Deploy Horizon VDI on Hitachi Virtual Storage Platform One SDS Block

Reference Architecture Guide

MK-SL-314-00
December 2024

Feedback

Please send comments to doc.feedback@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Revision history

Changes	Date
Initial release	December 2024

Reference Architecture Guide

This document describes the reference architecture for the implementation of Horizon VDI on Hitachi Virtual Storage Platform One SDS Block (VSP One SDS Block) storage. It combines VSP One SDS Block, Horizon VDI, Hitachi Advanced Server and VMware Hypervisor technology to enable an enterprise level solution for Virtual Desktop Infrastructure.

VSP One SDS Block is a comprehensive software-defined storage solution that unifies block, file, and object storage, providing a single data plane across on-premises and cloud environments. It offers flexibility and scalability, ensuring 100% data availability for hybrid cloud deployments (using Hitachi Vantara hardware). This integrated approach enables enterprises to efficiently manage their data across diverse infrastructures while optimizing performance and reliability.

Extensive testing was executed to evaluate the performance and capabilities of VSP One SDS Block, although the VSP One SDS Block cluster used in this guide was configured with the minimum number of storage nodes. VSP One SDS Block can scale out the number of nodes in the storage cluster, expanding the amount of I/O the whole storage cluster can process.

This paper presents the results of a series of tests conducted using Login VSI Enterprise, the industry standard benchmarking tool for VDI workloads. For these tests, Knowledge Worker user workloads was selected.

Note: These practices were developed in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

Product features

The following information describes the hardware and software features used in this reference architecture.

VSP One SDS Block

Hitachi Virtual Storage Platform One SDS Block (VSP One SDS Block) is a software-defined storage solution designed to manage both traditional and modern block-based workloads. Here are some key features and benefits:

- Unified Storage Platform

  Part of the larger Hitachi Virtual Storage Platform One Block platform, the VSP One SDS Block solution Integrates block, file, and object storage into a single solution. VSP One SDS Block simplifies data management across different storage types, reducing complexity and operational overhead.

- Single Data Plane

  This solution provides a consistent data management layer across on-premises and cloud. VSP One Block enables seamless data mobility and consistent management across hybrid cloud deployments.

- Software-Defined Architecture

  Runs on certified industry-standard hardware and supports various deployment options including on-prem and in the AWS cloud. If offers flexibility in hardware choices and deployment locations, reduces costs, and increases flexibility.

- Scalability

  Supports growth from terabytes to petabytes with independent scaling of performance and capacity. VSP One Block allows organizations to easily adapt to growing data needs without major infrastructure changes.

- Cloud Integration

  Available in AWS Marketplace and supports hybrid cloud deployments. VSP One Block facilitates cloud adoption and hybrid cloud strategies while leveraging existing on-prem investments.

- Automated Management

  The SDS solution from Hitachi includes support for ansible playbooks and AI-powered management tools. This capability not only simplifies operations and reduces manual management tasks but provides support for use cases including application development and test.

- Asynchronous Replication

  This feature enables data replication across multiple sites or availability zones without impacting primary storage performance. VSP One SDS Block with Async Replication allows application dev and test teams to use real data without impacting the current processes or applications.

- Data Reduction

  VSP One SDS Block offers 2:1 data reduction for block and file storage which improves storage efficiency and reduces storage costs.

- Enhanced Security

  Includes FIPS 140-3 level 1 data-at-rest encryption which protects sensitive data, shrinks the attack surface and helps meet compliance requirements.

This platform is designed to support a wide range of applications and use cases, from legacy databases to cloud-native applications, making it a versatile choice for modern enterprises.

See Virtual Storage Platform One, Software-Defined Block Storage - Ebook (hitachivantara.com) for details.

Horizon VDI

Virtual Desktop Infrastructure (VDI) is a cutting-edge technology that utilizes virtual machines to manage and provide virtual desktops. Rather than being confined to a specific physical device, a virtual desktop is a preconfigured image of an operating system, and its applications can be accessed from any compatible device. With VDI, desktop environments can be hosted on a centralized server and deployed to end-users on demand.

Within this lab validation report, we have leveraged VMware Horizon 8 connection server backed by core infrastructure which includes VMware vCenter 8 as well as Microsoft SQL database server and Microsoft active directory services to support VDI. Optionally, end users have the ability to use published applications, but in the Hitachi lab, validation was done using instant clones. In summary, VDI allows organizations to deploy Horizon VDI desktops that can be accessed from anywhere, at any time, making it an ideal solution for remote and on-prem users.

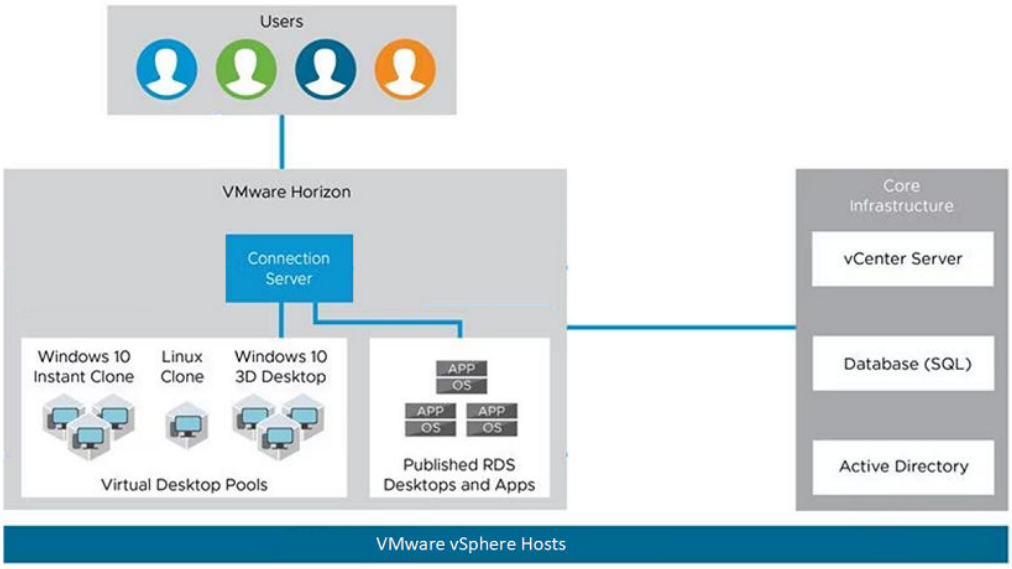

Connection Server

The Horizon Connection Server is a crucial component of VMware Horizon, acting as a broker for client connections. It authenticates users through Windows Active Directory, manages remote desktop and application sessions, and sets policies. It ensures secure connections between users and resources, supporting both internal and external connections.

