume:

PersistentVolumes > PersistentVolume details								
PV pvc-986b06e0-4599-4169-bf89-ef6fff7534d5 @ Bound								
Details YAML								
1 kind: PersistentVolume								
2 apiVersion: v1								
3 metadata:								
4 name: pvc-986b06e0-4599-4169-bf89-ef6fff7534d5								
5 uid: ba935ad3-406c-44a1-bb53-816f8b9c77cc								
6 resourceVersion: '15879800'								
7 creationimestamp: 2023-08-14107:24:032*								
8 annotations:								
9 pv.kubernetes.io/provisioned-by: nspc.csi.nitacni.com								
10 volume.kubernetes.io/provisioner-deletion-secret-name: secret-vsp3200								
12 finalizance								
13 - kubernetes in/nv-protection								
14 - external-attacher/hspc-csi-hitachi-com								
15 > managedFields: ···								
82 spec:								
83 capacity:								
84 storage: 100Gi								
85 csi:								
86 driver: hspc.csi.hitachi.com								
87 volumeHandle: 01scsi900000040028632spc-f0a529d756								
88 fsType: ext4								
89 volumeAttributes:								
90 nostModeUption:								
91 5126: 10051								
93 nickname: snc_f0a529d756								
94 norts: Cl1-C								
95 ldevIDHex: '02:78'								
96 connectionType: iscsi								

The volume 00:02:78 was assigned to the restored application in AWS.

pc-32309/d0f2d3f1a4d8d807b1fdf32 (0C)													
<u>59-520-5V10[5/N:40028]</u> > <u>Porty/Host Groups/iSCSI Tarcets</u> > <u>CL1-C</u> > spc-32309d0f2d3f1s4d84807b1fdf32													
Volume Migration *													
ISCSI Target Alias	spc-32309d0f2d3f1a4d	spc-32309d0f2d3f1a4d8d807b1fdf32 (0C)					Host Mode				00 [Standard]		
ISCSI Target Name	ign.1994-04.jp.co.hitach	n.1994-04.jp.co.hitachi:rsd.r90.t.40028.1c00c									Enabled		
Port ID	CL1-C	2L1-C						Method			Comply with Host Settin	9	
Virtual Storage Machine	VSP 5200, 5600 / 40028 Mutual CHAP							Disabled					
								User Name					
Hosts LUNs Host Mode Options CHAP U	sers												
Add LUN Paths Copy LUN Paths Edit Comman	nd Devices More Action	is 🔻										S	elected: 0
\$Filter ON OFF Select All Pages Column Set	ttings										c	ptions 🗸 候 🗧	/ 1
Port ID LUN 1 LDEV ID LDEV	/ Pool Name	Emulation	Capacity				Used Capacit	γ		Capacity	aving	Capacity Saving	Provisionin
ID IN LOEVID Nam	e (ID)	Type	Total	Reserved	Used	Used (%)	Tier 1	Tier 2	Tier 3	capacity a		Status	Type
CL1-C 223 00:02:78 spc-	f0a dr_pool(0)	OPEN-V CVS	100.00 GB	0.00 GB	5.12 GB	5	-	-	-	Disabled		Disabled	DP

Test 5: Recover from a Ransomware Attack

This test case demonstrates how a VSP snapshot combined with immutability feature from Data Retention Utility program product can be used to recover a stateful application affected by a ransomware attack. For this test case, volume snapshot of the persistent volume used in the application in Azure Red Hat OpenShift cluster has already been taken and the DRU writedisable attribute is set on the snapshot volume.

Assume that the application is affected by a ransomware attack and we must restore clean data from the snapshot. This recovery process can be carried out either in Azure or in AWS.

Recovering from a ransomware attack consists of the following high-level steps:

- Creating PVC with the snapshot volume (where the DRU attribute Write Disabled is set).
- Creating a cascaded snapshot of this volume because write is disabled.
- Using the cascaded snapshot (snap-on-snap) to recover the application data in any cluster.
- Creating a clone PVC and using that PVC as data volume to restore the MySQL application because snapshot volumes
 must not be directly used in a POD.
- Verifying that the data ingested from the ARO cluster is available.

