

Implementing SMB File Services on Cisco HyperFlex Using NexentaStor

June 2019

Contents

Executive summary	3
Audience	3
Purpose of this document	3
Introduction to Cisco HyperFlex Systems	3
Implementing SMB file services on Cisco HyperFlex using NexentaStor Virtual Storage Appliance (VSA) ...	5
Introduction to SMB file services	5
NexentaStor VSA SMB file share	5
Nexenta HA (High-Availability) Cluster for File Services	6
NexentaFusion	6
NexentaStor VSA sizing considerations	7
VM configuration best practices	7
Deploy NexentaStor VSA on Cisco HyperFlex	7
Engineering validation	8
Solution Under Test (SUT) environment	8
Virtual Desktop Infrastructure (VDI) workload	9
Performance results	9
VDI-like workload results	9
Test results for Medium VM size	9
Test results for Large VM size	10
HyperFlex Connect performance chart	10
High availability testing	11
Conclusion	11

Executive summary

This paper describes the best practices to deploy NexentaStor Virtual Storage Appliance (VSA) and NexentaStor High Availability Solutions Software–based file services on Cisco HyperFlex™. Server Message Block Protocol (SMB) file-sharing solutions can be architected in several ways; this paper addresses a NexentaStor VSA, which is a single Virtual Machine (VM) configuration with a high-availability option. Cisco HyperFlex is a purpose-built hyperconverged infrastructure (HCI) for private and hybrid cloud data centers. When deploying hyperconverged solutions, there are no prerequisite operations, such as Logical Unit Number (LUN) provisioning, Host Bus Adapter (HBA) compatibility checks, or Fibre Channel (FC) LUN zoning operations. Cisco HyperFlex delivers extreme performance, VM density, and a wide variety of powerful data management features for an enterprise-scale private cloud environment.

Using NexentaStor provides an extremely efficient means of getting both the block storage already available from Cisco HyperFlex and Network-Attached Storage (NAS) features by merely adding another virtual machine. A VM hosting NexentaStor provides all of the NAS capabilities while consuming very little overhead from the HyperFlex nodes.

The solution described in this paper has been tested by Cisco in our lab only for assisting with best practices when used with Cisco HyperFlex. The licensing and support of SMB file services using NexentaStor is handled directly by Nexenta.

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers deploying SMB file services on Cisco HyperFlex. External references are provided wherever applicable, but readers are expected to be familiar with Nexenta-specific technologies and services, infrastructure concepts, networking connectivity, and the security policies of the customer's installation.

Purpose of this document

This document describes the steps required to deploy NexentaStor VSA, NexentaStor high-availability solutions and configure SMB file services using Nexenta Fusion on Cisco HyperFlex Systems. The document is based on all known best practices using the software, hardware, and firmware revisions specified in the document. As such, recommendations and best practices can be amended in later versions.

Introduction to Cisco HyperFlex Systems

Cisco HyperFlex Systems provide an all-purpose virtualized server platform, with hypervisor hosts, networking connectivity, and virtual server storage across a set of Cisco UCS® HX-Series x86 rack mount servers. Legacy data center deployments have relied on a disparate set of technologies, each performing a distinct and specialized function, such as network switches connecting endpoints and transferring Ethernet network traffic, and Fibre Channel (FC) storage arrays providing block-based storage via a dedicated Storage Area Network (SAN). Each of these systems has unique requirements for hardware, connectivity, management tools, operational knowledge, monitoring, and ongoing support. A legacy virtual server environment was often divided up into areas commonly referred to as silos, within which only a single technology operated, along with their correlated software tools and support staff.

Silos could often be divided between the x86 computing hardware, the networking connectivity of those x86 servers, SAN connectivity and storage device presentation, the hypervisors and virtual platform management, and finally the guest VMs themselves along with their OS and applications. This model proves to be inflexible and difficult to navigate, and is susceptible to numerous operational inefficiencies.