Horizon Agent

Horizon Agent is an essential component that gets installed on virtual desktop and facilitates communication and monitoring.

Horizon Client

The VMware Horizon Client allows secure access to virtual desktops and applications from any device. It is compatible with multiple platforms, offers enhanced performance, and has a user-friendly interface.

SQL Database

SQL Database is installed as a separate component to maintain event DB configuration.

VMware OS Optimization Tool

The Windows OS Optimization Tool for VMware Horizon optimizes golden images by automating the removal of unnecessary settings and disabling unnecessary features typically present in a physical Windows machines to improve VDI performance and user experience. The following image shows the process of using the VMware optimization tool.

Within the context of this report, the VMware optimization tool was used to achieve better login VSI scores as outlined in the following sections.

Networking

Cisco Nexus 9000 Series switches are designed to tune datacenter networks into modern, automated environments that deliver exceptional application performance and effective IT operations. Cisco Nexus 9000 Series provide the foundation for the Cisco Application Centric Infrastructure (ACI) and deliver savings in capital expenditures (CapEx) and operating expenses (OpEx) to achieve an agile IT environment. As you plan your Cisco ACI implementation, Cisco Nexus 9000 Series can be deployed in standalone mode.

Hitachi Storage Modules for Red Hat Ansible

Red Hat Ansible is an open-source tool for automation, configuration management, application software deployment, and for Infrastructure as Code (IaC). When used for IaC, Ansible manages endpoints and infrastructure components in an inventory file, formatted in YAML or INI.

You can enhance your day 1 storage management experience with the integration of Hitachi Vantara Storage Modules into Ansible. These modules include configuration and management modules for on-premises storage and public cloud storage. Additional information on Hitachi Storage Modules can be found here: VSP One Block Storage Modules for Ansible User Guide • Viewer • Welcome to the Hitachi Vantara Documentation Portal.

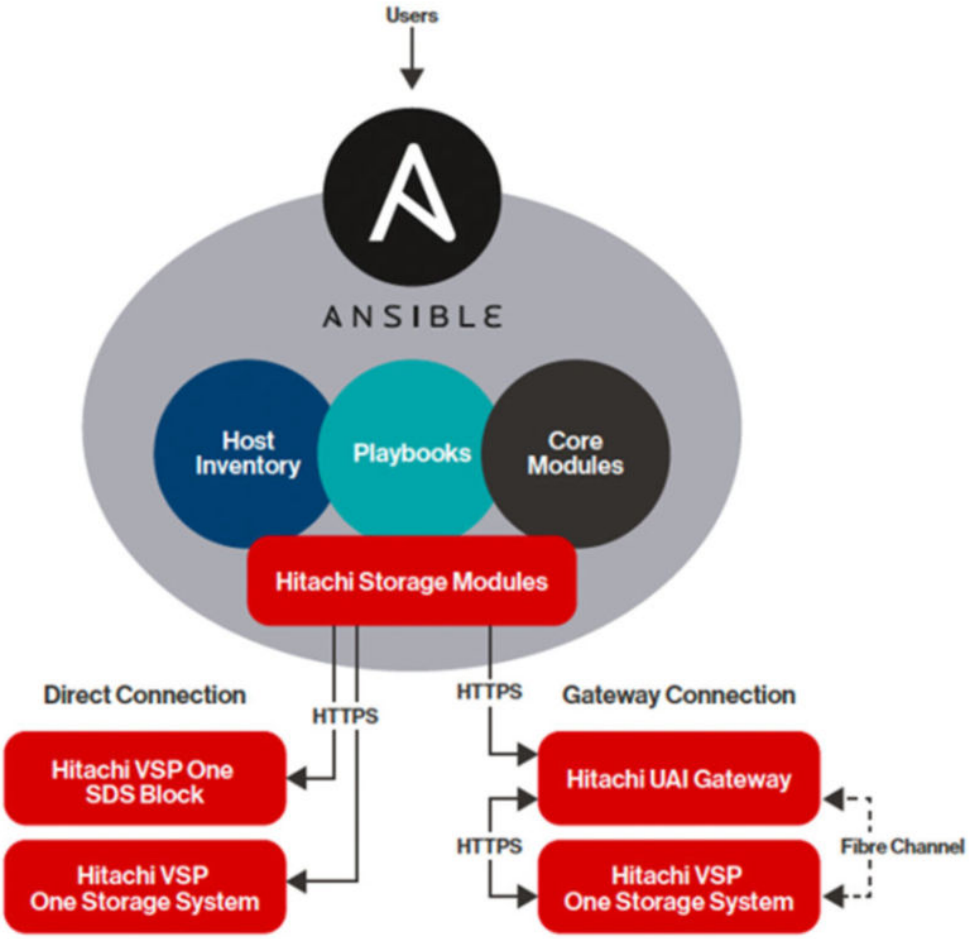

Solution components

This section describes the Hardware and Software components used to validate this reference architecture.

Hardware components

The following table lists the details of the hardware configuration of the solution used to run the different test cases of this reference architecture.

Vendor	Hardware	Detail Description
Hitachi Vantara	Hitachi Advanced Server HA810 G2	▪ 2-Node Compute Cluster • 2 × Intel(R) Xeon(R) Gold 6348 CPU @ 2.60GHz • Memory : 64 GB × 16 (1 TB) ▪ 2-Node Management Cluster: • 2 × Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz • Memory : 64 GB × 8 (512 GB)
Hitachi Vantara	Hitachi Virtual Storage Platform One SDS Block	▪ VSP One SDS Block 6-Node Cluster ▪ Hitachi Advanced Server HA820 G3 Specifications: • 2 × Intel Xeon Gold 5418Y CPU @ 2.00Ghz • Memory : 512 GB DDR5 • HPE MR216i-o Gen11
Cisco	Top-of-Rack 10 GbE switch	Cisco Nexus 93180YC-EX NX-OS Version

Software components

The following table lists the software components of the solution used to run the different test cases of this reference architecture.

