

Cisco HyperFlex iSCSI Overview

Version 1.0

12/4/2022

Contents

Document information	3
Intended use and audience	3
Legal notices	3
Overview	3
HyperFlex iSCSI use cases	3
Traditional applications	3
Cloud-native applications	4
Data center Infrastructure	4
HyperFlex iSCSI overview	5
Supported HyperFlex clusters	7
Supported initiator drivers	7
Supported HyperFlex topologies	8
License model for HyperFlex iSCSI	9
HyperFlex iSCSI workflow configuration	9
HyperFlex iSCSI snapshots	10
Bias statement	10
For more information	11

Document information

Document summary	Prepared for	Prepared by
First Release	Cisco Public	Joost van der Made

Intended use and audience

This document contains confidential material that is proprietary to Cisco. The materials, ideas, and concepts contained herein are to be used exclusively to assist in the configuration of Cisco® software solutions.

Legal notices

All information in this document is provided in confidence and shall not be published or disclosed, wholly or in part, to any other party without Cisco's written permission.

Overview

Cisco HyperFlex is a hyperconverged solution that combines compute, storage, and networking into a single scalable platform. Various needs, such as external storage and other protocols, require different workloads and applications.

The HyperFlex iSCSI feature can provide storage outside of the HyperFlex cluster and give applications the storage they need with the advantages, quality, and redundancy of the HyperFlex system.

This document will describe use cases of the HyperFlex iSCSI feature, an architectural overview of HyperFlex iSCSI, and configurations that are supported.

The HyperFlex iSCSI feature is available from Cisco HyperFlex HX Data Platform Release 4.5 and higher.

For more information about HyperFlex:

<https://www.cisco.com/site/us/en/products/computing/hyperconverged-infrastructure/index.html>.

HyperFlex iSCSI use cases

What is a typical application that uses iSCSI block storage devices? The following are some use cases and examples.

Traditional applications

Traditional applications are virtualized applications called Virtual Server Infrastructure (VSI).

Windows failover clustering

Windows failover clustering (FOC) provides servers working together for optimal availability. There is an active and standby server, and in case of a failure, the standby server can become the active server and use the HyperFlex iSCSI Logical Unit Number (LUN) connected to the failed server. This way, no data is lost.

You can find an example of a Windows FOC in the document: [Cisco HyperFlex All-NVMe Systems for Deploying Microsoft SQL Server 2019 Database with VMware ESXi](#).

Oracle Database and Oracle RAC (Real Application Clusters)

Databases are critical in today's digital world. Oracle Database is an enterprise database application with much power and scalability. Getting high availability for this database can be complicated, and it is a must to have the databases up and running in case of a failure event. All the data should be available all the time.

By using shared disk access across different Oracle RACs, virtual machines ensure that the database has high availability. The HyperFlex iSCSI feature accomplishes this without high license costs.

For more information about the implementation of this solution, please read the white paper: [Cisco HyperFlex All-NVMe Systems with iSCSI Support for Oracle Real Application Clusters White Paper](#).

Microsoft Exchange

Microsoft Exchange does not support NFS storage protocol. Configuring mailboxes and log files using HyperFlex iSCSI LUNs eliminate support restrictions. With this model, there are no shared physical disks.

HyperFlex iSCSI LUN Cloning for testing and development

HyperFlex ready clone is a feature that can quickly and rapidly clone VMs. The HyperFlex iSCSI LUN Cloning feature can clone a HyperFlex iSCSI LUN and create a new HyperFlex target IQN. The LUN clone is built on top of HyperFlex instant and space-efficient snapshots. Another initiator, such as a development environment server, can then connect to the cloned LUN and use the data for development purposes, because there is no loss during the cloning process. HyperFlex iSCSI LUN Cloning works with application consistency, which is a significant advantage of HyperFlex iSCSI LUN Cloning feature.

Cloud-native applications

Cloud-native applications are different from VSI, and they can scale quickly and develop in a Kubernetes environment.

Container Storage Interface (CSI)

Nowadays, persistent volumes are a must in containerized applications. Container Storage Interface (CSI) can provide the container with persistent volumes. HyperFlex has a CSI plug-in; thus, the containers in a HyperFlex cluster can use persistent volumes.

Cloning for testing and developing

Cloning of persistent volumes is part of the Kubernetes specifications. The advantages of cloning are in DevOps environments.