Hyperconverged infrastructures coalesce the computing, memory, hypervisor, and storage devices of servers into a single platform for virtual servers. There is no longer a separate storage system, because the servers running the hypervisors also provide the software-defined storage resources to store the virtual servers, effectively storing the virtual machines on themselves. Nearly all the silos are now gone, and a hyperconverged infrastructure becomes something almost completely self-contained, simpler to use, faster to deploy, easier to consume, yet still flexible and with very high performance. Many hyperconverged systems still rely on standard networking components, such as onboard network cards in the x86 servers, and top-of-rack switches. Cisco HyperFlex combines the convergence of computing and networking provided by Cisco UCS, along with next-generation hyperconverged storage software, to uniquely provide the compute resources, network connectivity, storage, and hypervisor platform to run an entire virtual environment, all contained in a single uniform system.

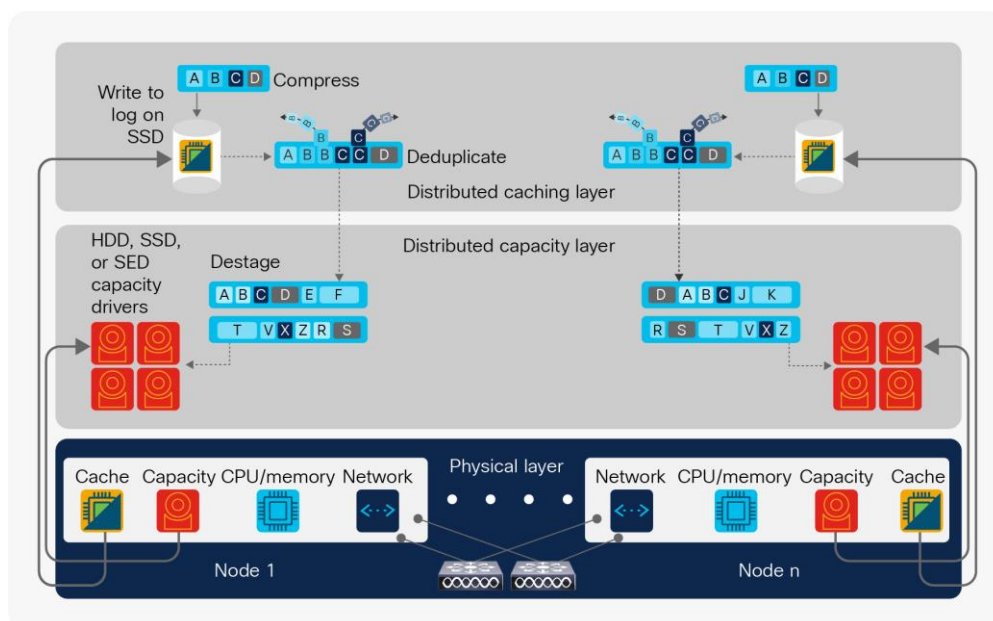
Some key advantages of hyperconverged infrastructures are the simplification of deployment and day-to-day management operations, and increased agility, thereby reducing operational costs. Since hyperconverged storage can be easily managed by an IT generalist. This can also, going forward, reduce technical debt that is often accrued by implementing complex systems that need dedicated management teams and skill sets.

Cisco HyperFlex supports VMware ESXi and Microsoft Hyper-V hypervisors.

For more information, see the [Cisco HyperFlex 3.5 for Virtual Server Infrastructure with Microsoft Hyper-V and Cisco HyperFlex 3.0 for Virtual Server Infrastructure with VMware ESXi](#).

Cisco HyperFlex provides a fully contained virtual server platform, with compute and memory resources, integrated networking connectivity, a distributed high-performance log-based file system for VM storage, and the hypervisor.

Figure 1. HyperFlex Architecture



Implementing SMB file services on Cisco HyperFlex using NexentaStor Virtual Storage Appliance (VSA)

This solution guide describes how to implement an SMB file service using a NexentaStor VSA to provide a genuinely heterogeneous file system for organizations with geographically dispersed sites that must uphold data availability without sacrificing data access and performance.