Snapshot Operation

Complete the following steps in Azure Red Hat OpenShift Cluster:

- 1. Create a new project "druapps" and deploy a stateful MySQL application with a persistent volume of 200 GB from the VSP 5200 storage system, as shown in the section <u>Test 2: Deploy a Stateful Application in Azure Red Hat OpenShift Cluster</u>.
- 2. Access the stateful MySQL application.
 - a. Log in to the pod azapps3-0 and verify whether the 200 GB persistent volume is mounted on "/var/lib/mysql".



b. Log in to MySQL and verify whether the database "devmysqldb" is created.

c. Create a table "employee" and ingest new records to the table.

roas > Poo	d details					
🕑 azap	ops3-0	C Running				
Details	Metrics	YAML	Environment	Logs	Events	Terminal
connecting	to 🙂 azapp	53				
mysql> mysal> s	how databas	es:				
+		+				
+	se 	+				
informa	ation_schem	a				
devmys	qldb					
nerfor	mance schem	al				
sys						
+		+				
5 rows 1	n set (0.00	sec)				
mysql≻ s	how tables;					
+		+				
Tables	_1n_devmysq	Idb				
employ	ee	i				
+		+				
1 row in	set (0.00	sec)				
mvsal> s	elect * fro	m employee	:			
++		+	+			
id	name	emai	1			
1 101	Prahin Ran	man nhan	mansadd com			
102	John Thang	li jtha	ngli@add.com			
103	Juliana K	jk@a	dd.com			
	Firoz Akht	ar fakh	tan@add.com			
104		nnod	dugadd com			
104 105	N Reddy	1 nreu	uyeauu.com			

- 3. Create a snapshot.
 - a. Create a snapshot of the PVC used in the MySQL application. From the Red Hat OpenShift console, navigate to **Storage**, click **VolumeSnapshots**, and then click **Create VolumeSnapshot**. In the Create VolumeSnapshot menu, enter the required information such as PVC, snapshot name, and snapshot class, and then click **Create**.

		PersistentVolumeClaim details
Create VolumeSnapshot	Edit YAML	Name
PersistentVolumeClaim *		PVC dev-vol-azapps3-0
PVC dev-vol-azapps3-0	•	Namespace
lame *		
dev-vol-azapps3-0-snapshot		Status Bound
		- Dound
inapshot Class *		StorageClass
VSC snapshotclass-sample	•	SC sc-vsp5200
		Requested capacity
Create Cancel		200 GiB
		Access mode
		Single user (RWO)
		Volume mode
		Filesystem

b. Verify whether the snapshot is created.

Project: druapps 🔻							
VolumeSnapshots							Create VolumeSnapshot
▼ Filter • Name • Search	h by name						
Name I	Status I	Size I	Source I	Snapshot content I	VolumeSnapshotClass 1	Created at I	
Visite dev-vol-azapps3-0-snapshot	Ready	200 GiB	even dev-vol-azapps3-0	view snapcontent-f829/c23-4d50-4ae2- b9d4-5dd7dbbd7223	vso snapshotclass-sample	Aug 21, 2023, 11:39 AM	I

c. In Storage Navigator, verify whether the snapshot volume 00:00:D2 is created successfully.

History	istory								
Сору Туре: ТІ 🔍									
TI History (Page.1)			_	_	_				
SFilter ON OFF									
Data and Time	Primary Volu	ume	Secondary Volume		Mirror	Bool ID	Diff. Compare Volume	Description	Description
Date and Time	LDEV ID	Provisioning Type	LDEV ID	Provisioning Type	Unit	POOLID	bin compare voidine	Code	Desciption
2023/08/21 06:10:11	00:00:D1	DP	00:00:D2	DP	3	0	-	2011	PSUS
2023/08/21 06:10:08	00:00:D1	DP	00:00:D2	DP	3	0	-	2001	PAIR

4. Set DRU write-disabled attribute to snapshot volume 00:00:D2 (LDEV# 210 in decimal format).

```
[root@linuxnfscl2 etc]# raidvchkdsp -g grp0 -fd -v gflag -I1
                                  Seq# LDEV# GI-C-R-W-S PI-C-R-W-S
Group
        PairVol Device File
                                                                      R-Time
                                          210 E E E E E
grp0
       pair0
                 Unknown
                                 540028
                                                           EEEEE
                                                                            0
[root@linuxnfscl2 etc]#
[root@linuxnfscl2 etc] # raidvchkset -g grp0 -vg wtd 5 -I1
[root@linuxnfscl2 etc] # raidvchkdsp -g grp0 -fd -v gflag -I1
                                  Seq# LDEV# GI-C-R-W-S
               Device File
Group
        PairVol
                                                         PI-C-R-W-
                                                                      R-Time
                                                                   5
grp0
       pair0
                 Unknown
                                 540028
                                          210
                                               EEEDE
                                                           ΕE
                                                                  D
                                                                   E
                                                                            5
                                                               F
[root@linuxnfscl2 etc]#
```

Restore Operation

This section shows the restoration procedure when an application in the ARO cluster is affected by ransomware. Restore operation is performed in Red Hat OpenShift cluster in AWS.