Software	Version	Function
Hitachi Virtual Storage Platform One SDS Block	1.15	Software Defined Storage
VMware vCenter Server	8.0.3 Build 24022515	Management Console
VMware ESXi	VMware ESXi, 8.0.3, Build 24022510	Hypervisor
Microsoft Windows Server 2019	Datacenter	Operating System
Microsoft Windows 10	Enterprise Edition	Operating System
VMware Horizon View	8.8	Management console
VMware Horizon Client	2309	
Login Enterprise	5.10.2	Benchmarking tool

Solution design

The infrastructure servers for VMware Horizon and Login VSI used for this solution were placed on separate infrastructure clusters with dedicated resources, also called management cluster. The compute cluster was dedicated for Horizon desktops only.

The following illustration shows the key hardware and software components used for the reference architecture tests.

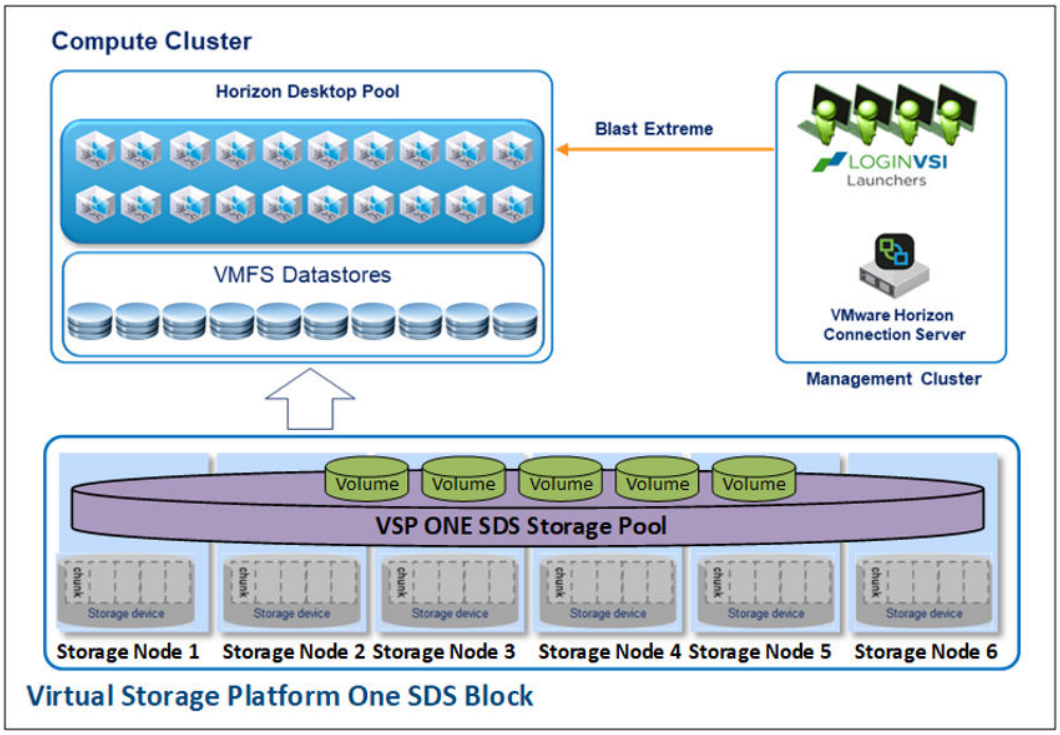

Test environment configuration

This section provides details of the test environment.

Solution infrastructure

The following illustration provides a high-level overview of the solution infrastructure. A 2-node compute cluster was used to host the Horizon desktops, a 2-node cluster was used to host Login VSI and Horizon management VMs, and for storage the Hitachi Virtual Storage Platform One SDS Block was configured using a 5-node cluster. See the Hardware Components section and the Hitachi Virtual Storage Platform One SDS Block configuration table.

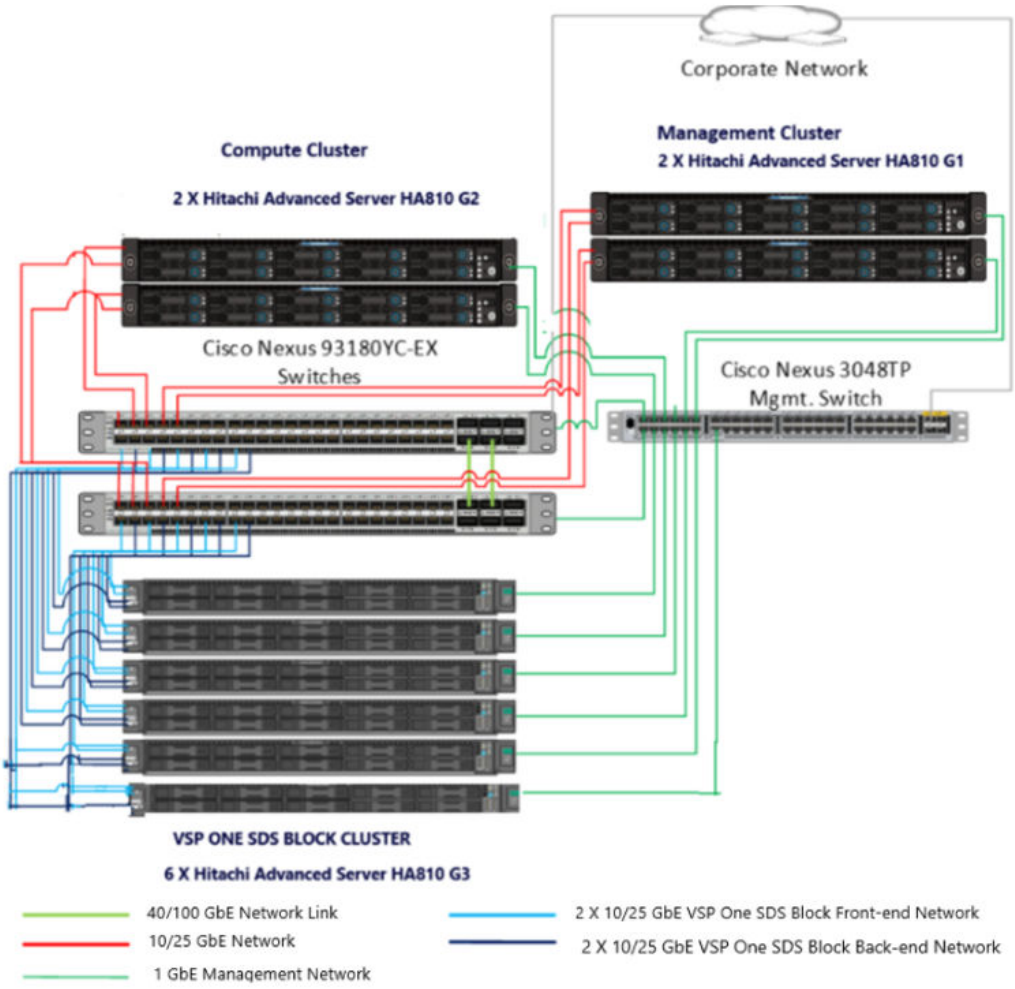

Hitachi Virtual Storage Platform One SDS Block

The following illustration provides an overview of Hitachi Virtual Storage Platform One SDS Block where multiple storage nodes are aggregated to provide a single Storage Pool. Volumes carved from the SDS Storage Pool will be presented to the hosts in the compute cluster using iSCSI protocol.