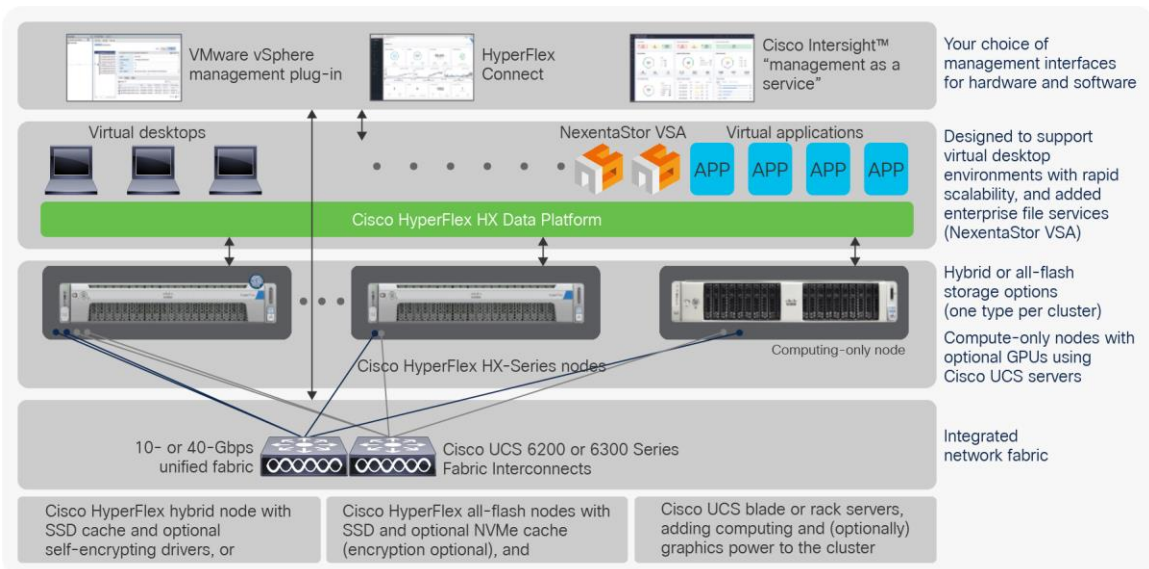
The solution described in this paper has been tested on Cisco HyperFlex with 4-node Cisco HyperFlex HXAF220c-M4S cluster all-flash disks for fast data access. Use this document for a baseline configuration when designing and planning your solution to better meet your requirements and budget.

Introduction to SMB file services

The Server Message Block (SMB) protocol is a network file-sharing protocol that allows applications on a computer to read and write to files and request services from server programs in a computer network. The SMB protocol can be used on top of its TCP/IP protocol or other network protocols. Using the SMB protocol, an application (or the user of an application) can access files or other resources at a remote server. This allows applications to read, create, and update files on the remote server. It can also communicate with any server program that is set up to receive an SMB client request.

NexentaStor VSA SMB file share

The NexentaStor software utilizes the existing storage contained in the Cisco HyperFlex systems by leveraging the physical elements that connect the disk media. VMware provides the hosting mechanism for running various Virtual Machine (VM) workloads, and NexentaStor operates in one of those VMs. Using variable capacities presented using VMware's VMDK disks, NexentaStor can create any mixture of NFS or SMB storage targets. To learn more about SMB file sharing in NexentaStor VSA, visit the [NexentaFusion UserGuide](#).



Cisco tested a single-instance and HA cluster SMB configuration. A NexentaStor VSA (which is a single VM configuration) was deployed using NexentaStor OVF and managed using NexentaFusion. This section describes the SMB file-sharing configurations that were tested with Cisco HyperFlex. For more information on how to create SMB file sharing using NexentaFusion, visit the following: [Nexenta website](#).

Nexenta HA (High-Availability) Cluster for File Services

The Nexenta HA (High-Availability) Cluster consists of two NexentaStor appliances and provides a storage volume-sharing service. You manage both systems actively for shared storage, although only one system provides access to a shared volume at a time. After you create a volume on one server and bring it under cluster control, then, when the high availability cluster detects a system failure, it transfers ownership of the shared volumes to the other server in the cluster pair.