Raw-block volume mode

Containerized applications may need to read/write directly to a raw storage volume device. The prominent use cases are databases. When containerizing database applications that are highly latency-sensitive, some specialized applications may prefer accessing storage when raw-block volumes overuse filesystems to reduce overhead latency.

Data center Infrastructure

Data centers cannot afford downtime. In a stateless computing environment, there are no dependencies on local hardware. With the best of HyperFlex, the servers can use iSCSI LUNs as networked boot devices and remain stateless.

HyperFlex iSCSI overview

This section explains the Cisco HyperFlex HX Data Platform's critical features before diving into the HyperFlex iSCSI overview.

- HyperFlex is a system that has a Replication Factor (RF) of two or three. Every data block is written multiple times across the cluster on separate nodes.
- Data is striped across all nodes, improving the performance of the cluster.
- Because of the log-structured file system, HyperFlex can use spinning-, flash-, and NVMe drives with stable latency.
- Applications depend on the latency, and if the latency is consistent, which is the case on a HyperFlex cluster, the application runs smoothly.



Figure 1.
HyperFlex benefits overview

The HyperFlex iSCSI feature uses all the key elements of HyperFlex. This feature eliminates all storage silos. Applications running inside or outside the hyperconverged solution can take advantage of HyperFlex as iSCSI storage.

Consolidate workloads

Besides virtual machines and desktops, you can provision iSCSI LUNs from the same storage pool on a HyperFlex cluster to service applications requiring block storage. Even exposing storage to "the outside world" is possible; external applications can consume storage using iSCSI from the HyperFlex cluster without running on the cluster itself.

Simplicity

HyperFlex consolidates and simplifies storage, compute, and network configuration and management. The HyperFlex iSCSI network creation workflow automatically configures the HyperFlex iSCSI network on Cisco UCS Fabric Interconnects. When an iSCSI initiator requires a connection to the HyperFlex iSCSI Target, discovery can be performed using the HyperFlex iSCSI Cluster IP as a centralized portal; thus, the HyperFlex system will be responsible for load balancing and failure handling. No additional configuration is required on the initiator.

Efficient, scale-out storage

As support for the iSCSI protocol is implemented at the top of the HyperFlex I/O stack, iSCSI workloads benefit from the rich data services available from the HyperFlex log-structured filesystem, such as inline compression and deduplication, and space-efficient instant snapshots and clones. HyperFlex IOvisor software distributes I/O from iSCSI workloads evenly across the cluster, in a similar fashion as it does for VM workloads, so that there are no performance hot spots.



Figure 2.
HyperFlex iSCSI feature overview

iSCSI is a protocol, and as seen in the use cases, some applications perform better with an iSCSI LUN.

The HyperFlex iSCSI feature is a day-2 operation. After installing HyperFlex, the iSCSI network should be configured first through HyperFlex Connect UI, APIs, or hxcli commands, before the HyperFlex cluster acts as a target for the initiators.

The HyperFlex iSCSI LUNs of the different targets are written, with a replication factor, across all nodes. There is no need to create datastores for the iSCSI LUNs; a default datastore is automatically created to contain all the iSCSI LUNs during the configuration of the iSCSI feature.

Supported HyperFlex clusters

HyperFlex data-center FI-based (non-stretched) clusters running VMware ESXi can enable the HyperFlex iSCSI feature. Other HyperFlex deployment types, such as stretched, DC-no-FI, and edge clusters, and clusters running non-ESXi hypervisors (Microsoft Hyper-V or Cisco Intersight Workload Engine) do not support iSCSI.

All hardware supported for a HyperFlex cluster can have the iSCSI feature enabled, and it does not depend on the types of storage in the HyperFlex cluster.

Table 1. Supported HyperFlex cluster

	HyperFlex iSCSI feature
HyperFlex Edge	X
HyperFlex ESXi	✓
HyperFlex DC-no-FI	X
HyperFlex Hyper-V	X
HyperFlex Stretched cluster	X

Supported initiator drivers

An initiator is an iSCSI driver on an Operating System (OS). The OS can be running on bare metal or any hypervisor, such as ESXi, Hyper-V, XenServer, KVM, and other hypervisors.

The OS initiator connects to the HyperFlex iSCSI Target, and then the OS can use the LUN as a drive.

The following operating systems are supported to connect to a HyperFlex iSCSI target:

Table 2. Supported OS initiators

OS initiators	HyperFlex iSCSI feature
Windows 2016/2019	✓
Oracle Linux 8	✓
Ubuntu 18.04 and 20.04	✓
RedHat Linux 7	✓
Open-iSCSI 2.0874	✓

Booting from a HyperFlex iSCSI target needs a hardware initiator, which provides a LUN to the server, making it possible to use this LUN as a boot drive.