The SDS cluster used during the tests consisted of 5-nodes with the configuration shown in the Virtual Storage Platform One SDS Block Configuration table.

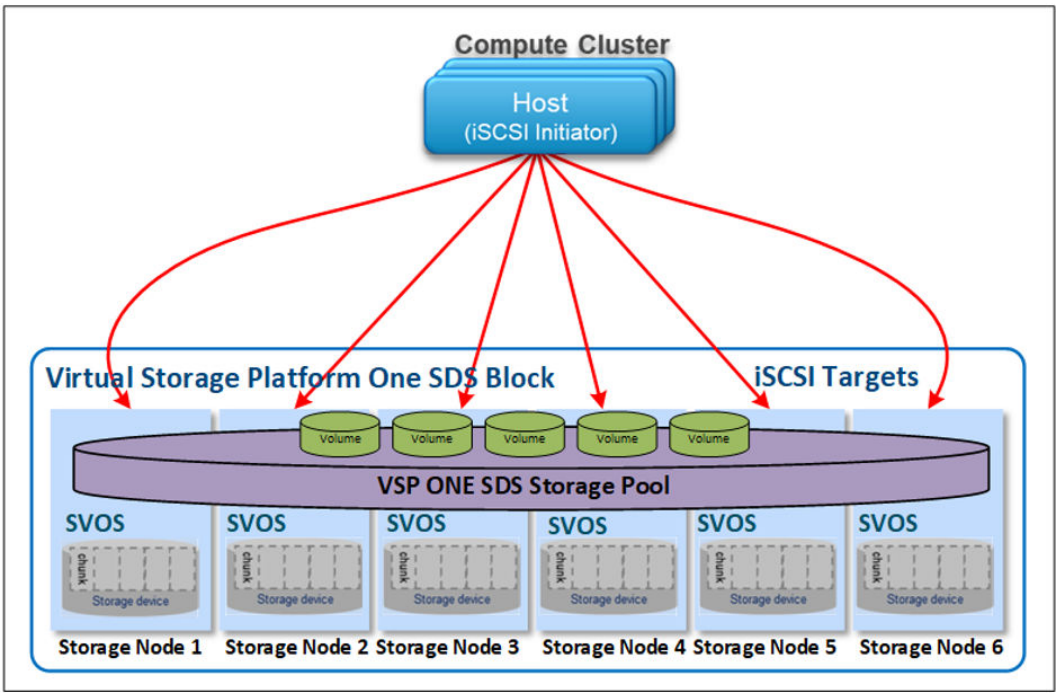

The dashboard window shows information about the Hitachi Virtual Storage Platform One SDS Block system.

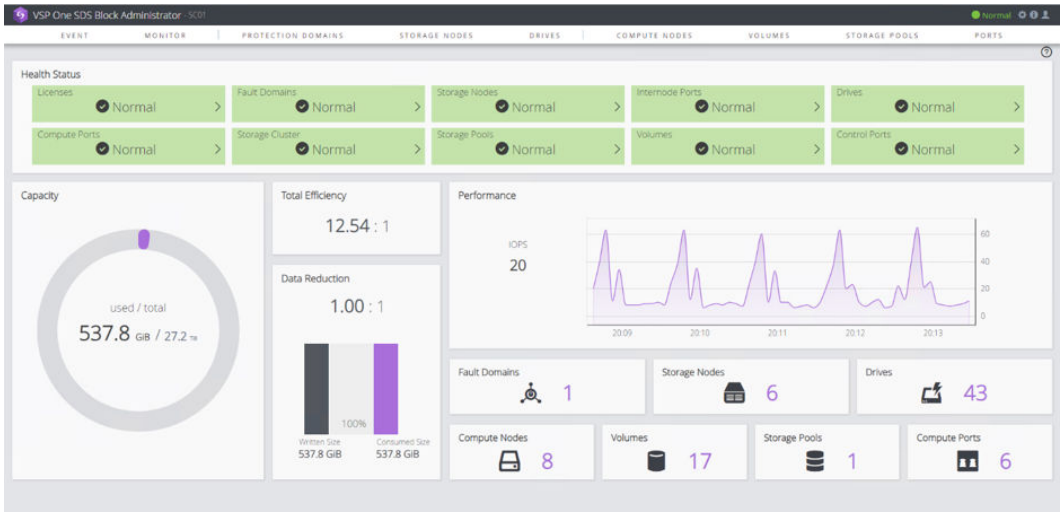

Network configuration

The compute cluster was configured with vSphere Distributed Switch and each host with two 25 GbE network interfaces for the front-end network (iSCSI and SDS management traffic), and two 25GbE network interfaces for the back-end network (SDS internode traffic). Distributed port groups for management traffic, vSphere vMotion, virtual machine and two distributed port groups for iSCSI traffic were created. Also, the jumbo frames option was enabled on the vMotion and iSCSI vmkernel ports, and physical network devices. The round-robin policy was used for iSCSI load balancing.

The following illustrations provide additional details about the network configuration on the compute cluster.