The Nexenta HA Cluster provides server monitoring and failover. Protection of services, such as iSCSI, involves cooperation with other modules such as the iSCSI target plugin.

The Nexenta HA Cluster includes:

- **NexentaStor Appliances**

These run a defined set of services and monitor each other for failures. The HA Cluster connects these NexentaStor appliances through various communication channels, through which they exchange heartbeats that provide information about their states and the services that reside on them.

- **Cluster Service**

A transferable unit that consists of:

- Application startup and shutdown codes
- Network identity and appliance data

You can migrate services between cluster appliances manually, or automatically, if one appliance fails.

For more information on how to setup a Nexenta HA Cluster, visit the following [Nexenta High availability cluster guide](#).

NexentaFusion

NexentaFusion provides simplified scaling and management through a single pane for reporting, monitoring, analytics, and orchestration for your storage assets. From tens of terabytes to petabyte scale enterprise and service provider solutions, it complements Nexenta deployments by managing Nexenta Software-Defined Storage solutions. NexentaFusion simplifies deployment, operation, and scaling of Nexenta Software-Defined Storage. From an easy-to-use central console, Nexenta customers gain immediate insights into key attributes of their Nexenta infrastructure and benefit from advanced analytics to make informed decisions. The multisystem dashboard provides detailed status on multiple Nexenta arrays, on premises or in the cloud, and has the ability to drill into your complete Nexenta Software-Defined Storage infrastructure. A single screen provides quick insights and a starting point to drill down into any system within the infrastructure.

NexentaFusion is highly customizable, with an interactive analytics dashboard for historical analysis, service monitoring, and event troubleshooting to provide speed and simplification when getting intelligence from a Nexenta Software-Defined Storage deployment. NexentaFusion allows customers to collect extensive long term statistics on their systems. To learn more about NexentaFusion, visit the following: [NexentaFusion Guide](#).

NexentaStor VSA sizing considerations

The configuration and specification of the VM that is used for the NexentaStor VSA SMB file-sharing services plays a critical role in performance and the ability to handle the user's workload. The following table is intended to provide general guidelines; however, additional adjustments may be necessary to better fit your environment and workload. If you choose to deploy NexentaStor VSA using OVF, there are configuration profiles already available for the user.

VM size	vCPU	vRAM	User capacity
ROBO	1	8GB	8TB
Small	1	8GB	16TB
Medium	2	24GB	32TB
Large	2	32GB	64TB

For workload-specific sizing guidelines, use the [Cisco HyperFlex sizer](#).

VM configuration best practices

The following section describes best-practices considerations for configuring the VM:

- Use VMware's thick provision eager zeroed provisioning format for all virtual machine disks (VMDKs) used for SMB file sharing. This ensures that the disk is formatted and "zeroed out" prior to configuring the SMB share.

Deploy NexentaStor VSA on Cisco HyperFlex

To install NexentaStor VSA on Cisco HyperFlex, there are three simple steps to get started:

Step 1. Download and install the NexentaStor VSA VMware OVA file. Best practices can be found at the following:

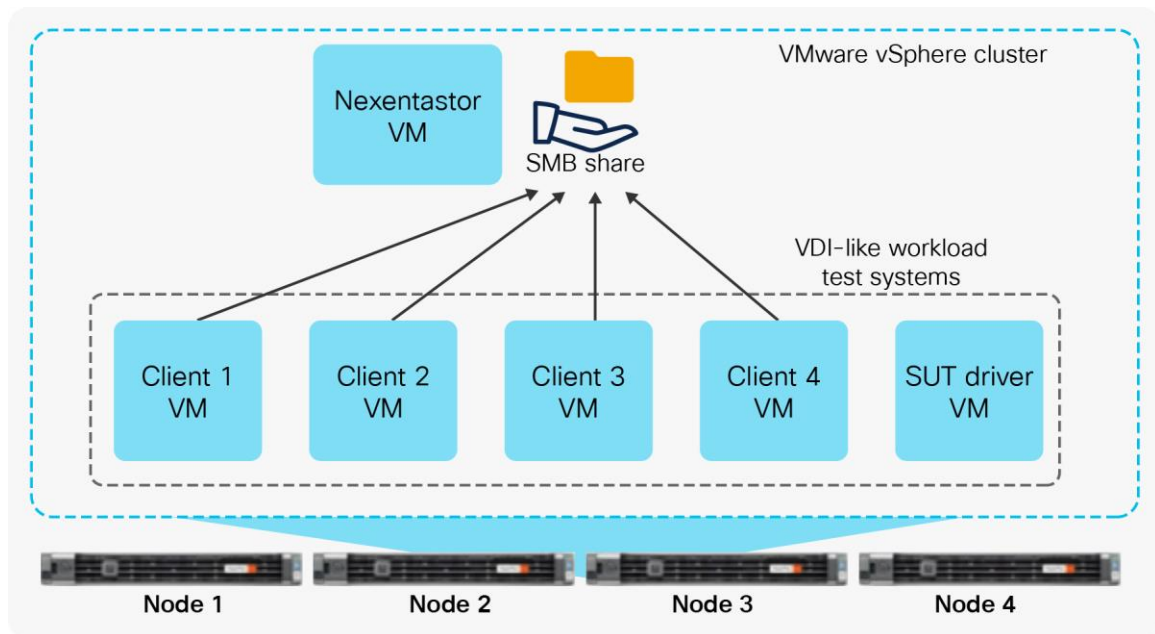
- Nexenta YouTube video, https://youtu.be/xc_cAgTAILE
- Nexenta Blog posting, <https://blog.nexenta.com/2018/07/25/taking-the-ez-pass-lane-to-a-hybrid-storage-cloud>

Step 2. Get and install a License Key. For trials, you can request a 45-day trial key using [NexentaStor-5-Enterprise Register](#)

Step 3. Start creating file systems and sharing with your users via the VMware vCenter Plugin or NexentaFusion.

Engineering validation

This section describes the testing and validation of the SMB file services option for Cisco HyperFlex.



Solution Under Test (SUT) environment

The testing and performance validation for this solution uses a common Virtual Desktop Infrastructure–like (VDI-like) workload to ensure that performance and reliability of the SMB file share access are met when running on Cisco HyperFlex.

Test environment	Version
Cisco HyperFlex (HX)	3.5.2b
NexentaStor Virtual Storage Appliance	5.2.0
NexentaFusion	1.2.1

The test environment consists of five VMs, with four as test clients to simulate a distributed workload. The SUT VMs are described below:

- VM1: SUT driver VM for controlling of the test clients
- VM2: Test client 1
- VM3: Test client 2
- VM4: Test client 3
- VM5: Test client 4

Test results reported in the section below for the Medium VM configuration profile are based on a NexentaStor VSA (a single VM configuration).

Test results reported in the section below for the Large VM configuration profile are based on a NexentaStor HA (High-Availability) Cluster.

Virtual Desktop Infrastructure (VDI) workload

This workload simulates a steady-state high-intensity knowledge worker in a VDI environment that uses full clones. This workload does not simulate a linked-clone environment. This is the behavior that was seen in traces between the hypervisor and storage when the VMs were running on ESXi, Hyper-V, KVM, and Xen environments.