- 1. Create a project "druapps" in OpenShift cluster in AWS.
- 2. Create a PV for snapshot volume 00:00:D2.
 - a. Identify the volume handle string for the snapshot volume 00:00:D2. The volume handle string for this volume is "60060e80089c5c0000509c5c0000**00d2--spc-38105307cc**".
 - b. In the string, the volume ID is "00d2" and the volume name is "spc-38105307cc". The volume name is automatically assigned by HSPC.
 - c. Create a manifest file for PV using the volume handle string. This ensures that storage class does not dynamically create a new volume; instead, it uses the existing volume for preserving the snapshot data.
 - d. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumes**, and then click **Create PersistentVolume**. In the Create PersistantVolume menu, populate the YAML file with the required information and click **Create**.

Create I	te PersistentVolume y manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.
1	apiversion: vi
2	kind: PersistentVolume
3	metadata:
4	name: drusnapshotpv
5	namespace: druapps
6	spec:
7	capacity:
8	storage: 200Gi
9	accessModes:
10	- ReadWriteOnce
11	persistentVolumeReclaimPolicy: Retain
12	storageClassName: sc-vsp5200
13	
14	driver: hspc.csi.hitachi.com
15	volumeHandle: 60060e80089c5c0000509c5c00000d2spc-38105307cc
16	claimRef:
17	name: drusnapshotpvc
18	namespace: druapps
19	

e. Verify whether the PV is created as per manifest.

PersistentVolumes						Create PersistentVolume
Name 👻 Search by name /						
Name 1	Status 🚦	Claim 1	Capacity 1	Labels 1	Created ↓	
🕑 drusnapshotpv	Available	(PVC) drusnapshotpvc	200Gi	No labels	Aug 21, 2023, 7:56 AM	I

f. Create a manifest file for PVC using the PV which you created. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumeClaims**, and then click **Create PersistentVolumeClaim**. In the Create PersistantVolumeClaim menu, populate the YAML file with the required information and click **Create**.



g. Verify whether the PVC is created as per manifest.

Project: druapps 🔹						
PersistentVolumeCla	ims					Create PersistentVolumeClaim
▼ Filter ▼ Name ▼ S	▼ Filter Name Search by name_ [/]					
Name 1	Status 💲	PersistentVolumes 1	Capacity 1	Used 💲	StorageClass 🗍	
(PVC) drusnapshotpvc	Sound 8	ev drusnapshotpv	200 GiB		SC sc-vsp5200	1

- 3. Create a cascaded snapshot from the volume 00:00:D2.
 - a. From the Red Hat OpenShift console, navigate to Storage, click VolumeSnapshots, and then click Create VolumeSnapshot. In the Create VolumeSnapshot menu, enter the required information such as PVC, Snapshot Name, and Snapshot Class, and click Create. Select the PVC you created in step 2.

Project: druapps 🔹		
		PersistentVolumeClaim details
Create VolumeSnapshot	Edit YAML	Name
		PVC drusnapshotpvc
PersistentVolumeClaim *		
(PVC) drusnapshotpvc	•	Namespace NB druanos
Name *		and approximation of the second
Name		Status
drusnaponsnap		Bound
Snapshot Class *		StorageClass
VSC snapshotclass-sample	•	Se sc-vsp5200
		Requested capacity
Create		200 GiB
		Access mode
		Single user (RWO)
		Volume mode
		Filesystem

b. Verify whether the cascaded volumesnapshot is created.

Project: druapps 🔹	Project druapps • VolumeSnapshots									
VolumeSnapshots							Create VolumeSnapshot			
▼ Filter ▼ Name ▼	Search by name									
Name I	Status 1	Size 1	Source 1	Snapshot content 🕴	VolumeSnapshotClass	Created at 💲				
🕼 drusnaponsnap	👁 Ready	200 GiB	(PVC) drusnapshotpvc	VSC snapcontent-18c08c06-5721-44e1- 9cce-cd8737c828f1	VSC snapshotclass-sample	Aug 21, 2023, 8:02 AM	1			

c. In Storage Navigator, verify whether the cascaded volume 00:00:D3 is created successfully.

