

# Clouidian Hyperstore on VMware Cloud Foundation with Tanzu

Enterprise-grade, fully native S3-compatible object storage  
with the VMware vSAN Data Persistence platform

## Table of Contents

Executive Summary	4
Business Case .....	4
Business Value .....	4
Key Results .....	4
Audience .....	5
Technology Overview	5
VMware Cloud Foundation with Tanzu .....	5
VMware vSphere .....	5
VMware vSAN .....	6
VMware vSAN Data Persistence platform .....	6
VMware NSX Data Center .....	8
Cloudian HyperStore .....	8
Solution Configuration	8
Architecture Diagram .....	9
Hardware Resources .....	9
Software Resources .....	10
Network Configuration .....	10
Solution Validation	11
Overview .....	11
Deploying Cloudian HyperStore using vSAN Direct Configuration .....	12
Deploying Cloudian HyperStore using vSAN-SNA .....	12
Test Tools .....	12
Failure and Maintenance Scenarios .....	13
Best Practices	14
VMware Cloud Foundation Infrastructure .....	14
Cloudian HyperStore .....	14
HyperStore Storage Policies .....	15
Use Cases	16
Use Case 1: Cloud Native Apps .....	16

Use Case 2: Data Protection .....	16
Use Case 3: Big Data Apps.....	16
Use Case 4: File Tiering .....	17
Conclusion .....	18
Reference .....	18
About the Author .....	18

## Executive Summary

### Business Case

As enterprises increasingly adopt and deploy modern, cloud-native applications, there is a need to ensure associated storage is also modernized and up to date. Applications are shifting away from traditional storage and are turning towards the S3 API as the standard for cloud object-based storage.

Cloudian®'s enterprise-grade, fully native, S3-compatible object storage software integrated with the VMware vSAN™ Data Persistence platform enables new efficiencies and savings by allowing enterprises to run both modern, cloud-native applications, and traditional applications on a single, shared storage environment, at any scale, on-premises and in private cloud. The solution extends Cloudian HyperStore®'s simple-to-deploy and easy-to-manage, exabyte scalable, highly secure, multi-tenant storage to any application, while also reducing total cost of ownership (TCO). The integration supports Cloudian HyperStore with VMware Cloud Foundation™ with VMware Tanzu™, and combines the industry-leading VMware hyperconverged infrastructure (HCI) and Cloudian HyperStore into a single, shared-nothing data platform.

In this solution, we provide deployment procedures, design and sizing guidance, best practices for enterprise infrastructure administrators and application owners to run Cloudian HyperStore object storage on VMware Cloud Foundation with Tanzu.

### Business Value

Here are the top 6 benefits to deploy Cloudian HyperStore with vSAN Data Persistence platform on VMware Cloud Foundation with Tanzu:

- **Enterprise-grade object storage on vSAN Data Persistence platform:** Leverage VMware vSAN Direct Configuration for object storage to improve storage efficiency and address multiple new use case needs for object storage.
- **Automated deployment and scaling through Kubernetes:** Easily install, deploy, and scale Cloudian HyperStore on VMware vSphere clusters straight from the VMware vCenter® UI dynamically and in a self-service manner through Kubernetes APIs.
- **Service-aware infrastructure operations:** Minimize the disruptions and impact of infrastructure operations such as planned maintenance and lifecycle management through a deeply integrated solution.
- **vSphere-integrated capacity and health monitoring:** Seamlessly manage multiple HyperStore clusters across data centers, edge environments, and public clouds without leaving vSphere.
- **CAPEX savings:** Leverage Cloudian data policies (replication and erasure coding) with vSAN Direct Configuration to lower costs and improve storage efficiency, while also ensuring data durability.
- **OPEX savings:** Consolidate both traditional workloads and modern, cloud-native workloads on the same platform, eliminating infrastructure silos and capitalizing on common skills and training.

### Key Results

This reference architecture is a showcase of VMware Cloud Foundation with Tanzu for operating and managing Cloudian HyperStore, the industry-leading, enterprise-grade, fully native S3-compatible object storage platform. HyperStore leverages VMware's vSAN Data Persistence platform in a fully integrated SDDC environment. Key results can be summarized as follows:

- Cloudian brings highly secure, exabyte-scalable enterprise-grade S3-compatible object storage to VMware Cloud Foundation, allowing customers to deploy Cloudian HyperStore on VMware vSphere clusters and leverage underlying vSAN disks for storage of object data.

- The integrated solution allows IT operators to modernize their storage infrastructure to meet the requirements of any modern, cloud native or traditional application on a unified storage platform, at any scale, on-premises, and in the cloud, delivering better performance and availability, at lower TCO.
- The integration supports Cloudian HyperStore on VMware Cloud Foundation with Tanzu, and combines industry-leading VMware HCI platform and Cloudian HyperStore into a single, shared-nothing data platform for all applications.

### Audience

This solution is intended for IT administrators, IT operators, and storage experts who are involved in the planning, design, and deployment of containerized workloads on VMware Cloud. It is assumed that the reader is familiar with the concepts and operations of S3 compatible object storage and VMware Cloud Foundation related components.

### Technology Overview

Solution technology components are listed below:

- VMware Cloud Foundation
  - VMware Cloud Foundation with Tanzu
  - VMware vSphere
  - VMware vSAN
  - VMware vSAN Data Persistence platform
  - VMware NSX Data Center
- Cloudian HyperStore

#### VMware Cloud Foundation with Tanzu

VMware Cloud Foundation with Tanzu is the best way to deploy Kubernetes at scale. VMware Cloud Foundation is a ubiquitous hybrid cloud platform for both traditional enterprise apps and modern apps, providing a complete set of secure software-defined services for compute, storage, network security, Kubernetes management, and cloud management. VMware Cloud Foundation with Tanzu automates full-stack deployment and operation of Kubernetes clusters through integration with VMware Tanzu Kubernetes Grid. This helps eliminate manual steps for configuring hosts, creating logical relationships, managing hypervisors for faster deployment of applications at scale. The most exciting feature added to the VMware Cloud Foundation architecture is the integration of Kubernetes directly into the vSphere hypervisor, which delivers an entirely new set of VMware Cloud Foundation Services, a new Kubernetes and RESTful API surface that empowers developers to have self-service access to Kubernetes clusters, vSphere Pods, virtual machines, persistent volumes, stateful services, and networking resources. The result is an agile, reliable and efficient hybrid cloud platform that bridges the gap between app developers and IT administrators.

#### VMware vSphere

VMware vSphere is the next-generation infrastructure for next-generation applications, which provides a powerful, flexible, and secure foundation for business agility that accelerates the digital transformation to cloud computing and promotes success in the digital economy. VMware vSphere embeds containers and Kubernetes into vSphere, unifying them with virtual machines as first-class citizens. This enables all VI admins to become Kubernetes admins and easily deliver new services to their developers. VMware vSphere addresses key challenges faced by the IT admins in areas of lifecycle management, security, and performance and resiliency needed by business-critical applications, AI/ML applications and latency sensitive applications. With VMware

vSphere, customers can run, manage, connect, and secure both traditional and cloud-native applications in a common operating environment, across clouds and devices.

#### VMware vSAN

VMware vSAN is the industry-leading software powering VMware's software-defined storage and HCI solution. vSAN helps customers evolve their data center with reduced risk, control IT costs, and scale to tomorrow's business needs. vSAN, native to the market-leading hypervisor, delivers flash-optimized, secure storage for all of your critical vSphere workloads, and is built on industry-standard x86 servers and components that help lower TCO in comparison to traditional storage.

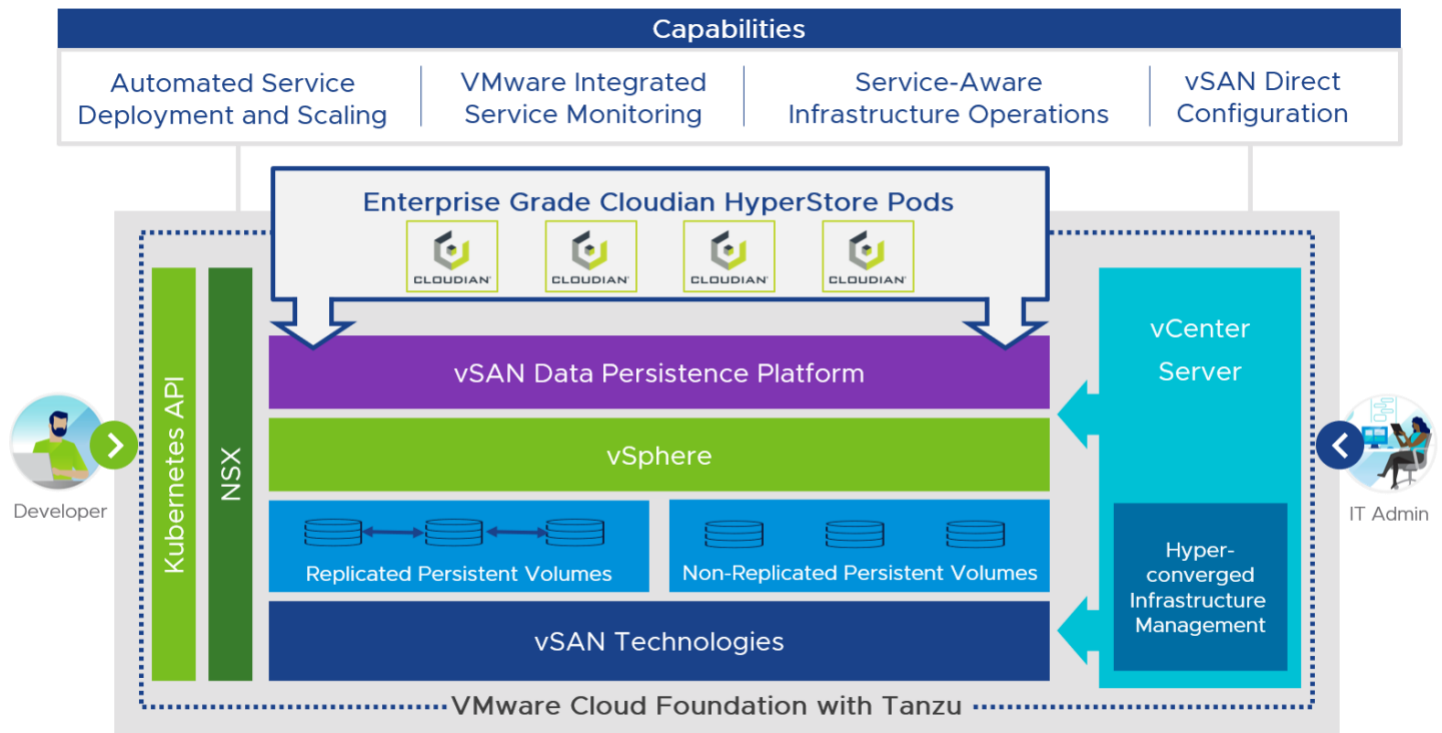
vSAN simplifies Day 1 and Day 2 operations, and customers can quickly deploy and extend cloud infrastructure and minimize maintenance disruptions. Together with [Cloudian HyperStore](#), vSAN modernizes HCI by providing admins with a unified storage control plane for all block, file and object protocols, and provides significant enhancements that make it a great solution for VMs as well cloud native applications. vSAN helps reduce the complexity of monitoring and maintaining infrastructure and enables admins to rapidly provision storage for Kubernetes-orchestrated cloud native applications.