Table 3 shows the list of supported hardware initiators.

Table 3. Supported hardware initiators

Hardware initiators	HyperFlex iSCSI feature
UCS VIC 1300 series	✓
UCS VIC 1400 series	✓

All other initiators are not supported.

Supported HyperFlex topologies

During the installation of HyperFlex, the installer automatically configures the network interface cards, virtual switches, and Quality of Service (QoS) settings on the Fabric Interconnects to simplify the network configuration. The HyperFlex iSCSI network configuration automatically configures all iSCSI network settings up to the northbound switches.

The best practice is to have a separate network for iSCSI traffic. If VLANs virtually separate this traffic, the best approach is to configure Quality of Service (QoS) on the switches for the iSCSI VLAN.

The OS with the initiator can be a virtual machine on the HyperFlex cluster or any other hypervisor located anywhere on the iSCSI network. Install the OS on a bare-metal server in the same Cisco UCS domain or anywhere on the iSCSI network.

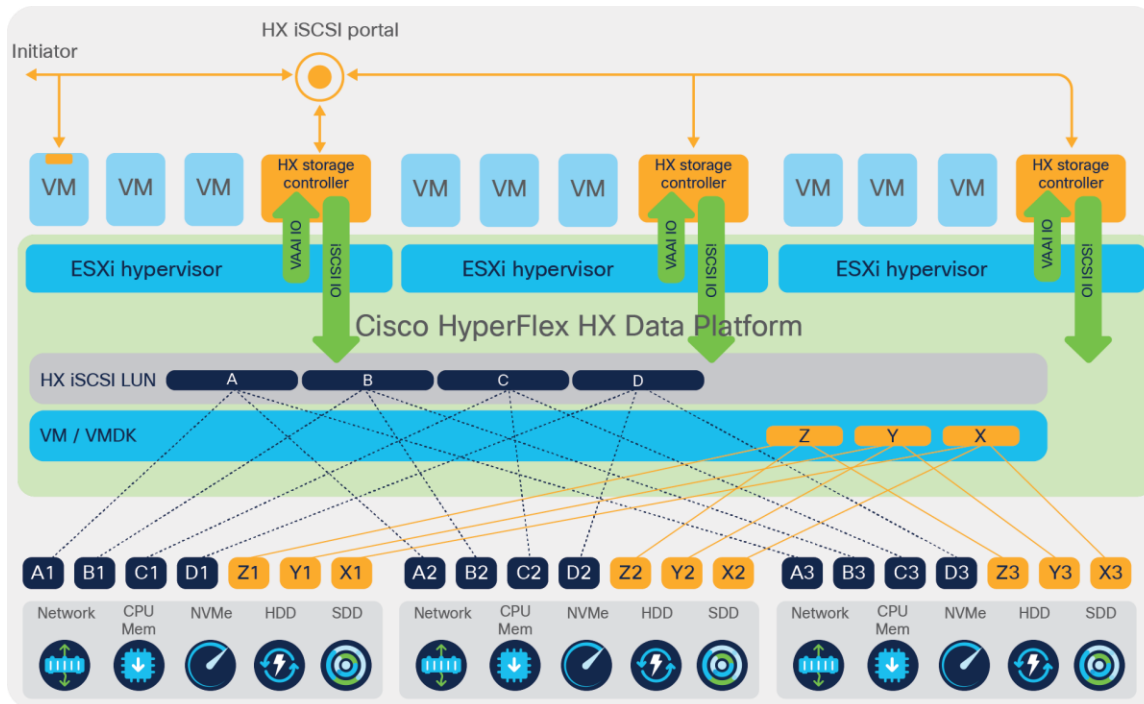


Figure 3. HyperFlex architecture diagram with iSCSI

The HyperFlex iSCSI cluster IP address (CIP) is a parameter in the HyperFlex iSCSI network configuration workflow. When an initiator discovers their targets through this HyperFlex iSCSI CIP, the HyperFlex cluster controls failover and load balances the nodes. The initiator will get another HyperFlex iSCSI node IP for the connection. HyperFlex writes the data across all nodes, and the number of copies depends on the RF.