H	listory										
	Date and Time Primary Volume Secondary Volume Mirror Linit Pool ID Diff Compare Volume Description Code Description Descript										
	TI History (Page.1)	_		_			_		_		
	☆Filter ON OFF										
	Primary	Primary Volu	Primary Volume Secondary V		olume	Mirror	De al ID	Diff Come on Maluera	Description	Description	
	Date and nine	LDEV ID	Provisioning Type	LDEV ID	Provisioning Type	Unit	P0011D	bin compare volume	Code	Description	
	2023/08/21 08:03:02	00:00:D2	DP	00:00:D3	DP	3	0	-	2011	PSUS	
Date and T 2023/08/2: 2023/08/2:	2023/08/21 08:03:01	00:00:D2	DP	00:00:D3	DP	3	0	-	2001	PAIR	

The DRU attribute is not set on the new cascaded snapshot volume 00:00:D3.

```
[root@linuxnfscl2 etc]# raidvchkdsp -g grpl -fd -v gflag -I1
Group PairVol Device_File Seq# LDEV# GI-C-R-W-S PI-C-R-W-S R-Time
grpl pair0 Unknown 540028 211 E E E E E E E E E 0
[root@linuxnfscl2 etc]#
```

- 4. Use the snap-on-snap volume to create a clone volume to restore the application.
 - a. Create a project "drusnaponsnap".
 - b. Create a manifest file for PV using the volume handle string of snap-on-snap volume 00:00:D3. This way, storage class does not dynamically create a new volume; instead, it uses the existing volume for preserving the snapshot data.
 - c. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumes**, and then click **Create PersistentVolume**. In the Create PersistantVolume menu, populate the YAML file with the required information and click **Create**.



d. Verify whether the PV is created as per manifest.

2						
PersistentVolumes						Create PersistentVolume
Name V Search by name_ /						
Name 1	Status 1	Claim 1	Capacity 1	Labels 🗄	Created 1	
ev drusnaponsnappv	Available	eve drusnaponsnappvc	200Gi	No labels	Aug 21, 2023, 9:20 AM	1

e. Create a manifest file for PVC using the PV created in step 4. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumeClaims**, and then click **Create PersistentVolumeClaim**. In the Create PersistantVolume menu, populate the YAML file with the required information and click **Create**.

Project: o	Project: drusnaponsnap 🔻									
Creat	Create PersistentVolumeClaim									
Create by	Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.									
1	apiVersion: v1									
2	kind: PersistentVolumeClaim									
3	metadata:									
4	name: drusnaponsnappvc									
5	namespace: drusnaponsnap									
6	spec:									
7	accessModes:									
8	- ReadWriteOnce									
9	resources:									
10	requests:									
11	storage: 200Gi									
12	volumeName: drusnaponsnappv									
13	storageClassName: sc-vsp5200									

f. Verify whether the PVC is created as per manifest.

(PVC) drusnaponsnappvc	Bound	ev drusnaponsnappv	200 GiB	-	SC sc-vsp5200	ŧ				
Name 1	Status 🗍	PersistentVolumes 1	Capacity 1	Used 1	StorageClass 1					
▼ Filter ▼ Name ▼ Search	▼ Filter ▼ Name ▼ Search by name 7									
PersistentVolumeClaims	;					Create PersistentVolumeClaim				
Project: drusnaponsnap 👻	voject drusnaponsnap 🔹									

g. Create a clone PVC using the snapshot PVC created in step 4e as dataSource. From the Red Hat OpenShift console, navigate to **Storage**, click **PersistantVolumeClaims**, and then click **Create PersistentVolumeClaim**. In the Create PersistantVolumeClaim menu, populate the YAML file with the required information and click **Create**.

Project	: drusnaponsnap 👻					
Create b	te PersistentVolumeClaim ay manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.					
1	apiversion: v1					
2	kind Persistentvolumetidim					
2						
-	name: drusnaponsnapcione					
2	themespace. In using pointing					
7	spec.					
8	dataSource:					
9						
10	kind: PersistentvolumeClaim					
11	apiGroup: ""					
12	accessModes:					
13	- ReadWriteOnce					
14	resources:					
15	requests:					
16	storage: 200Gi					

h. Verify whether the clone PVC is created.

Project: drusnaponsnap 💌						
PersistentVolumeClaims						Create PersistentVolumeClaim
▼ Filter ▼ Name ▼ Search by na	me					
Name 1	Status 1	PersistentVolumes 1	Capacity 1	Used 1	StorageClass 1	
(EVC) drusnaponsnapcione	🕲 Bound	pvc-6076ee41-dcc4-44bf-bef6- 5426flaf05c1	200 GiB	-	SC sc-vsp5200	ŧ
(PVC) drusnaponsnappvc	Sound 8	(20) drusnaponsnappv	200 GiB	-	SC sc-vsp5200	:

i. Creating the PVC also creates a persistent volume. Verify the YAML of this persistent volume. In this scenario, the volume ID is 00:D6.