#### VMware vSAN Data Persistence platform

The vSAN Data Persistence platform provides an as-a-service framework for VMware partners that offer modern stateful services to integrate with the underlying virtual infrastructure, allowing you to run stateful services with high velocity scaling, simplified IT operations, and optimized TCO. You can deploy a stateful service alongside traditional applications on a regular vSAN cluster with vSAN-SNA (vSAN support for Shared Nothing Architecture) policy, or deploy it on a dedicated vSAN cluster with VMware vSAN Direct Configuration™, a technology enabling direct access to the underlying direct-attached hardware which can be optimized for the application needs. Both options benefit from optimal storage efficiency for stateful services by leveraging service-level replication, as well as unified management of services in vCenter.

The platform offers a way for the IT admin to enable, manage and monitor all aspects of the stateful service from vSphere Interfaces (API/UI) while the developers get public cloud-like simple self-service consumption experience.

The regular vSAN cluster with vSAN-SNA policy or a dedicated vSAN cluster with vSAN Direct Configuration makes it easy for cloud-native services such as Cloudian integrated into the vSAN Data Persistence platform to co-locate its compute and storage on the same physical ESXi host. This host-local placement then allows us to do replication only at the service layer and not at the storage layer.



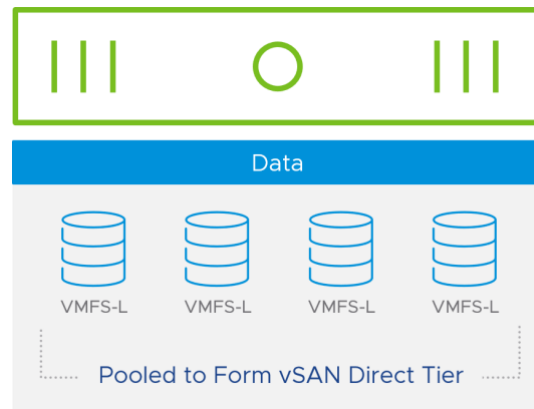
- vSAN with SNA Storage Policy

With this technology, you can use a distributed replicated vSAN datastore with the vSAN host-local SNA policy. The technology makes it easy for the stateful service to co-locate its compute instance and a storage object on the same physical ESXi host. The compute instance such as a pod has to come up first on one of the nodes in the vSAN cluster and then the vSAN object created with vSAN-SNA policy (vsan-sna storage class) will automatically have all of its data placed on the same node where the pod is running.

Although vSAN with SNA storage policy can place data local to the compute, we still leverage a distributed vSAN data path between the application and the raw physical disk. Also, the application can only specify affinity at the node granularity and cannot use different disks attached to the same node as independent fault domains.

- vSAN Direct Configuration

vSAN Direct Configuration provides optimal host local storage to shared nothing cloud native services, it creates an independent datastore on every disk attached to a physical host and makes it available as a placement choice to modern applications. vSAN Direct Configuration extends the simplicity of HCI management to host local VMFS disks. It manages and monitors VMFS-L disks and provides insights into health, performance, and capacity of these disks. The VMFS datastores that vSAN Direct Configuration manages are exposed as storage pools in Kubernetes. vSAN Direct Configuration allows modern applications direct access to the disks while giving the ease of management to the VI admin, these disks are consumed just as vSAN with minimal overhead.



#### VMware NSX Data Center

VMware NSX Data Center is the network virtualization and security platform that enables the virtual cloud network, a software-defined approach to networking that extends across data centers, clouds, and application frameworks. With NSX Data Center, networking and security are brought closer to the application wherever it's running, from virtual machines to containers to bare metal. Like the operational model of VMs, networks can be provisioned and managed independent of underlying hardware. NSX Data Center reproduces the entire network model in software, enabling any network topology—from simple to complex multi-tier networks—to be created and provisioned in seconds. Users can create multiple virtual networks with diverse requirements, leveraging a combination of the services offered via NSX or from a broad ecosystem of third-party integrations ranging from next-generation firewalls to performance management solutions to build inherently more agile and secure environments. These services can then be extended to a variety of endpoints within and across clouds.

#### Cloudian HyperStore

Cloudian HyperStore is highly secure, enterprise-grade, fully native, S3-compatible object storage software. HyperStore-C (HSC) is a containerized version of HyperStore, designed to run on the VMware vSAN Data Persistence Platform. It enables new efficiencies and savings by allowing enterprises to run both modern, cloud-native applications and traditional applications on a single, shared storage environment, at any scale, on-prem and in both private and public clouds. The solution extends HyperStore's simple-to-deploy and easy-to-manage, exabyte scalable, highly secure, multi-tenant storage to any application, while also reducing costs by 60% and more. The integration supports Cloudian HyperStore with VMware Cloud Foundation with Tanzu, and combines industry-leading VMware HCI platform and Cloudian HyperStore into a single, shared-nothing data platform.

HyperStore supports modular growth with flexible deployment options and scales non-disruptively across multiple locations to exabytes. HyperStore stays simple with a single storage fabric and a single, global namespace enabling a unified view and control of all data across locations. HyperStore supports this geo-distribution of data with flexible storage-level data protection using erasure coding and replication to ensure the desired level of data durability. With the industry's highest S3 API compliance, interoperability with S3-compatible applications is assured. For hybrid and multi-cloud needs, data tiering and replication to remote locations including public cloud are feature options. Secure multi-tenancy features are native to HyperStore and include QoS and billing. HyperStore data security certifications include Common Criteria, FIPS 140-2, SEC 17a-4(f), among others, providing data immutability, encryption for data-at-rest and data-in-flight, among other capabilities. HyperStore is also hardened storage with access controls, secure shell, and integrated firewall.

#### Solution Configuration

This section introduces the resources and configurations:

- Architecture diagram
- Hardware resources



- Software resources
- Network configuration

Architecture Diagram

The following sections describe how Cloudian HyperStore components are configured.

Cloudian HyperStore is typically deployed on a workload domain1 with a recommended N+1 minimum of 4-node cluster. The following management domain example consists of a 4-node cluster shared across all workload domains.

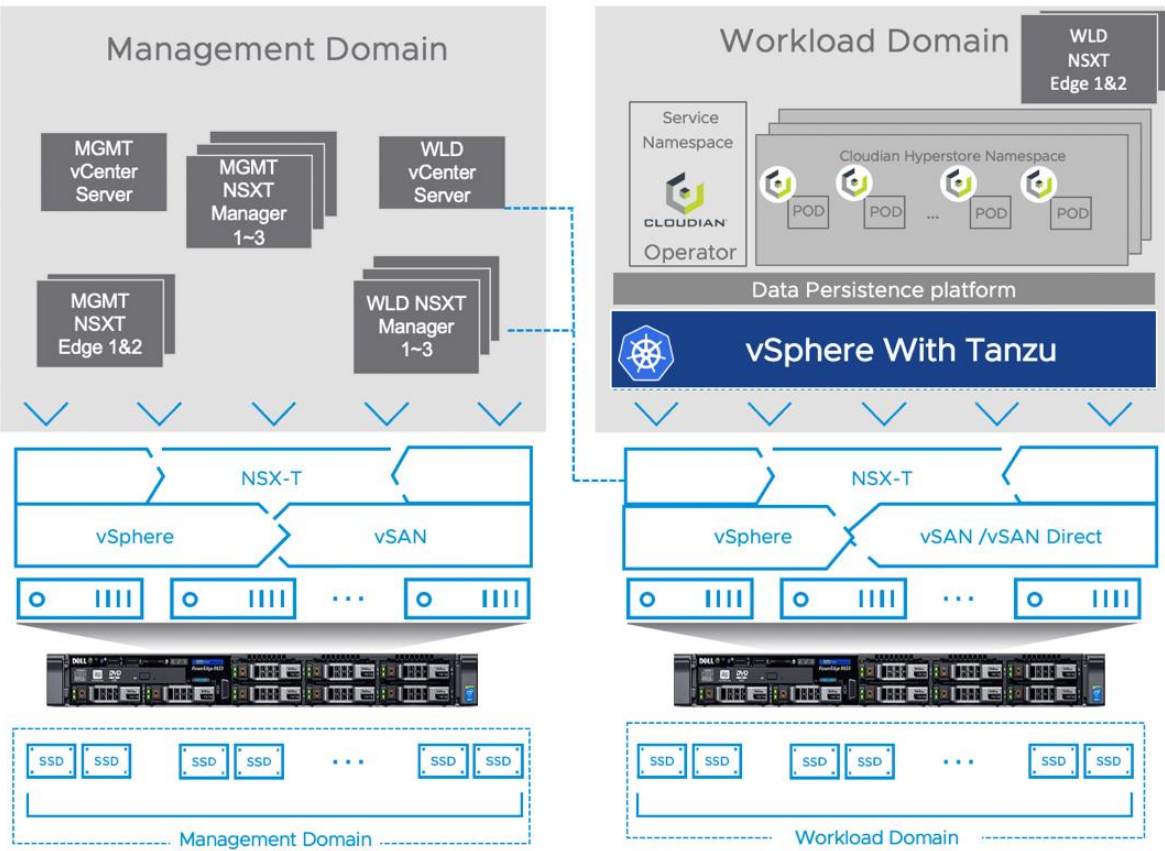


Figure 1. Architectural Diagram

Hardware Resources

Table 1 shows the minimum hardware requirements for each ESXi host in the workload domain cluster.

Table 1. Minimum Hardware Configuration

1. In VMware Cloud Foundation, a workload domain is a policy-based resource construct with specific availability and performance attributes. It combines compute (vSphere), storage (vSAN), networking (NSX), and cloud management (vRealize Suite) into a single consumable entity. Workload domains greatly speed up the instantiation of Kubernetes clusters, deploying the underlying infrastructure in an automated fashion. Workload domains also allow IT admins and developers to securely sandbox and allocate the right infrastructure for containers alongside VMs.

PROPERTY	SPECIFICATION
CPU	1 Socket, 12 core CPUs
RAM	128GB
Network adapter	Min 2 X 10Gb. Recommended 25Gb.
Storage adapter	Must be listed on vSAN Hardware Compatibility List (HCL)
Disks	Required minimum flash drives for metadata storage class is 1 per node. The recommended number of minimum data drives in vSAN Direct object storage class is 10 per node.

### Software Resources

Table 2 shows the software resources used in this solution.

**Table 2. Software Resources**

SOFTWARE	VERSION	PURPOSE
VMware Cloud Foundation	4.2	VMware Cloud Foundation provides integrated cloud infrastructure (compute, storage, networking, and security) and cloud management services to run both modern, cloud native and traditional workloads. See <a href="#">VMware Cloud Foundation</a> for details.
HyperStore	HSC V1.0	HyperStore brings the advantages of a software defined s3-compliant object storage solution to VMware Tanzu's container ecosystem.

### Network Configuration

**Table 3. Virtual Distributed Switch Teaming Policy**

PORT GROUP	TEAMING POLICY	VMNIC0	VMNIC1
Management network	Route based on Physical NIC load	Active	Active <sup>2</sup>
VM network	Route based on Physical NIC load	Active	Active
vSphere vMotion	Route based on Physical NIC load	Active	Active
vSAN	Route based on Physical NIC load	Active	Active

<sup>2</sup>The Supervisor Cluster that is configured with VMware NSX-T Data Center, uses the software-based networks of the solution as well as an NSX Edge load balancer to provide connectivity to external services and DevOps users. The Cloudian Hyperstore pods are placed in a namespace, Hyperstore instance pods are placed in namespaces. For more information, see [Namespaces Networking](#).

VXLAN VTEP	Route based on the originating virtual port	Active	Active
------------	---	--------	--------

## Solution Validation

### Overview

With the vSAN Data Persistence platform, service operators can be deployed with a single click by the VI admin straight from the vSphere UI. In brief, the platform operator brings up an operator of the service that is enabled from vSphere UI, and the service can register a vCenter UI plugin which offers a seamless user experience to manage these services right within vSphere.

The platform operator creates optimized storage policies for this service, which can be used to deploy service instances. Once the service operator is up, developers can provision and scale instances of the service either through UI or Kubernetes APIs in a self-service manner.

Leveraging VMware integration, Cloudian HyperStore services can be deployed in a single click:

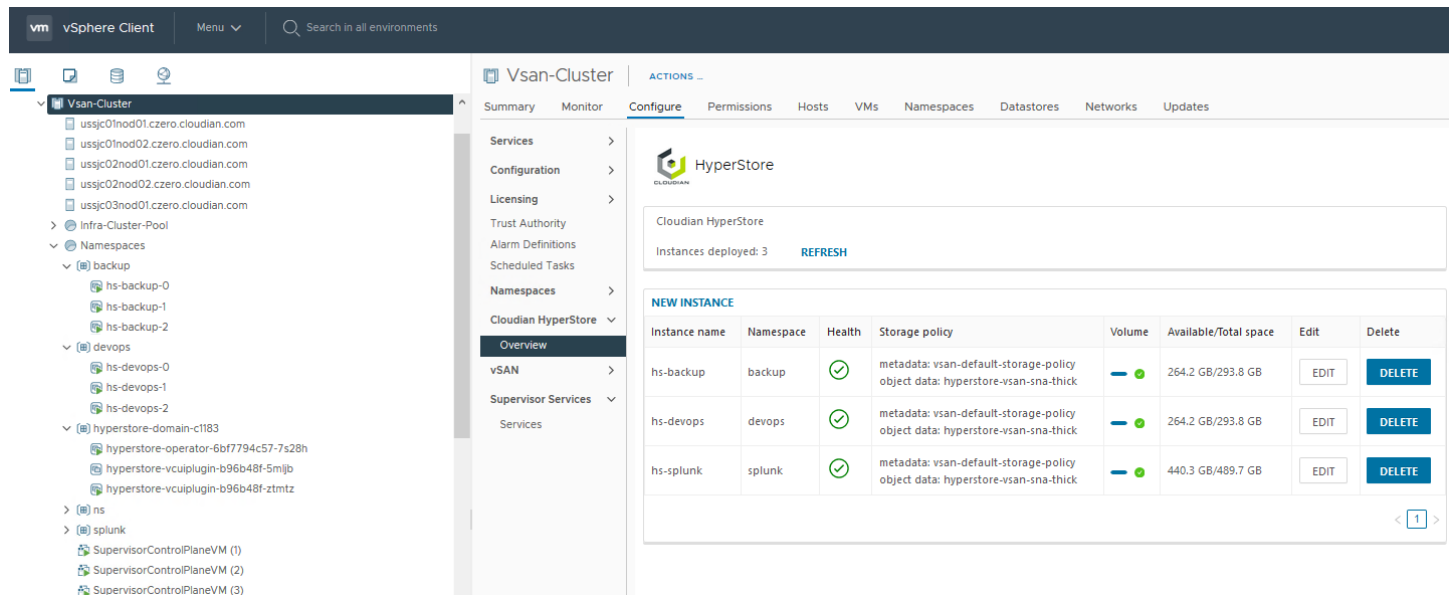
The screenshot shows the vSphere UI configuration page for a compute cluster. The left sidebar contains a navigation menu with categories like Licensing, Namespaces, vSAN, and Supervisor Services. The 'Services' option under Supervisor Services is selected. The main panel shows the 'Services' tab with an 'ENABLE' button and a table of services.

Service	Status	Current Version	Desired Version
Cloudian HyperStore	Disabled	--	v1.0.0
MinIO	Disabled	--	1.0
Velero vSphere Operator	Disabled	--	--

3 Items

HyperStore instance can be deployed quickly via UI:

**Hosts and Clusters -> Compute Cluster -> Configure -> Cloudian HyperStore – Overview -> New Instance**



## Deploying Cloudian HyperStore using vSAN Direct Configuration

HyperStore manages persistent storage and persistent volume claims. HyperStore issues commands to vSAN to create persistent volumes that map back to the HyperStore pods. A persistent volume exists on a single hard drive. Data protection is handled by HyperStore.

HyperStore storage policies are set with Replica Factor (RF) or Erasure Coding (EC). The default setting includes RF3 and EC4+2 in HSC v1.0 depending upon the number of nodes.

Using a small number of persistent volumes will limit performance. As PVs are directly related to the number of drives in the ESX host, having more drives in the host will significantly improve performance.

## Deploying Cloudian HyperStore using vSAN-SNA

HyperStore can leverage the vSAN-SNA storage policy and provide data protection based on HyperStore storage policy.

The default HyperStore storage policies are RF3 and EC4+2.

## Test Tools

We used the following monitoring tools and benchmark tool in the solution testing:

- Monitoring tools

### vSAN Performance Service

[vSAN Performance Service](#) is used to monitor the performance of the vSAN environment, using the vSphere web client. The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of problems.

### vSAN Skyline Health Check

[vSAN Skyline Health Check](#) delivers a simplified troubleshooting and monitoring experience of all things related to vSAN. Through the vSphere web client, it offers multiple health checks specifically for vSAN including cluster, hardware compatibility, data, limits, physical disks, and HyperStore instance health.

### vCenter UI plugin for Cloudian HyperStore

While vSAN Skyline Health Check provides a quick health overview of all HyperStore instances when the UI plugin is enabled, the plugin also allows more in-depth examination into the system stats and monitoring of each instance within the vCenter interface.

- Object storage performance tool

### Cosbench

[Cosbench](#) is a benchmarking tool to measure the performance of Cloud Object Storage services such as Cloudian HyperStore.

### Failure and Maintenance Scenarios

HyperStore responds automatically to events associated with the nodes on which a HyperStore instance is running.

Any failure or maintenance scenarios will change the cluster's health status from green to yellow or red.

The overall HyperStore cluster health is displayed and can be monitored via:

### Hosts and Clusters -> Compute Cluster -> Monitor -> Cloudian HyperStore – Overview

Instance name	Namespace	Health	Storage policy	Volume	Available/Total space	Alarms
hs-backup	backup	✓	metadata: vsan-default-storage-policy object data: hyperstore- <i>vsan-sna-thick</i>	264.2 GB/293.8 GB	No Alarms	
hs-devops	devops	✓	metadata: vsan-default-storage-policy object data: hyperstore- <i>vsan-sna-thick</i>	264.2 GB/293.8 GB	No Alarms	
hs-splunk	splunk	✓	metadata: vsan-default-storage-policy object data: hyperstore- <i>vsan-sna-thick</i>	440.3 GB/489.7 GB	No Alarms	

The table below lists some common failure or maintenance scenarios. Full details of HyperStore failover behavior can be found in the [Cloudian HyperStore-C Quick Start Guide](#).

Failure Type	Spare Component Available	Description	Note
Host Decommission	Yes	Host entering Maintenance Mode with “Full Data Migration” for decommission	HyperStore makes sure that the data for every HyperStore instance on the host is moved to a spare host.
	No		If there is no spare host available, HyperStore will reject the Enter Maintenance Mode (EMM) request and generate an alert.
Host Maintenance	N/A	Host entering Maintenance Mode with “Ensure Accessibility” for temporary maintenance purposes, or for	If the host exits Maintenance Mode before the rebuildTimerEMM period expires, the HyperStore pod recovers on that host and a data repair is automatically triggered.
	Yes		If the host has not yet exited Maintenance Mode and the rebuildTimerEMM period expires, and if terminating the pod will not lead to data unavailability (for example, the HyperStore instance health will be YELLOW after terminating the Pod, not RED), then

		implementing a rolling upgrade where every host goes down in sequence	HyperStore starts a new pod on a spare host that is not being used by that instance yet, creates PVs, and runs data repair.
	No		If the host has not yet exited Maintenance Mode and the rebuildTimerEMM period expires, but there is no spare host available, then HyperStore does nothing. In other scenarios in which HyperStore cannot allow the EMM to proceed -- such as if doing so would put the HS instance into RED health status -- HyperStore will generate an alert.
Host Failure	Yes	Host failure with spare host	When there is a failed host, HyperStore will attempt to start the pod on a spare host that is not being used by that instance yet, create PVs, and run data repair.
	No	Host failure with no spare host	If there is no spare host available, HyperStore does nothing except generate an alert.
Pod Stuck in Pending Status	Yes	Stuck pod with spare host	If a HyperStore pod stays in Pending status for a duration of time longer than the rebuildTimerEMM period, HyperStore will start the pod on a spare host, create PVCs, and run data repair.
	No	Stuck pod with no spare host	If there is no spare host available, HyperStore does nothing except generate an alert.
Drive Failure	Yes	Drive failure with spare drive on the same host	If a drive fails and there is an extra disk on the same host with enough available capacity, HyperStore deletes the inaccessible PVC, recreates the PVC using the extra disk on the same host, and restarts the pod.
	No	Drive failure with no spare drive on the same host	If there is no extra disk on the same host with enough capacity, HyperStore generates an alert. After the disk is replaced by the admin and a corresponding PV is created, then the PVC will be recreated on the same host and available to the pod.
			If the administrator does not replace the disk and there is a spare host with enough capacity, the administrator can trigger an EMM with "Full Data Migration" on the host with the failed disk. This will cause the pod to be rebuilt on a spare host.

## Best Practices

In this solution, the following recommendations provide the best practices and sizing guidance to run Cloudian HyperStore on VMware Cloud Foundation.

### VMware Cloud Foundation Infrastructure

Follow the general sizing guide: [Cloud Foundation Kubernetes Sizing Guide](#)

### Cloudian HyperStore

Refer to “Minimum Hardware Configuration” in table 1 for ESXi host sizing requirements.

When creating a new HyperStore instance, the following parameters specified in the wizard affect the overall sizing of compute and storage requirements:

- **Namespace:** The target deployment namespace for the cluster. This namespace needs to be pre-created in vCenter and assigned the appropriate vSAN storage policies and access permissions to users who can create HyperStore instance in it.
- **CPU required:** Number of vCPUs associated per pod. 16 vCPUs as minimum, enlarge vCPU based on workload to keep CPU usage under 80%.
- **Memory required:** RAM per pod. Minimum of 64GB is recommended.
- **Number of nodes:** Number of nodes (pods) in a HyperStore cluster. It cannot exceed the number of physical hosts in the vSAN cluster. The larger clusters provide better performance and data efficiency.
- **Metadata storage class:** Set the associated storage class for object metadata. Typically set to SSD or NVMe based class.
- **Object data storage class:** Set the associated storage class of object data. This class is the bulk of where HyperStore data is stored, so vSAN Direct or SNA policy is preferred.

- **Metadata volume size:** The amount of capacity dedicated to object metadata. 100TB capacity supports approximately 68 million objects.
- **Total object data storage size:** The amount of RAW capacity for object data in the cluster. To calculate the usable capacity, refer to *HyperStore Storage Policies* below.
- **Object Data volume size:** The size of PV the nodes consume to meet the total object data storage size value. Because each persistent volume exists on a single disk drive, this value should not exceed the unused capacity of the drive.

The screenshot shows the Vsan-Cluster management console. The left sidebar contains navigation links: Summary, Monitor, Configure (selected), Permissions, Hosts, VMs, Namespaces, Datastores, Networks, and Updates. The main content area displays the HyperStore configuration page. A 'NEW INSTANCE' dialog is open, showing the following fields:

- Name: ds-devops2
- Namespace: devops
- CPU required: 16
- Memory required: 64
- Number of nodes: 3
- Metadata storage class: vsan-default-storage-policy
- Object data storage class: hyperstore-nsa-thick
- Metadata volume size: 30
- Total object data storage size: 100
- Object data volume size: 100

The background table shows existing HyperStore instances:

Instance name	Namespace	Health	Storage policy
hs-backup	backup	✓	metadata: vsan-default-storage-policy object data: hyperstore-nsa-thick
hs-devops	devops	✓	metadata: vsan-default-storage-policy object data: hyperstore-nsa-thick
hs-splunk	splunk	✓	metadata: vsan-default-storage-policy object data: hyperstore-nsa-thick

## HyperStore Storage Policies

**Note:** The HyperStore storage policy is unrelated to the vSphere/vSAN storage policy

HyperStore storage policies are ways of protecting data so that it is durable and highly available to users. The HyperStore system lets you pre-configure one or more storage policies. Two options exist for creating storage policies. A default storage policy is created based on the number of nodes in a cluster at the creation time (for example: 3 nodes=RF3, 6 nodes=EC4+2). Additional policies can be created in the Cloudian Management Console.

After a new storage bucket is created, users can choose which pre-configured storage policy to use to protect data in that bucket.

Users cannot create buckets until you have created at least one storage policy.

For each storage policy that you create, choose from either of two data protection methods:

- **Replication** — With replication, a configurable number of copies of each data object are maintained in the system, and each copy is stored on a different node. For example, with 3X replication 3 copies of each object are stored, with each copy on a different node.
- **Erasur coding** — With erasure coding, each object is encoded into a configurable number (known as the "k" value) of data fragments plus a configurable number (the "m" value) of redundant parity fragments. Each of an object's "k" plus "m" fragments is unique, and each fragment is stored on a different node. The object can be decoded from any "k" number of fragments. To put it another way: the object remains readable

even if the "m" number of nodes are unavailable. For example, in a 4+2 erasure coding configuration (4 data fragments plus 2 parity fragments), each object is encoded into a total of 6 unique fragments which are stored on 6 different nodes, and the object can be decoded and read so long as any 4 of those 6 fragments are available.

In general, erasure coding requires less storage overhead -- the amount of storage consumption above and beyond the original size of the stored objects, to ensure data persistence and availability -- than replication. Put differently, erasure coding is more efficient in utilizing raw storage capacity than is replication.

#### Number of Nodes and Storage Policy

The number of nodes in a HyperStore cluster will affect the available storage policy. A 3-node cluster can only support Replica Factor (RF) 3. A cluster with 6 nodes can support RF3 as well as an Erasure Coding scheme of 4+2. The larger the cluster, the more efficient the EC scheme can be.

Data protection efficiency is calculated as a ratio of usable capacity to raw capacity as follows:

- 3 nodes: RF3 = 33%,
- 6 nodes: EC4+2 = 67%,
- 8 nodes: EC6+2 = 75%,
- 10 nodes: EC8+2 = 80%

## Use Cases

Cloudian HyperStore object storage is used for a broad range of use cases and for any traditional or modern application that benefits from S3-compatible storage. Cloudian HyperStore can be deployed and shared for multiple application workloads and use cases. For this document, we will highlight four key use cases:

#### Use Case 1: Cloud Native Apps

Proven VMware HCI platform supporting modern applications and data, with Cloudian HyperStore as scalable datastore for container-based application data. Public cloud-like storage access for flexibility and agility. S3-APIs for data and Kubernetes APIs for management. Start small and grow non-disruptively, as needs evolve. Secure multi-tenancy with encryption (data-at-rest and data-in-flight), WORM, IAM/RBAC, and secure shell. Eliminate resource silos, optimize storage economics, and simplify operation. Meet SLAs with confidence and respond to the needs of the business faster.

#### Use Case 2: Data Protection

Proven VMware HCI platform with Cloudian HyperStore as scalable backup/archive and DR target for both traditional (for example, SharePoint, SAP, SQL Server, and Oracle) as well as modern, container-based applications. Seamlessly compatible with S3 data protection solutions such as Veeam, Veritas, Rubrik, Commvault, Dattrium, and Velero. Data protection capabilities can include S3 Object Lock enabled protection for data immutability, important to ensure tamper-proof storage and protect against data loss due to cyber-attacks or rogue data breaches. Also support data retention and compliance needs. Policy-based replication to remote site or cloud, 100X faster access than tape and up to 70% less cost than legacy backup appliances. Per workload storage, backup, and management.

#### Use Case 3: Big Data Apps

A single platform for enterprise applications and data, Cloudian HyperStore as scalable data lake for applications requiring massive datastore (hundreds-to-thousands of terabytes). Examples include Enterprise Splunk, running Splunk SmartStore indexers with vSphere, hot bucket with vSAN, and warm and cold buckets with Cloudian HyperStore. Ingest GBs of data per day with quick searching across entire datasets, with up to 60% lower storage costs. Self-contained deployment for data security and privacy with a single global namespace, searchable across the entire dataset. Other application examples include Elasticsearch, Hadoop as well as Artificial Intelligence, Machine Learning, and other data analytics applications.



#### Use Case 4: File Tiering

Cloudian HyperStore as lower cost capacity storage for aging primary data. Typically, 60% of files are not accessed after 12 months. Nonetheless, they consume expensive storage capacity and data backup resources every day. Cloudian HyperStore integrates with data lifecycle management solutions and enable immediate reclaim of that space by offloading this infrequently accessed data from primary storage to a more cost-effective object storage pool to improve and balance infrastructure economics.

## Conclusion

Cloudian HyperStore on VMware Cloud Foundation with Tanzu is ready to provide S3 object storage services for your private or hybrid cloud infrastructure. New efficiencies and savings can be achieved by leveraging the ability to mix modern cloud applications with traditional applications using a single shared storage environment. A combination of VMware Cloud Foundation, VMware Tanzu, and VMware vSAN allows for the extension of Cloudian HyperStore's ease of deployment, high scalability, and strong security to any S3 enabled application.

There are many benefits of this solution. An enterprise grade S3 object storage service based on vSAN Data Persistence platform allows for improved storage efficiency. Automated deployment through Kubernetes makes installation and scaling easy. Service-aware infrastructure operations help minimize the impact of disruptions and infrastructure maintenance operations. Capacity and health monitoring is improved via the vSphere interface. Finally, CAPEX and OPEX savings can be realized by utilizing Cloudian data policies and consolidating traditional and cloud workflows on a single converged platform.

## Reference

- [VMware Cloud Foundation](#)
- [VMware vSphere](#)
- [VMware vSAN](#)
- [VMware NSX Data Center](#)
- [VMware vRealize Suite](#)
- [Cloudian HyperStore](#)

## About the Author

This paper was co-authored by the following people:

- Ka Kit Wong, Staff Solutions Architect in the Solutions Architecture team of the of the Cloud Platform Business Unit in VMware
- Ting Yin, Senior Solutions Architect in the Solutions Architecture team of the of the Cloud Platform Business Unit in VMware
- Henry Golas, Director of Technology, Canada, at Cloudian

The following reviewers also contributed to the paper contents:

- Siddharth Agrawal, Product Management, at Cloudian
- Grant Jacobson, VMware Alliance Director, at Cloudian
- Scott Ekstrom, Technical Marketing Engineer, at Cloudian



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 [www.vmware.com](http://www.vmware.com).

Copyright © 2021 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at [vmware.com/go/patents](http://vmware.com/go/patents). VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: vmw-wp-tech-temp-word-102-proof 5/19

©2021 Cloudian, Inc. Cloudian, the Cloudian logo, HyperFile, HyperScale, and HyperStore are registered trademarks or trademarks of Cloudian, Inc. All other trademarks are property of their respective holders. SB-vSAN-0920