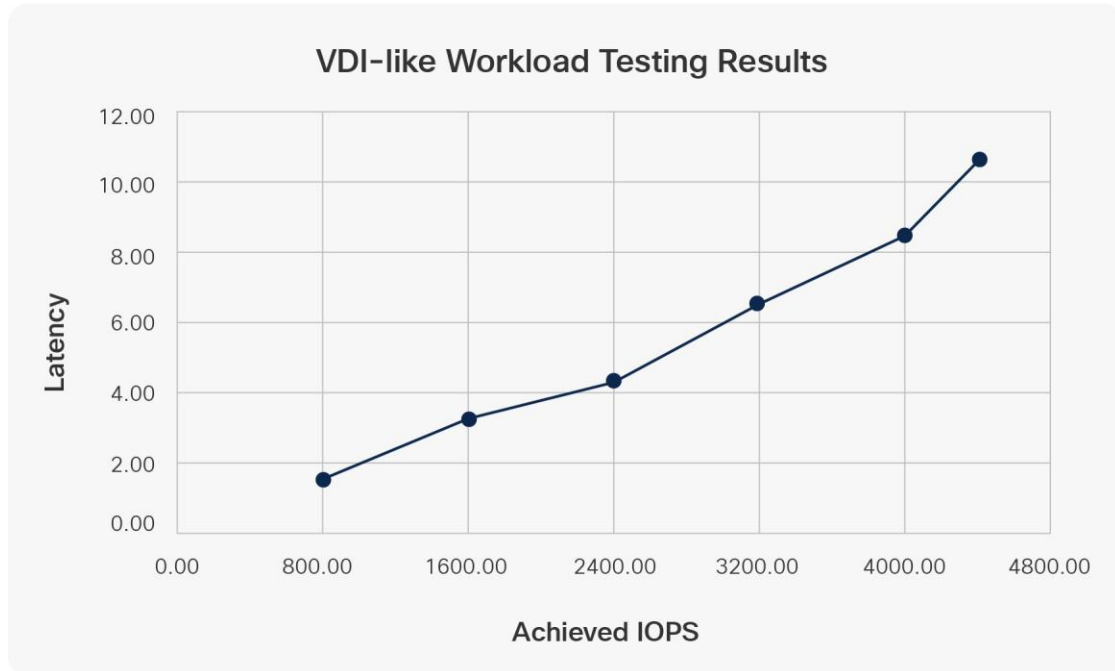
Performance results

The following sections describe the test results.


VDI-like workload results

The table below shows the results for a mixed I/O workload that simulates a home directory environment in a VDI-like environment when running on the SMB share for each of the VM sizes indicated earlier in the paper. Please note: these numbers are meant for validation only. For complete end-to-end sizing, use the Cisco [HyperFlexizer](#).

Test results for Medium VM size



Test results for Large VM size



HyperFlex Connect performance chart

The following charts show the backend storage performance during each test.

These performance charts show that the underlining storage scaled as the workload increased during the VDI-like workload stress testing. This ensures that the storage performance can keep up with the VDI workloads as user counts increase.

The following chart shows the backend performance on the HyperFlex Connect screen during the Medium VM test.



The following chart shows the backend performance on the HyperFlex Connect screen during the Large VM test.



High availability testing

The solution was tested to ensure that it meets Cisco's reliability requirements for use in a production environment. The test simulated a Cisco HyperFlex node failure while the system was under test to mimic an unplanned production outage; the VM experienced a 2ms outage and was able to recover without user intervention. The active workload testing on the VM was not affected during the outage event.

The following test scenarios were performed to ensure the reliability of the VM for SMB file sharing:

- Cisco HyperFlex node failure
- Network failure for the Cisco HyperFlex node and SMB VM
- Cisco HyperFlex node disk failure
- Enabling vMotion and checking if the VM can be migrated from the original to another node during active workloads
- NexentaStor VSA failure test to check if the VM restarts automatically on another node after the original node crashes

For further guidance and best practices of VMware HA (High Availability) and NexentaStor HA, visit the following [VMware KB article](#), [Nexenta website](#).

Conclusion

NexentaStor VSA with SMB file sharing on Cisco HyperFlex provides an enterprise-scale file system for corporate and remote-site data platforms that extends the use of a Cisco HyperFlex hyperconverged environment. The solution can be deployed as a NexentaStor VSA or as a NexentaStor High-Availability (HA) configuration using NexentaStor High Availability (HA) features to achieve a highly available and redundant SMB file-sharing solution to support business Service Level Agreements (SLAs).




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)