PersistentVol	PersistentVolume > PersistentVolume details								
PV pvo	-6076ee41-dcc4-44bf-bef6-5426f1af05c1 Bound Bound Bound Bound Bound Bound								
Details	YAML								
1 ki	nd: PersistentVolume								
2 ap	iversion: v1								
3 me	tadata:								
4	name: pvc-6076ee41-dcc4-44bf-bef6-5426f1af05c1								
5	uld: d6f9321c-cc69-4e26-96a2-ef27c762c423								
6	resourceversion: 20623136								
2	Creditoriimestamp: 2023-08-21109:27:272								
9	ambutations.								
10	volume.kubernets.io/provisioner-deletion-secret-name: secret-vsp5200								
11	volume.kubernetes.io/provisioner-deletion-secret-namespace: default								
12	finalizers:								
13	- kubernetes.io/pv-protection								
14 >	managedFields:								
72 sp	ec:								
73	capacity:								
74	storage: 200Gi								
75	csi:								
76	driver: hspc.csi.hitachi.com								
77	VolumeHandle: 01SCS190000040028214SpC-C1C0707a4e								
78	TSIVE: EXL4								
80	host MadeDation: ''								
81	size 2005								
82	portrs: "								
83	nickname: spc-cicd707a4e								
84	ports: CL1-C								
85	ldevIDHex: '00:D6'								
86	connectionType: iscsi								
87	storage.kubernetes.io/csiProvisionerIdentity: 1689941886921-8081-hspc.csi.hitachi.com								
88	ldevIDDec: '214'								

j. In Storage Navigator, verify whether the clone volume 00:00:D6 is created successfully.

nisu	ory											
Cop	ру Туре: ТІ										€ € 1	/ 1
TI History (Page.1)												
A Filter ON OFF												
D)ate and Time	Primary Volu	ime	Secondary V	olume	Mirror	Real ID	Diff Compare Volume	Description	Description		
Dat	Jace and Time	LDEV ID	Provisioning Type	LDEV ID	Provisioning Type	Unit	1 001 10	bin compare volume	Code	Description		
2	2023/08/21 09:31:18	00:00:D3	DP	00:00:D6	DP	3	0	-	2092	CLONE END		
2	2023/08/21 09:27:34	00:00:D3	DP	00:00:D6	DP	3	0	-	2091	CLONE START		
2	2023/08/21 09:27:33	00:00:D3	DP	00:00:D6	DP	3	0	-	2001	PAIR		

- 5. Restore the MySQL application in AWS using the clone PVC.
 - a. Create a MySQL service. From the Red Hat OpenShift console, navigate to **Networking**, click **Services**, and then click **Create Service**. In the Create Service menu, populate the YAML file with the required information and click **Create**.



b. Verify whether the MySQL service is created.

Project: drusnaponsnap 👻				
Services				Create Service
Name Search by name /				
Name I	Labels I	Pod selector I	Location I	
S azapps3	(app=azapps3)	Q app=azapps3	None	I

c. Create a MySQL statefulset application. From the Red Hat OpenShift console, navigate to **Workloads**, click **StatefulSets**, and then click **Create StatefulSet**. In the Create StatefulSet menu, populate the YAML file with the required information and click **Create**. In the volume section, use the claim "drusnaponsnapclone" which you created in step 4g. This ensures that the MySQL application uses the clone PVC for persistent data.

Project:	drusnaponsnap 👻
Creat	e StatefulSet
Create by	manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.
1	apiversion: apps/v1
2	kind: StatefulSet
З	metadata:
4	namespace: drusnaponsnap
5	name: azapps3
6	spec:
7	selector:
8	matchLadels:
9	app: azapps3
10	scivilerame. a200053
12	replicas: 1
13	template:
14	metadata:
15	labels:
16	app: azapps3
17	spec:
18	terminationGracePeriodSeconds: 30
19	containers:
20	- name: azapps3
21	image: mysql:5.7
22	args:
23	- "ignore-db-dir=lost+found"
24	env:
25	- name: MYSQL_ROOT_PASSWORD
20	Value: passizs
27	value: devmvsoldb
29	- name: MYSQL USER
30	value: admin
31	- name: MYSQL_PASSWORD
32	value: secret
33	ports:
34	- containerPort: 3306
35	name: mysql
36	volumeMounts:
37	- name: dev-vol
38	mountPath: /var/lib/mysql
39	volumes:
40	- Name: dev-voi
41	persistentvolumetidim:
43	

d. Verify whether the statefulset is running.