The iSCSI load-balancing makes the setup of the HyperFlex iSCSI feature straightforward on the initiator side. If a HyperFlex node fails or the connected node is very busy, the initiator will signal that the iSCSI connection will be moving to another node. From the standpoint of the initiator, nothing will change. This configuration is the easiest to set up.

The initiator can also use direct login, which directly points to a HyperFlex iSCSI node. The drawback of this topology: HyperFlex cannot move the initiator to another node when a path fails.

Multi-path I/O (MPIO) is a driver on the OS that has multiple paths from the initiators to the same iSCSI target. The MPIO driver will handle path failover during a failure event.

License model for HyperFlex iSCSI

HyperFlex iSCSI is a HyperFlex Datacenter Advanced license feature on a HyperFlex ESXi cluster.

Enabling HyperFlex Boost Mode is a best practice for the HyperFlex iSCSI feature, and this does not require an additional license.

You can find more information about the HyperFlex Boost Mode and how to enable it in this document: [HyperFlex Boost Mode White Paper](#).

HyperFlex iSCSI workflow configuration

Configuring the HyperFlex iSCSI network is a day-2 operation, and therefore it is configured after the installation of HyperFlex. During the configuration of the iSCSI network, the UCS Manager, vSwitches, and HyperFlex controller VMs have configured automatically. Only upstream switches, if any, must be configured manually.

Each Initiator Group (IG) has its own HyperFlex Target. There are different orders for configuring a HyperFlex target, and here are the workflows.

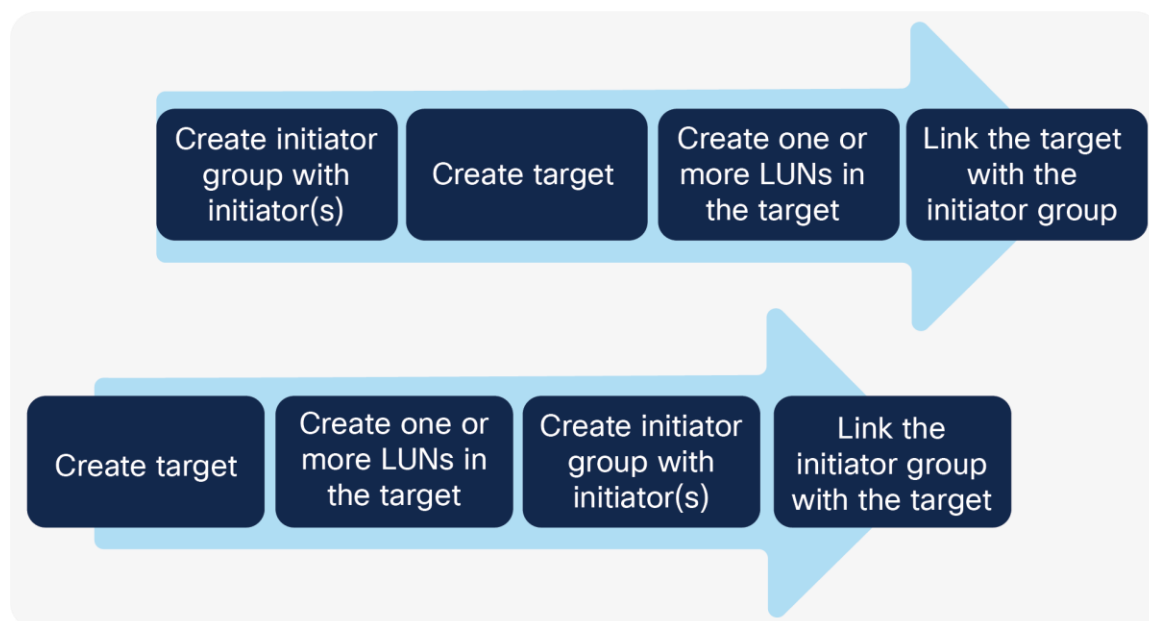


Figure 4.
HyperFlex iSCSI workflows

Detailed information about the HyperFlex iSCSI configuration is in the [Cisco HyperFlex Admin Guide 4.5 section Managing iSCSI](#).

HyperFlex iSCSI snapshots

In Cisco HyperFlex HX Data Platform (HXDP) Release 5.0(1a) and higher, it is possible to create consistency groups with HyperFlex LUNs and create snapshots with the same timestamp across all LUNs within the group using HyperFlex Data Platform REST APIs. These are crash-consistent snapshots, and HyperFlex does not handle application consistency. Quiescing and un-quiescing application I/O before and after group snapshot creation is the application's or the backup software responsibility.