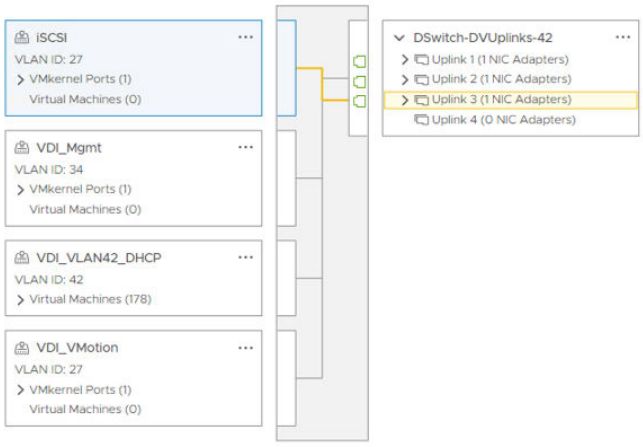

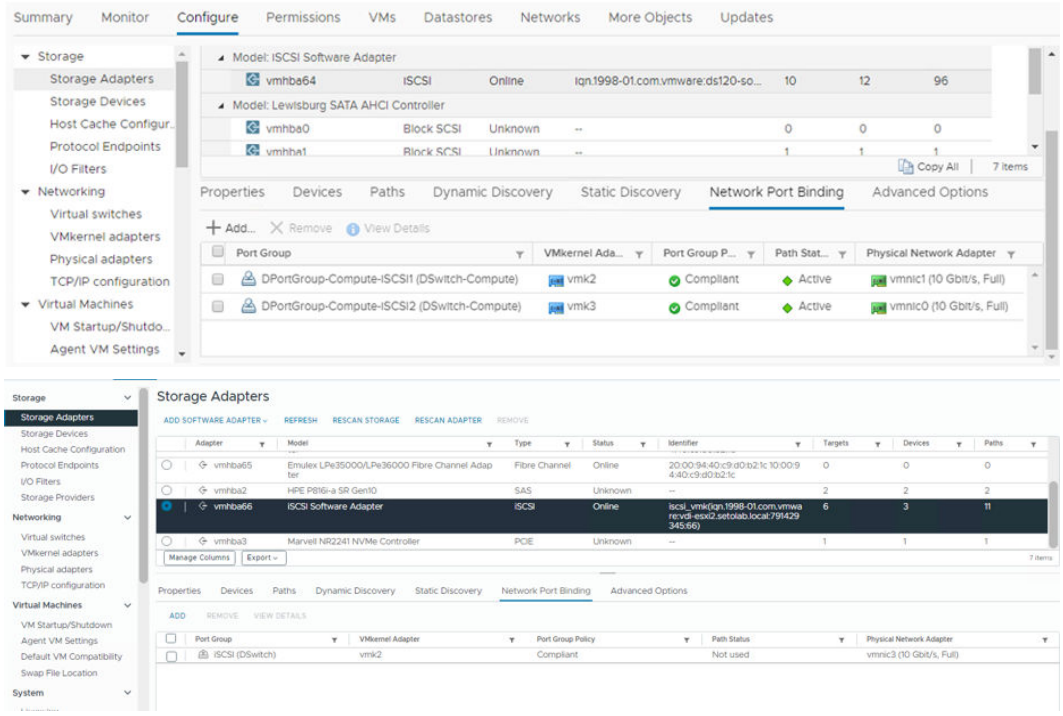

Horizon infrastructure

The following table lists the VMware Horizon View components.

Server Name	vCPU	Memory	Disk Size	Disk Type	Operating System
View Connection Server	8	16 GB	80 GB	Eager Zeroed Thick	Windows Server 2019
Domain Controller	4	8 GB	80 GB	Eager Zeroed Thick	Windows Server 2016

The domain controller was deployed to support user authentication and domain services for the VMware Horizon infrastructure and Login VSI.

The number of Horizon virtual desktops deployed is based on the CPU/Memory resources on the servers of the compute cluster, see Test cases (on page 15) for more details.

Storage configuration

The following video demonstrates storage configuration:

- Add compute nodes to VSP One SDS Block clusters
- Create and assign volumes to compute nodes

https://bcove.video/49p0Aj7

Test methodology

This section outlines the test methodology for the solution.

Login VSI

Login VSI is dedicated to shaping a future where End User Computing (EUC) operations and technology never impede employees' day-to-day productivity. The Login Enterprise platform (https://www.loginvsi.com) combines validation, testing, and monitoring capabilities for VDI and DaaS environments, published applications, and physical devices (desktops, laptops, kiosks and point-of-sale). By delivering unparalleled insights through automation into performance, scalability, and application experiences, the platform helps unlock the maximum value in any EUC environment or deployment.

For over a decade, Login VSI has been at the forefront, providing application-centric workloads and industry-standard benchmarks, including the End User Experience (EUX) Score and VSImax. Widely adopted by the EUC vendor community, these benchmarks serve as objective measures for performance evaluation and validation. This established frame of reference offers critical guidance and insights for EUC technologists, empowering them to make informed decisions and optimize the performance of their environment.

Load Generation Tool

Login VSI was used to generate light and heavy workload on the desktops. Login VSI launchers were each configured to initiate up to 20 sessions with the VMware Blast protocol to the VMware Horizon View Connection Server to simulate end-to-end execution of the entire VMware Horizon View infrastructure stack.

The Login VSI Knowledge Worker profile is used for test scenarios. For all the tests we used the VMware Blast Extreme display protocol for the desktop sessions.

The following table lists the applications that Login VSI exercised during workload testing for knowledge workload profiles.

- Outlook
- Microsoft Edge
- Word
- PowerPoint

See https://support.loginvsi.com/hc/en-us/article_attachments/15202963063452 for details.

Desktop VM configuration

The following table lists the configuration profiles of the Windows 10 VM templates used for the Task, Knowledge, and Power worker workloads. Each image was optimized with the VMware OS Optimization tool in accordance with Creating an Optimized Windows Image for a VMware Horizon Virtual Desktop and in conformance with the Login VSI test standards.

VM Profile	VM OS	Disk	vCPU	VM Memory (configured)	Resolution
Knowledge Worker	Windows 10 Enterprise 64-bit	40 GB (Thin)	2	4 GB	1920 × 1080

Test cases

The series of tests that were run for this reference architecture are designed to capture the performance capabilities and user experience (based on Login VSI workloads) when using Virtual Storage Platform One SDS Block. The following table provides a summary of the test cases, and the total of desktops tested in each case.

Test Case	VM Profile/Login VSI Workload	Density per host	Total VMs in the Cluster
Test Case 1	Knowledge Worker	175	350

Login VSI results

Performance testing was conducted on 350 Windows 10 Instant Clones, achieving an EUX Score of 6 out of 10 which is considered as a fair score, especially when using iSCSI and a CPU Score of 9.2. The test ran for 68 minutes, with the initial 48 minutes dedicated to logging on sessions, the next 20 minutes for full load testing, and the remaining time for logging off sessions and recreating instant clones.