StatefulCate) StatefulCat dataile								
stateduser of stateduser details								Actions 👻
Details Metrics YAML	Pods Environment Events							
▼ Filter ▼ Name ▼ S	Search by name							
▼ Filter ・ Name ・ S Name †	Status	Ready 1	Restarts 1	Node I	Memory I	CPU I	Created	

e. Log in to the pod azapps3-0 and verify whether the 200 GB persistent volume is mounted on "/var/lib/mysql".

Project: drusnapo	nsnap ·	•						
Pods > Pod details								
P azapps3-	-0 C R	tunning						
Details Metri	ics Y	AML	Envir	ronment	Logs	Events	Terminal	
Connecting to 🧿	azapps3							
sh-4.2\$ df -h	Size	lised	Avail	Lice% Mo	unted on			
overlay	100G	29G	726	29% /				
tmpfs	64M	0	64M	0% /d	ev			
tmpfs	7.8G	0	7.8G	0% /s	ys/fs/cgro	oup		
shm	64M	0	64M	0% /d	ev/shm			
tmpfs	7.8G	56M	7.7G	1% /e	tc/passwd			
/dev/nvme0n1p4	4 100G	29G	72G	29% /e	tc/hosts			
/dev/sdf	196G	271M	186G	1% /v	ar/lib/mys	sql		
tmpfs	15G	20K	15G	1% /ri	un/secrets	s/kubernete	es.io/serviceaccount	
tmpfs	7.8G	0	7.8G	0% /p	roc/acpi			
tmpfs	7.8G	0	7.8G	0% /p	roc/scsi			
tmpfs	7.8G	0	7.8G	0% /s	ys/firmwa	re		
sh-4.2\$								

f. Log in to MySQL and verify whether the database "devmysqldb" is available.

Project: drusnaponsnap 🔻	
Pods > Pod details	
Details Metrics YAML Environment Logs Events	Terminal
Connecting to (g) azapps3 mysql> mysql> show databases;	
Database information_schema devmysqldb	
my544 performance_schema sys +	
mysql> use devmysqldb; Reading table information for completion of table and column You can turn off this feature to get a quicker startup with	names -A
Database changed mysql> show tables; t	
employee + 1 row in set (0.00 sec)	

g. Verify whether the ingested data from Azure RedHat OpenShift cluster is available here.



h. Verify whether HSPC automatically creates an iSCSI target on port CL1-C and mounts the clone volume 00:00:D6.

spc-32309d0f2d3f1a4d8d807b1fdf32 (0C)										Last Up	dated : 2023/08/21 09	
VSP-5200-SV10(S/N:40028) > Ports/Host Groups/iSCS)	I Targets > CL1-C > spc-32309d0f	2d3f1a4d8d807b	1fdf32									
Volume Migration												
iSCSI Target Alias spc-32309d0f2d3f1a4d8d807b1fdf32 (0C)				Host Mode				00 [Standard]				
iSCSI Target Name	iqn.1994-04.jp.co.hitachi:rsd.r90.t.40028.1c00c				Port Securi	Port Security				Enabled		
Port ID	CL1-C				Authentica	hentication Method			Comply with Host Setting			
Virtual Storage Machine	VSP 5200, 5600 / 40028				Mutual CHAP			Disabled				
							Use	er Name				
Hosts LUNS Host Mode Options CHAP Us	sers				^							
Add LUN Paths Copy LUN Paths Edit Command Devices More Actions 💌												
& Filter OI OFF Select All Pages Column settings												
	Pool Name	Emulation	Capacity U			Used Capac	ity			Capacity Saving	Pro	
LON ID LOEV ID LOEV	(ID)	Туре	Total	Reserved	Used 1▼	Used (%)	Tier 1	Tier 2	Tier 3	Capacity Saving	Status	Тур
CL1-C Ø 95 00:00:D6 spc-0	c1cd707a4e dr_pool(0)	OPEN-V CVS	200.00 GB	0.00 GB	9.31 GB	4	-		-	Disabled	Disabled	DP

i. Delete the snap-on-snap PV (drusnaponsnappv) and PVC (drusnaponsnappvc) created in step 4c and step 4e.