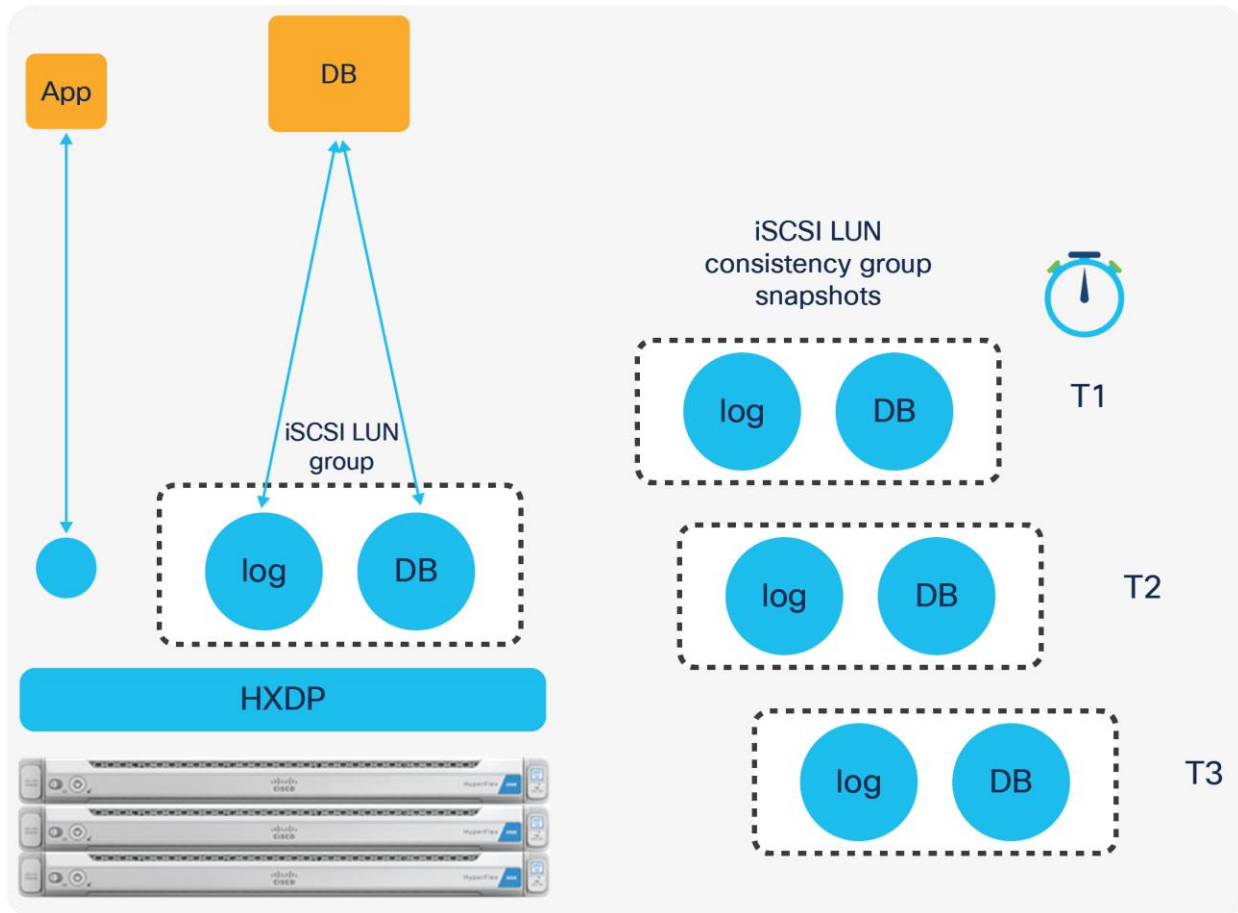


Figure 5.
HyperFlex iSCSI workflows

Bias statement

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

For more information

For additional information, please refer to the following resources

- <https://www.cisco.com/go/hyperflex>
- [HyperFlex 4.5 Administration Guide, Managing iSCSI](#)
- [Cisco HyperFlex All-NVMe Systems with iSCSI Support for Oracle Real Application Clusters White Paper](#)
- [Cisco HyperFlex 4.5 for Virtual Server Infrastructure with VMware ESXi](#)
- [Cisco HyperFlex All-NVMe Systems for Deploying Microsoft SQL Server 2019 Database with VMware ESXi](#)
- [HyperFlex Boost Mode White Paper](#)

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)