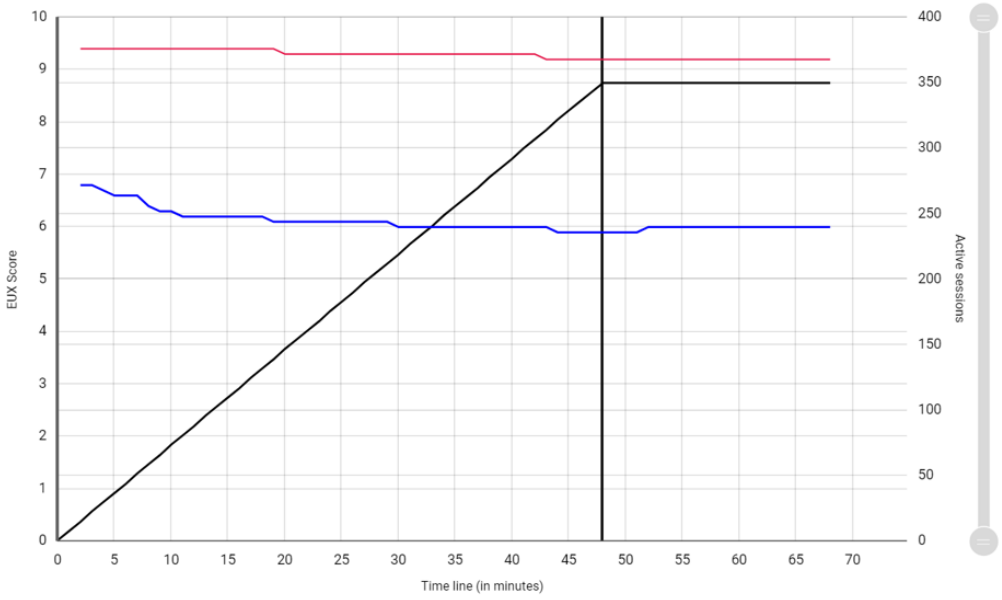

VSP One SDS Block performance

The following figures show IOPS, latency, and transfer rate usage at the VSP storage pool level during the log on, steady state, and log off phases of testing. During the log on phase of IOPS started around 5000 and reached to 30000 which indicates actively more sessions logging and IOPS reached to maximum of 33000 during full load which pushes the maximum capacity and there is gradually drop of IOPS due to sessions logoff.

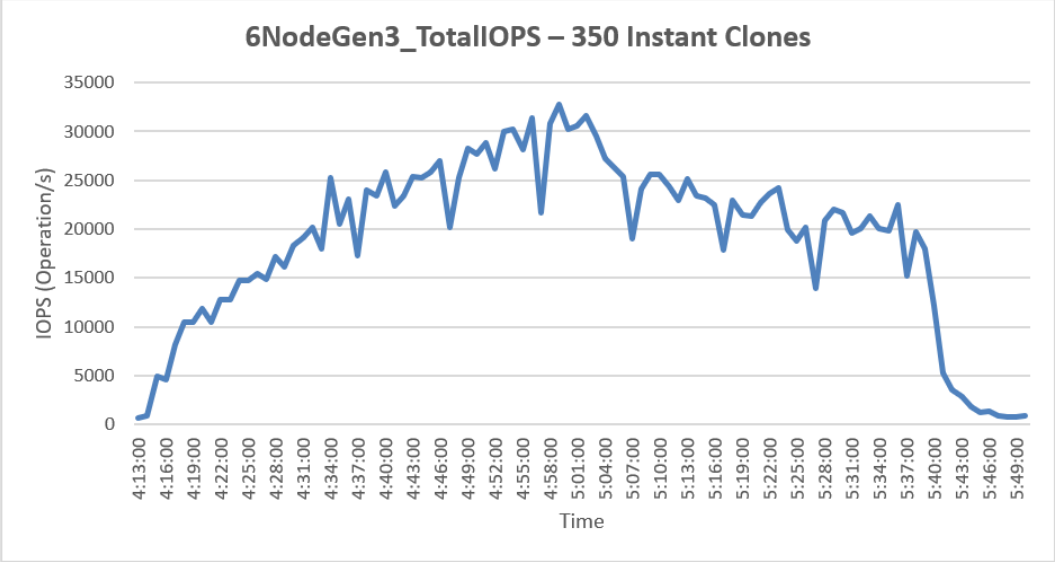

In the 350 sessions of knowledge user testing, the latency performance was quite impressive. The peak latency recorded was 4.7 milliseconds, while the average latency consistently ranged between 2.5 to 4 milliseconds throughout the tests. This indicates a stable and responsive system, ensuring smooth and efficient user interactions.

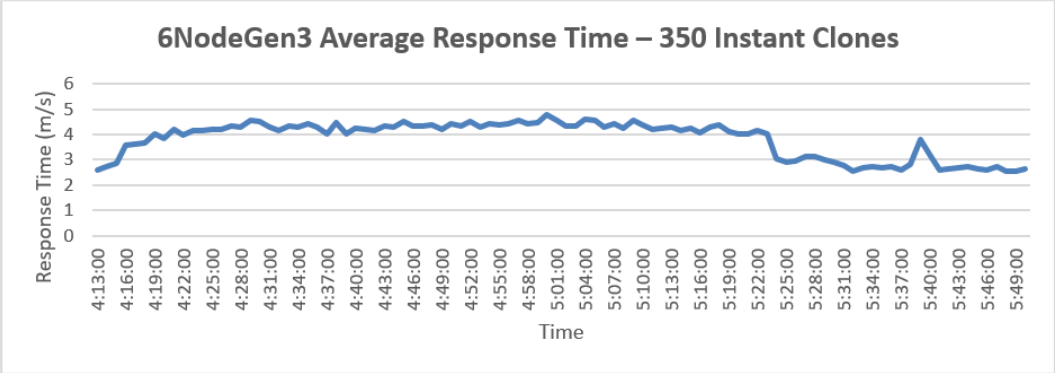

The following figure shows Total Transfer Rate of 350 sessions of Windows10 instant clones, during the log on phase of transfer rate started around 100 mb/s and reached to 600 which indicates actively more sessions logging and IOPS reached to maximum of 658 during full load which pushes the maximum capacity and there is gradually drop of IOPS due to sessions logoff.

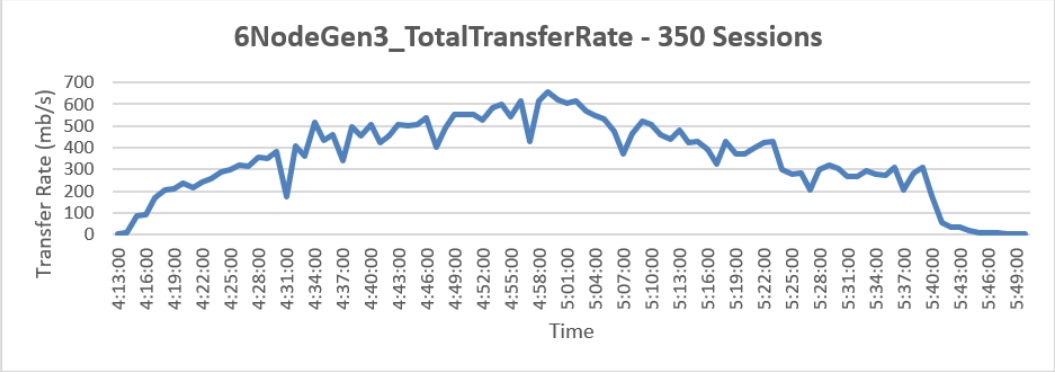

Conclusion

This architecture provides guidance to organizations implementing Horizon VDI using the new Hitachi Virtual Storage Platform One SDS Block system, and describes tests performed by Hitachi Vantara to validate and measure the performance and capabilities of the recommended solution, including third-party validated performance testing from Login VSI, the industry standard benchmarking tool for virtualized workloads.

Organizations looking for VDI solutions can confidently deploy them on top of VSP One SDS Block and Hitachi Advanced Server to ensure office, knowledge, and power users a good user experience.

The performance testing of Horizon VDI with VSP One SDS Block storage demonstrated impressive results, particularly with 350 Windows 10 instant clones. Achieving a Login VSI Score of 6 and an CPU Score of 9.2 highlights the system's robust performance and user experience.

The key benefits of using VSP One SDS Block storage in this setup include:

- Scalability: The SDS architecture makes it possible for scaling to be a seamless task that can hold even the most demanding of workloads without sacrificing the quality of performance.
- Cost Efficiency: SDS allows businesses to use their resources more efficiently and therefore to reduce the total costs associated with storage.
- Flexibility: SDS provides the flexibility to manage and allocate storage resources dynamically, ensuring optimal performance for various applications.
- High Performance: The combination of Horizon VDI and VSP One SDS Block storage ensures high throughput and low latency, critical for maintaining a smooth user experience.
- Simplified Management: Centralized management of storage resources simplifies administrative tasks and enhances operational efficiency.

Overall, the integration of Horizon VDI with VSP One SDS Block storage proves to be a powerful solution for delivering high-performance virtual desktop environments, ensuring both efficiency and excellent user experience.

Appendix

This appendix contains the following:

- Ansible deployment and configuration
- Hitachi Virtual Storage Platform One SDS Block provisioning with Ansible

Ansible deployment and configuration

Ansible by Red Hat is a popular open-source infrastructure and application automation tool, giving speed and consistency to deployments and configuration. Ansible is designed around these principles:

- Agent-less architecture — Low maintenance overhead by avoiding the installation of additional software across IT infrastructure.
- Simplicity — Automation playbooks use straightforward YAML syntax for code that reads like documentation. Ansible is also decentralized, using SSH existing OS credentials to access remote machines.
- Scalability and flexibility — Easily and quickly scale the systems you automate through a modular design that supports a large range of operating systems, cloud platforms, and network devices.
- Idempotence and predictability — When the system is in the state your playbook describes Ansible does not change anything, even if the playbook runs multiple times.

Ansible runs on many Linux platforms, Apple OSX, and MS Windows. Installation instructions will vary between platforms, and for this environment, the instructions and the control host used for playbook invocation will be for a RHEL 9 VM. Ansible can be installed at the system level, but convention and instructions from the Ansible site steer users to install with pip to create a user-specific instance.

With the base RHEL VM installed as a control host, Ansible was set up for the invoking user account with the following:

- `sudo dnf install pip`
- `sudo dnf install git`
- `pip install ansible`

Hitachi Virtual Storage Platform One SDS Block provisioning with Ansible

Product user documentation and release notes for Storage Modules for Red Hat Ansible are available on Hitachi Vantara Documentation and identify the available storage modules, which can be used to configure Hitachi Virtual Storage Platform One SDS Block using direct connection. Check the website for the most current documentation, including system requirements and important updates that may have been made after the release of the product.

See VSP One Block Storage Modules for Ansible User Guide • Viewer • Welcome to the Hitachi Vantara Documentation Portal for details.

In this demonstration, we will showcase an Ansible playbook designed for managing VSP One SDS Block. The playbook includes tasks to create a volume, set up a compute node, and retrieve the number of existing volumes. By automating these processes, we can streamline storage management, ensuring efficient and consistent operations. This approach not only saves time but also reduces the potential for human error, making it an invaluable tool for administrators managing complex storage environments.

Create a node with Ansible

To add a compute node named Esxi_Node2 to a VSP One SDS Block cluster, you need to create and execute a playbook named `create_volume.yml`. This playbook should include the necessary arguments to configure the new node correctly. The arguments typically include details such as the node name, IP address, storage configuration, and any specific parameters required by your environment. When the playbook is ready, you can run it using Ansible to automate the addition of Esxi_Node2 to your cluster, ensuring that it integrates seamlessly with the existing infrastructure. This process helps maintain consistency and reduces manual errors during the setup.

```
[root@localhost sds_block_direct]# cat create_compute.yml
- name: Compute Node Module
  hosts: localhost
  gather_facts: false
  collections:
    - hitachivantara.vspone_block.sds_block

  vars_files:
    - ../ansible_vault_vars/ansible_vault_storage_var.yml
    - sds_block_variables.yml

  vars:
    # Common connection info for all tasks
    connection_info:
      address: "{{ storage_address }}"
      username: "{{ vault_storage_username }}"
      password: "{{ vault_storage_secret }}"

  tasks:
##############################################################################
# Task 1 : Create compute node
##############################################################################
    - name: Create compute node
      hv_sds_block_compute_node:
        connection_info: "{{ connection_info }}"
        state: "{{ state_present }}"
        spec:
          name: Esxi_Node2
          os_type: VMware
      register: result

    - debug:
        var: result
```

```
[root@localhost sds_block_direct]# ansible-playbook create_compute.yml
[WARNING]: No inventory was parsed, only implicit localhost is available
[WARNING]: provided hosts list is empty, only localhost is available. Note that the implicit localhost does not match 'all'

PLAY [Compute Node Module] *****************************************************************************************************

TASK [Create compute node] *****************************************************************************************************
changed: [localhost]

TASK [debug] *******************************************************************************************************************
ok: [localhost] => {
    "result": {
        "changed": true,
        "data": {
            "compute_node_info": {
                "id": "c28b0acd-0879-4e31-9bb9-d7cc73699ac7",
                "name": "Esxi_Node2",
                "number_of_paths": 0,
                "number_of_volumes": 0,
                "os_type": "VMware",
                "paths": [],
                "total_capacity_mb": 0,
                "used_capacity_mb": 0,
                "vps_id": "(system)",
                "vps_name": "(system)"
            },
            "volume_info": []
        },
        "failed": false
    }
}

PLAY RECAP *********************************************************************************************************************
localhost                  : ok=2    changed=1    unreachable=0    failed=0    skipped=0    rescued=0    ignored=0

[root@localhost sds_block_direct]#
```

Provision volumes with Ansible

To create an SDS volume using Ansible, you can write a playbook named myplaybook.yml. In this playbook, you define a task to create a 100 MB volume with the name sds_ansible_test. After running the playbook, the volume will be created with the specified name. The playbook might look something like this: it specifies the host as EsxiNode1 and includes a task that uses the sds_volume module to create the volume with the desired size and name. Once executed, the playbook ensures that the SDS volume is created as intended.

```
[root@localhost sds_block_direct]# cat myplaybook1.yml
########################################################################
# Example : Volume Playbook
########################################################################
- name: Volume Module
  hosts: localhost
  gather_facts: false
  collections:
    - hitachivantara.vspone_block.sds_block

  vars_files:
    - ../ansible_vault_vars/ansible_vault_storage_var.yml
    - sds_block_variables.yml
#   - sds_block_variables_baclup.yml

  vars:
    # Common connection info for all tasks
    connection_info:
      address: "{{ storage_address }}"
      username: "{{ vault_storage_username }}"
      password: "{{ vault_storage_secret }}"

  tasks:
    ########################################################################
    # Task 1 : Create volume
    ########################################################################
    - name: Create volume
      hv_sds_block_volume:
        connection_info: "{{ connection_info }}"
        state: "{{ state_present }}"
        spec:
          name: sds_ansible_test
          capacity: 100MB
          pool_name: SP01
          capacity_saving: Compression
          compute_nodes: Esxi_Node1
      register: result

    - debug:
        var: result
```

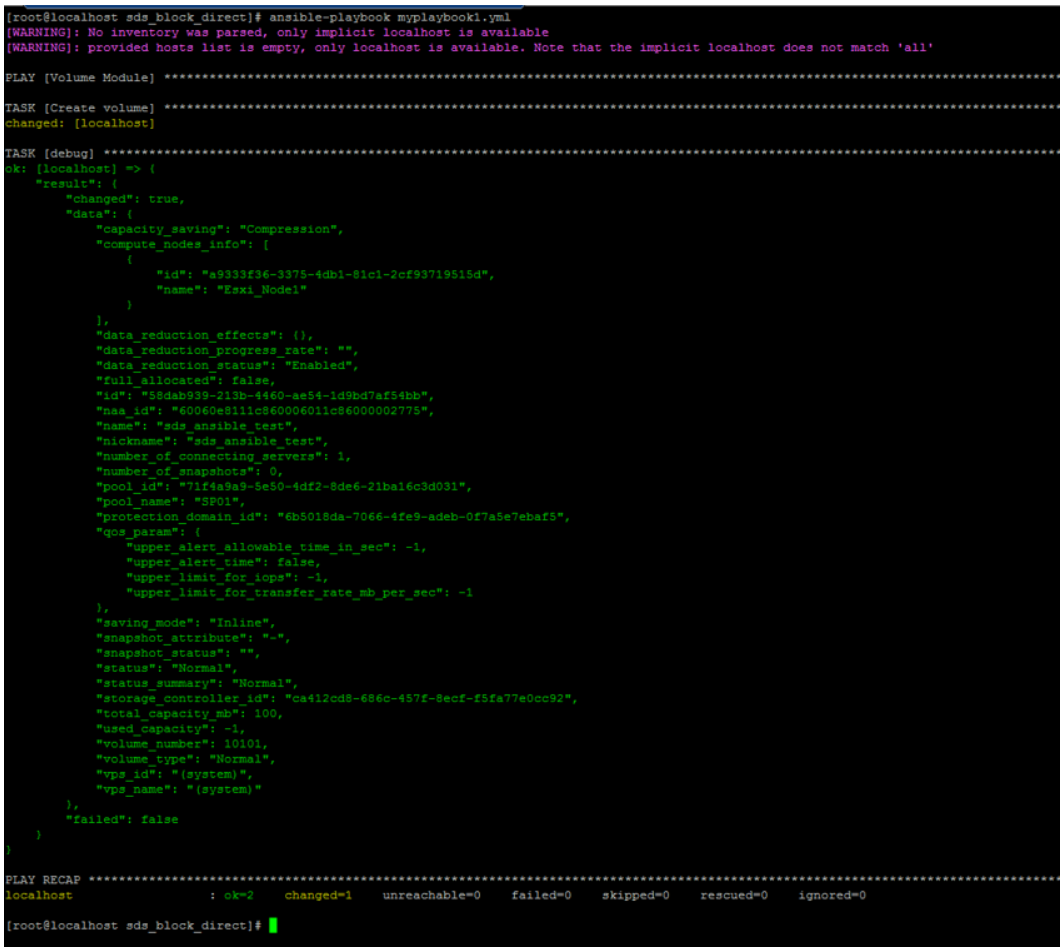

There are several Ansible playbooks available for automating tasks with VSP One SDS Block storage. These playbooks can help streamline various storage management tasks such as provisioning, configuration, and maintenance.

You can find these playbooks and modules on GitHub, specifically in GitHub - hitachi-vantara/vspone-block-ansible: Public repository for Hitachi Storage Modules for Red Hat® Ansible®. This repository includes comprehensive tools and documentation to help you get started with automating your storage solutions.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact