



Cisco HyperFlex with Veeam Availability Suite

Deployment Guide for Data Protection on Cisco HyperFlex Systems through Veeam Availability Suite 9.5 and Cisco UCS S3260 Storage Servers in Single Data Center

Last Updated: May 15, 2017



About the Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0809R)

© 2017 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary.....	6
Solution Overview	7
Introduction	7
Audience.....	7
Purpose of this Document.....	7
Solution Summary.....	7
Technology Overview	10
Cisco Unified Computing System	11
Cisco Unified Computing System Components	12
Cisco Fabric Interconnects	12
Cisco UCS S3260 Storage Server	13
Cisco HyperFlex HX-Series Nodes	15
Cisco HyperFlex HX220c-M4S Node.....	15
Cisco HyperFlex HX240c-M4SX Node	15
Cisco VIC 1227 MLOM Interface Card	16
Cisco HyperFlex Converged Data Platform Software.....	17
Cisco HyperFlex HX Data Platform Administration Plug-in.....	17
Cisco HyperFlex HX Data Platform Controller.....	18
Data Operations and Distribution.....	19
Veeam Availability Suite	21
Backup	21
Restore.....	22
Replication.....	23
Deployment Types	24
Failover and Failback.....	26
Backup Server	27
Components.....	27
Backup Proxy	31
Veeam Repository Sizing.....	32
Solution Design	34
Deployment Architecture.....	36
Deployment Hardware and Software.....	39
Software Versions.....	39
Configuration Guidelines.....	39
Network Switch Configuration.....	42
Physical Connectivity.....	42
Cisco Nexus Base.....	42
Set Up Initial Configuration.....	42

Cisco Nexus Switch Configuration	44
Enable Licenses	44
Set Global Configurations	45
Create VLANs	45
Add NTP Distribution Interface.....	45
Create Port Channels	46
Configure Port Channel Parameters	47
Configure Virtual Port Channels	47
Uplink into Existing Network Infrastructure	49
Cisco UCS S3260 Storage Server and HyperFlex Configuration	50
Cisco UCS Base Configuration.....	50
Perform Initial Setup of Cisco UCS 6248UP Fabric Interconnects	50
Cisco UCS Setup	52
Log in to Cisco UCS Manager.....	52
Upgrade Cisco UCS Manager Software to Version 3.1(2b).....	52
Anonymous Reporting.....	52
Add Block of IP Addresses for KVM Access.....	53
Synchronize Cisco UCS to NTP	54
Uplink Ports	55
Uplink Port Channels	56
Edit Chassis Discovery Policy	58
Server Ports	58
Server Discovery	59
HyperFlex Installation.....	61
Cisco UCS S3260 Configuration.....	63
Create Sub-Organization	63
Create Chassis Firmware Packages	64
Create Disk Zoning Policy.....	65
Setting S3260 Disk to Unconfigured Good	67
Create MAC Address Pools.....	67
Create UUID Suffix Pool.....	71
Create Server Pool.....	72
Create VLANs	73
Create Host Firmware Package	74
QoS Policy.....	75
Create Network Control Policy for Cisco Discovery Protocol	76
Create Power Control Policy	77
Create Server BIOS Policy	78
Create Maintenance Policy.....	79
Create Adaptor Policy.....	80

Create vNIC Templates.....	81
Create Disk Group Policy.....	85
Create Storage Profile.....	87
Create Chassis Profile Template.....	91
Create Service Profile Template.....	92
Create Chassis Profile.....	99
Associate Chassis Profile to S3260 Chassis	100
Create Service Profiles.....	102
Associate Service Profile to Server node of S3260 Chassis.....	103
Veeam Availability Suite 9.5 Installation	107
Install and Configure Windows 2016	107
Load Driver for S3260 RAID Controller	108
Update Cisco VIC Driver for Windows 2016.....	111
Update Intel ChipSet Driver for Windows 2016.....	115
Create NIC Teaming for Windows 2016.....	116
Create Disk Volume for Veeam Repository	118
Install Veeam Availability Suite 9.5	120
Configure Veeam Availability Suite 9.5.....	127
Validation.....	136
Validated Hardware and Software	139
Bill of Materials	141
References	144
About the Authors	145
Acknowledgements.....	145

Executive Summary

The Cisco HyperFlex™ Systems solution together with Veeam Availability Suite gives customers a flexible, agile, and scalable infrastructure that is protected and easy to deploy. Building on top of the Cisco HyperFlex HX Data Platform's built-in protection tools, Veeam Availability Suite expands the protection of your data with local and remote backups and VM-level replication. Today's data centers are heterogeneous, and most administrators want to avoid siloed data protection for each application or infrastructure stack. Customers need data protection to work across the enterprise and for recovery to be self-service, easy, and fast—whether the data is local or remote.

Cisco and Veeam have partnered to deliver a joint solution that enables backup, restore and replicate, virtualized workloads running on Cisco HyperFlex, utilizing Veeam Availability Suite deployed on Cisco UCS S3260 Storage Server. Ensure fast, reliable backup and recovery of virtual machines, critical data and applications with a data protection solution that provides enterprise availability, business agility, and operational efficiency. Veeam and Cisco's solution helps customers realize the full potential of virtualization and converged infrastructures by simplifying management to minimize risk, decrease downtime, and easily adapt to business demands. IT administrators can leverage policy-based controls for smarter data protection to recover the data they want, when they want it, enabling organizations to deploy a high performance, compatible solution that has been tested and validated by Cisco and Veeam experts.

The Veeam Availability Solution for Cisco UCS is a best-of-breed solution that is the perfect answer for customers requiring an advanced, enterprise-class data availability solution for their virtual environments that is simple to order, deploy and manage. In addition, it can easily expand over time as the need increases. It provides fast, flexible and reliable recovery of virtualized applications and data bringing virtual machine backup and replication together in a single solution with award-winning support.

This Cisco Validated Design (CVD), Data Protection for Cisco HyperFlex with Veeam Availability Suite, is a certified solution built on a modern architecture that delivers fast, reliable recovery, reduced total cost of ownership (TCO) and a better user experience, and addresses the challenge of delivering agile protection for Cisco HyperFlex platform. This solution uses Cisco components such as Cisco UCS Manager, Cisco Fabric Interconnect 6248UP, Cisco HyperFlex Data Platform, Cisco HyperFlex HX220c and HX240c nodes, Cisco Nexus 9000 series networking and Cisco UCS S3260 Storage Server.

A Cisco Validated Design (CVD) and pre-validated reference architectures facilitate faster, more reliable, and more predictable customer deployments:

- Each CVD has been extensively tested, validated, and documented by Cisco and partner experts
- CVD's minimize both integration and performance risks to ensure always-on availability in the enterprise

From design to configuration, instructions to bill of materials (BOMs) - CVDs provide everything businesses need to deploy the solutions in the most efficient manner.

Solution Overview

Introduction

Designed specifically for virtual environments, Data Protection for Cisco HyperFlex with Veeam Availability Suite is integrated with VMware vSphere, helping ensure consistent and reliable virtual machine recovery.

The Cisco HyperFlex solution delivers next-generation hyperconvergence in a data platform to offer end-to-end simplicity for faster IT deployments, unifying computing, networking, and storage resources. The Cisco HyperFlex solution is built on the Cisco Unified Computing System™ (Cisco UCS®) platform and adheres to a data center architecture supporting traditional, converged, and hyperconverged systems with common policies and infrastructure management. The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system delivering a wide range of enterprise-class data management and optimization services. This platform redefines distributed storage technology, expanding the boundaries of hyperconverged infrastructure with its independent scaling, continuous data optimization, simplified data management, and dynamic data distribution for increased data availability. This agile system is easy to deploy and manage, scales as your business needs change, and provides the first level of data availability. However, as with most systems, a second layer of protection that is equally agile is recommended. Veeam Availability Suite can meet this need.

Veeam is an industry leader within the data protection market. In the era of Digital Transformation, Veeam recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise by helping organizations meet today's service-level objectives, enabling recovery of any IT service and related applications and data within seconds and minutes. Veeam consistently pushes the envelope in bringing sophisticated backup and disaster recovery functionality to enterprises and cloud providers.

Veeam delivers efficient virtual machine (VM) backup and replication to dramatically lower the recovery time objective (RTO) and recovery point objective (RPO), for RTPO™ of <15 minutes for ALL applications and data. Veeam replication between HyperFlex clusters, both local and distributed, provides site-level DR. Veeam also provides backup and recovery at the VM- and item-level for instant recovery from more common, day-to-day problems. These isolated Veeam managed backups, stored on secondary storage, cloud or tape, allow organizations to meet both internal and external data protection and recovery requirements.

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 looking to provision backup and recovery of virtualized application on Cisco HyperFlex Clusters deployed across data centers or in several Remote Offices, across different geographies.

Purpose of this Document

This document elaborates on design, deployment and configuration procedures for backup and replication of virtualized server infrastructure on HyperFlex Cluster utilizing Veeam Availability Suite 9.5 and Cisco UCS S3260 as the backup target repository.

Solution Summary

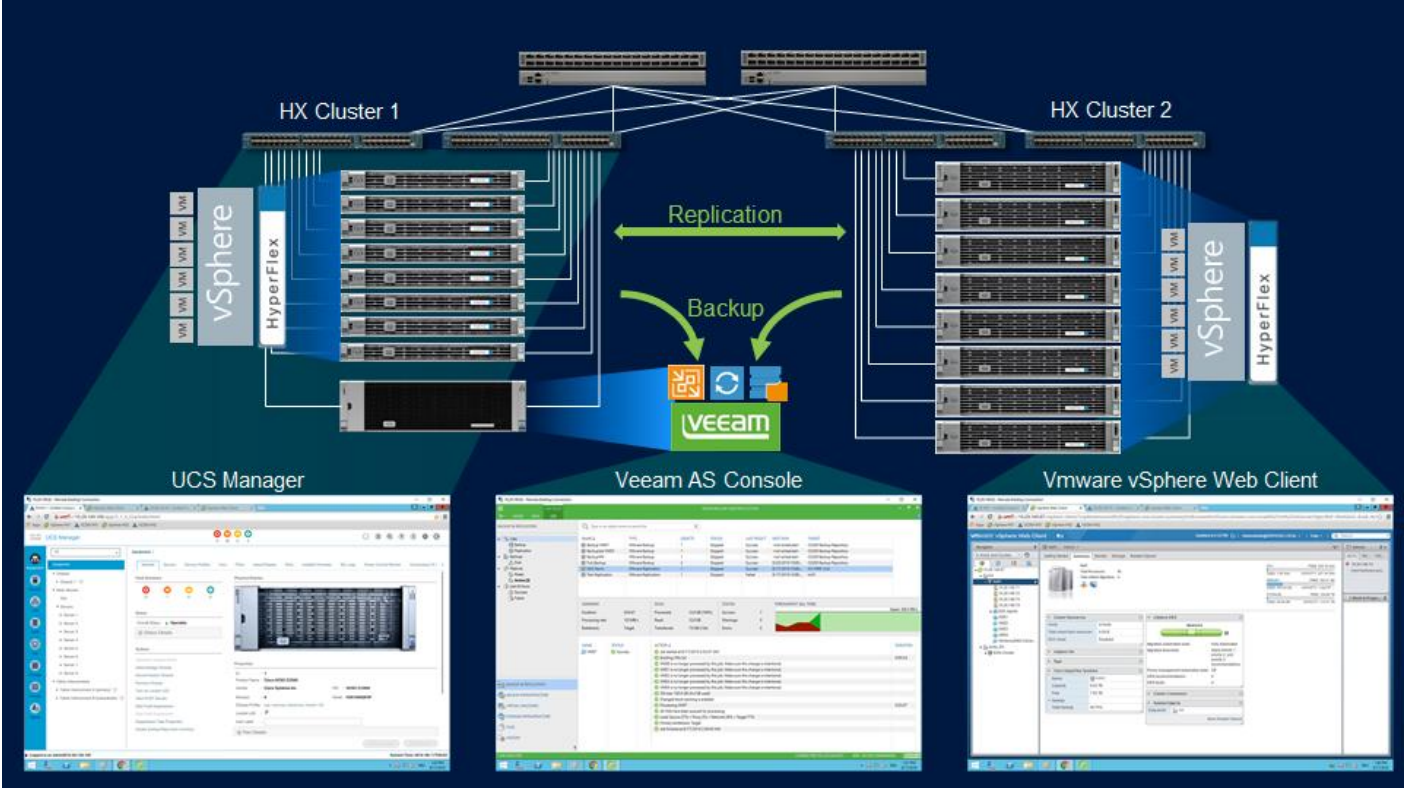
The solution for Cisco HyperFlex and Cisco UCS S3260 Storage Server with Veeam Availability Suite delivers reliable and fast Backup and Replication of application VMs residing on Cisco HyperFlex Cluster. The solution is flexible and can easily be extended across heterogeneous virtualized Data Center environments comprising of both converged and hyperconverged infrastructure. This solution can be accurately sized in accordance with present demands of enterprise deployments and thereafter can be scaled as per the future growth projections.

This design provides Backup and Replication of infrastructure VMs on Cisco HyperFlex cluster located in the same Data Center or Campus through Veeam Availability Suite. Veeam Availability Suite which comprises of Veeam Repository, Veeam Proxy and Veeam Backup Server all reside on a single Cisco UCS S3260 Storage Server which provides up to 600 TB of raw storage capacity. Replication of application VM is executed to a separate Cisco HyperFlex Cluster.

Figure 1 provides a high-level view of Cisco HyperFlex with Cisco S3260 Storage Server and Veeam Availability Suite and elaborates on the following:

- Replication of application VMs across Cisco HyperFlex Clusters through Veeam Availability Suite
- Backup of application VMs on Cisco S3260 Storage Server
- Management end points for Cisco HyperFlex, Cisco UCS S3260 Storage Server and Veeam Availability Suite

Figure 1 Cisco HyperFlex with Veeam Availability Suite and Cisco UCS S3260 Storage Server



Technology Overview

The Design guide for Cisco HyperFlex with Veeam Availability suite uses the following infrastructure and software components:

- Cisco Unified Computing System (Cisco UCS)
- Cisco HyperFlex Data Platform
- Cisco Nexus 9000
- Veeam Availability Suite 9.5
- Windows 2016 Datacenter Edition for Veeam Availability Suite

This Deployment guide uses the following models of above mentioned infrastructure components:

- Cisco UCS S3260 Storage Server
- Cisco UCS HX220c M4 Node
- Cisco UCS HX240c M4 node
- Cisco UCS 6200 Series Fabric Interconnects (FI)
- Cisco Nexus 9300 Series Platform switches

The other optional software and hardware components of this design solution are:

- Cisco UCS Central provides a scalable management platform for managing multiple, globally distributed Cisco UCS domains with consistency by integrating with Cisco UCS Manager to provide global configuration capabilities for pools, policies, and firmware. Cisco UCS Central platform eliminates disparate management environments. It supports up to 10,000 Cisco UCS servers (blade, rack, and Mini) and Cisco HyperFlex Systems. You can manage multiple Cisco UCS instances or domains across globally distributed locations.
- Veeam Enterprise Manager provides a single “pane of glass” management for a globally dispersed Backup and Replication environment. It collects data from the backup servers and enables you to run backup and replication jobs across the entire back-up infrastructure, edit, and clone jobs using a single job as a template. It also provides reporting data for various areas (for example, all jobs performed within the last 24 hours or 7 days, all VMs engaged in these jobs, and so on). Using indexing data consolidated on one server, Veeam Backup Enterprise Manager provides advanced capabilities to search for VM guest OS files in VM backups created on all backup server (even if they are stored in repositories on different sites), and recover them in a single click. Search for VM guest OS files is enabled through Veeam Backup Enterprise Manager itself; to streamline the search process, you can optionally deploy a Veeam Backup Search server in your backup infrastructure.
- Veeam WAN Accelerator is a dedicated component of the backup infrastructure responsible for global data caching and data deduplication. On each WAN accelerator, Veeam Backup & Replication installs the Veeam WAN Accelerator Service responsible for WAN acceleration tasks.

The above components are integrated using component and design best practices to deliver an integrated infrastructure for Enterprise and cloud data centers.

The next section provides a technical overview of the hardware and software components of the present solution design.

Cisco Unified Computing System

Cisco brings 30 years of breadth, leadership, and vision to guide businesses through networking and infrastructure challenges. Cisco Unified Computing System (UCS) continues Cisco's long history of innovation in delivering integrated systems that deliver business results. Cisco UCS integrated infrastructure solutions speed up IT operations today and create the modern technology foundation needed for the critical business initiatives of tomorrow.

The Cisco Unified Computing System is a next-generation data center platform that unites compute, network, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. All resources participate in a unified management domain in an integrated, scalable, multi chassis platform.

The main components of Cisco Unified Computing System are:

- **Computing**—The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers based on Intel Xeon Processors.
- **Network**—The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **Virtualization**—The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access**—The system provides consolidated access to both (Storage Area Network) SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access the Cisco Unified Computing System can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This provides customers with choice for storage access and investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.
- **Management**—The system uniquely integrates all system components that enable the entire solution to be managed as a single entity by the Cisco UCS Manager (UCSM). The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations.

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system that unifies the technology in the data center. The system is managed, serviced and tested as a whole.
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand.
- Industry standards supported by a partner ecosystem of industry leaders

Cisco Unified Computing System Components

Cisco Fabric Interconnects

The Cisco UCS 6200 Series Fabric Interconnect is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel functions.

The Cisco UCS 6200 Series provides the management and communication backbone for the Cisco UCS C-Series and HX-Series rack-mount servers, Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All servers and chassis, and therefore all blades, attached to the Cisco UCS 6200 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6200 Series provides both the LAN and SAN connectivity for all blades within its domain.

From a networking perspective, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, 1Tb switching capacity, 160 Gbps bandwidth per chassis, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over a lossless Ethernet fabric from a server through an interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6248UP Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a one rack unit (1 RU) 10 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 960-Gbps throughput and up to 48 ports. The switch has 32 1/10-Gbps fixed Ethernet, FCoE and FC ports and 1 expansion slot.

Figure 2 Cisco UCS 6248UP Fabric Interconnect



Cisco UCS S3260 Storage Server

The Cisco UCS® S3260 Storage Server is a modular, high-density, high-availability dual-node rack server well suited for service providers, enterprises, and industry-specific environments. It provides dense, cost-effective storage to address your ever-growing data needs. Designed for a new class of data-intensive workloads, it is simple to deploy and excellent for applications for big data, data protection, software-defined storage environments, scale-out unstructured data repositories, media streaming, and content distribution.

Some of the key features of Cisco UCS S3260 Storage Server are

- Dual 2-socket server nodes based on Intel Xeon processor E5-2600 v2 or v4 CPUs with up to 36 cores per server node
- Up to 512 GB of DDR3 or DDR4 memory per server node (1 TB total)
- Support for high-performance Non-Volatile Memory Express (NVMe) and flash memory
- Massive 600-TB data storage capacity that easily scales to petabytes with Cisco UCS Manager
- Policy-based storage management framework for zero-touch capacity on demand
- Dual-port 40-Gbps system I/O controllers with Cisco UCS Virtual Interface Card (VIC) 1300 platform embedded chip
- Unified I/O for Ethernet or Fibre Channel to existing NAS or SAN storage environments
- Support for Cisco bidirectional (BiDi) transceivers, with 40-Gbps connectivity over existing 10-Gbps cabling infrastructure

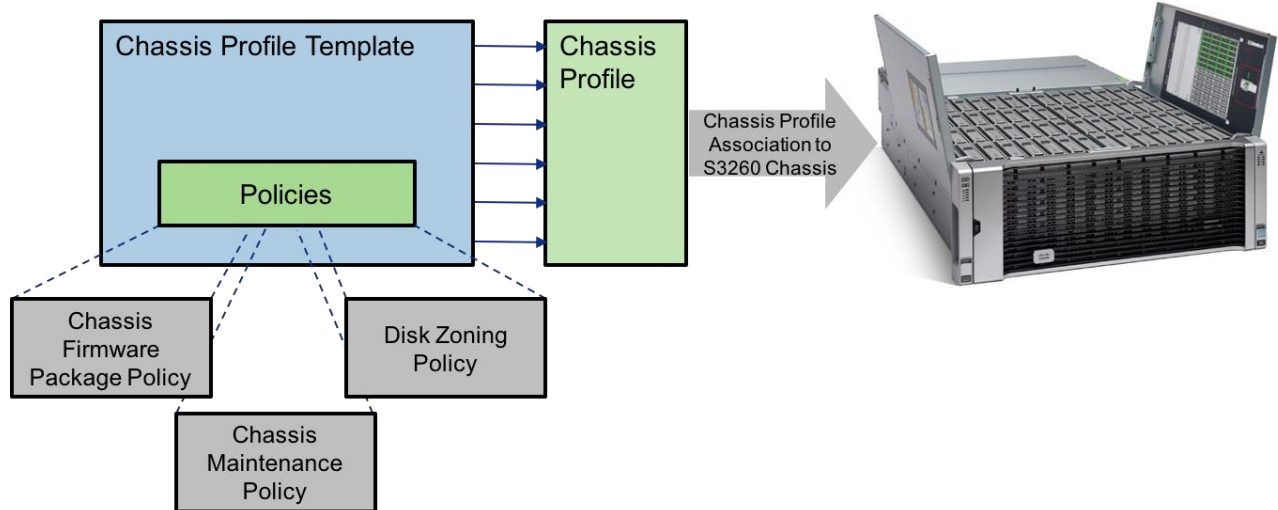
For more information, see: <http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-735611.html>

Figure 3 Cisco UCS C3260 M4 Rack Server



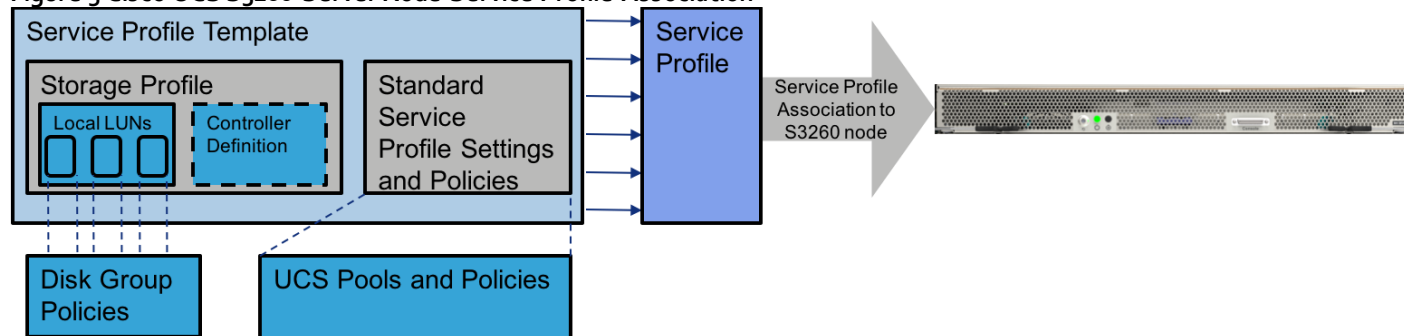
The S3260 can be CIMC managed, or Cisco UCS Manager managed as a registered Chassis with the Cisco UCS Fabric Interconnects. When the S3260 is Cisco UCS Manager managed, the Chassis will use a Chassis Profile that can be generated from template, and will contain specifications for Firmware and Maintenance policies as well as the Disk Zoning Policy. The Disk Zoning Policy will be used to set how disk slot allocation occurs between server nodes.

Figure 4 Cisco UCS S3260 Chassis Profile Association



Server Nodes in a Cisco UCS Manager managed S3260 are configured in nearly the same manner as standard Cisco UCS B-Series and Cisco UCS Manager managed Cisco UCS C-Series servers. The Server Nodes use Service Profiles that can be provisioned from a template. These nodes need to have a Storage Profile set within the Service Profile to be able to access the disk slots made available to it by the Disk Zoning Policy set within the Chassis Profile of the chassis the node is hosted within.

Figure 5 Cisco UCS S3260 Server Node Service Profile Association



Within the Storage Profile there are two main functions; Local LUN creation that is specified by Disk Group Policies, that are set within the Storage Profile. The LUNs created from the Disk Group Policies will have options of RAID 0, 1, 5, 6, 10, 50, or 60 and will allow the selection of type, quantity, or manual specification of slot the disk should be used from, as well as drive configuration policies of the LUN. S3260 M3 server nodes will use a Controller Definition, which will set how the PCH Controller should handle the rear facing SSDs of the S3260 Chassis. The Controller Definition is specific to the SSD drives allocated to the node by the PCH Controller and will have valid settings of RAID 0, 1, or no RAID.

Cisco HyperFlex HX-Series Nodes

A HyperFlex cluster requires a minimum of three HX-Series nodes. Data is replicated across at least two of these nodes, and a third node is required for continuous operation in the event of a single-node failure. The HX-Series nodes combine the CPU and RAM resources for hosting guest virtual machines, with the physical storage resources used by the HyperFlex software. Each HX-Series node is equipped with one high-performance SSD drive for data caching and rapid acknowledgment of write requests and is equipped with up to the platform's physical capacity of spinning disks for maximum data capacity.

Cisco HyperFlex HX220c-M4S Node

The Cisco HyperFlex HX220c-M4S rackmount server is one rack unit (1 RU) high and can mount in an industry-standard 19-inch rack. This small footprint configuration contains a minimum of three nodes with six 1.2 terabyte (TB) SAS drives that contribute to cluster storage capacity, a 120 GB SSD housekeeping drive, a 480 GB SSD caching drive, and two Cisco Flexible Flash (FlexFlash) Secure Digital (SD) cards that act as mirrored boot drives.

Figure 6 HX220c-M4S Node



Cisco HyperFlex HX240c-M4SX Node

The Cisco HyperFlex HX240c-M4S rackmount server is two-rack unit (2 RU) high and can mount in an industry-standard 19-inch rack. This capacity optimized configuration contains a minimum of three nodes, a minimum of fifteen and up to twenty-three 1.2 TB SAS drives that contribute to cluster storage, a single 120 GB SSD housekeeping drive, a single 1.6 TB SSD caching drive, and two FlexFlash SD cards that act as mirrored boot drives.

Figure 7 HX240c-M4SX Node



Cisco VIC 1227 MLOM Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1227 is a dual-port Enhanced Small Form-Factor Pluggable (SFP+) 10-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS HX-Series Rack Servers (**Error! Reference source not found.**). The mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, enabling a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile.

Figure 8 Cisco VIC 1227 mLOM Card



Cisco HyperFlex Converged Data Platform Software


The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology, exceeding the boundaries of first-generation hyperconverged infrastructures. The data platform has all the features that you would expect of an enterprise shared storage system, eliminating the need to configure and maintain complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:

- **Replication** of all written data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- **Deduplication** is always on, helping reduce storage requirements in which multiple operating system instances in client virtual machines result in large amounts of duplicate data.
- **Compression** further reduces storage requirements, reducing costs, and the log-structured file system is designed to store variable-sized blocks, reducing internal fragmentation.
- **Thin provisioning** allows large volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a “pay as you grow” proposition.
- **Fast, space-efficient clones** rapidly replicate virtual machines simply through metadata operations.
- **Snapshots** help facilitate backup and remote-replication operations: needed in enterprises that require always-on data availability.

Cisco HyperFlex HX Data Platform Administration Plug-in

The Cisco HyperFlex HX Data Platform is administered through a VMware vSphere web client plug-in. Through this centralized point of control for the cluster, administrators can create datastores, monitor the data platform health, and manage resource use. Administrators can also use this data to predict when the cluster will need to be scaled.

Figure 9 vCenter HyperFlex Web Client Plugin



Cisco HyperFlex HX Data Platform Controller

A Cisco HyperFlex HX Data Platform controller resides on each node and implements the Cisco HyperFlex HX Distributed Filesystem. The storage controller runs in user space within a virtual machine intercepting and handling all I/O from guest virtual machines. The storage controller VM uses the VMDirectPath I/O feature to provide PCI pass-through control of the physical server's SAS disk controller. This method gives the controller VM full control of the physical disk resources, utilizing the SSD drives as a read/write caching layer and the HDDs as a capacity layer for distributed storage. The controller integrates the data platform into VMware software through two preinstalled VMware ESXi vSphere Installation Bundles (VIBs):

- IOvisor:** The IOvisor is deployed on each node of the cluster and acts as a stateless NFS proxy that looks at each IO request and determines which cache vNode it belongs to and routes the IO to the physical node that owns that cache vNode. In the event of a failure, the IOvisor transparently handles it and will retry the same request to another copy of the data based on new information it receives. Decoupling the IOvisor from the controller VM enables access to the distributed filesystem and prevents hotspots. Compute-only nodes and VMs continue to perform storage IO in the event of a disk, SSD, or even a storage controller failure.
- VMware API for Array Integration (VAAI):** This storage offload API allows vSphere to request advanced file system operations such as snapshots and cloning. The controller implements these operations through the manipulation of metadata rather than actual data copying, providing rapid response, and thus rapid deployment of new environments.

Data Operations and Distribution

The Cisco HyperFlex HX Data Platform controllers handle all read and write operation requests from the guest VMs to their virtual disks (VMDK) stored in the distributed datastores in the cluster. The data platform distributes the data across multiple nodes of the cluster, and across multiple capacity disks of each node, per the replication level policy selected during the cluster setup. This method avoids storage hotspots on specific nodes, and on specific disks of the nodes, and thereby avoids networking hotspots or congestion from accessing more data on some nodes versus others.

Replication Factor

Enterprise class hyperconverged solutions should have three copies of data blocks across any three data nodes. This helps to ensure high availability during rare failure events such as single node failure and disk failure or during software and firmware upgrades, performed on a HX System. Thus three copies, or a replication factor of 3 (RF=3), is a default setting and a recommended best practice for HyperFlex systems.

- **Replication Factor 3:** For every I/O write committed to the storage layer, 2 additional copies of the blocks written will be created and stored in separate locations, for a total of 3 copies of the blocks. Blocks are distributed in such a way as to ensure multiple copies of the blocks are not stored on the same disks, nor on the same nodes of the cluster. This setting can tolerate simultaneous failures of 2 disks, or 2 entire nodes without losing data and resorting to restore from backup or other recovery processes.
- **Replication Factor 2:** For every I/O write committed to the storage layer, 1 additional copy of the blocks written will be created and stored in separate locations, for a total of 2 copies of the blocks. Blocks are distributed in such a way as to ensure multiple copies of the blocks are not stored on the same disks, nor on the same nodes of the cluster. This setting can tolerate a failure of 1 disk, or 1 entire node without losing data and resorting to restore from backup or other recovery processes.

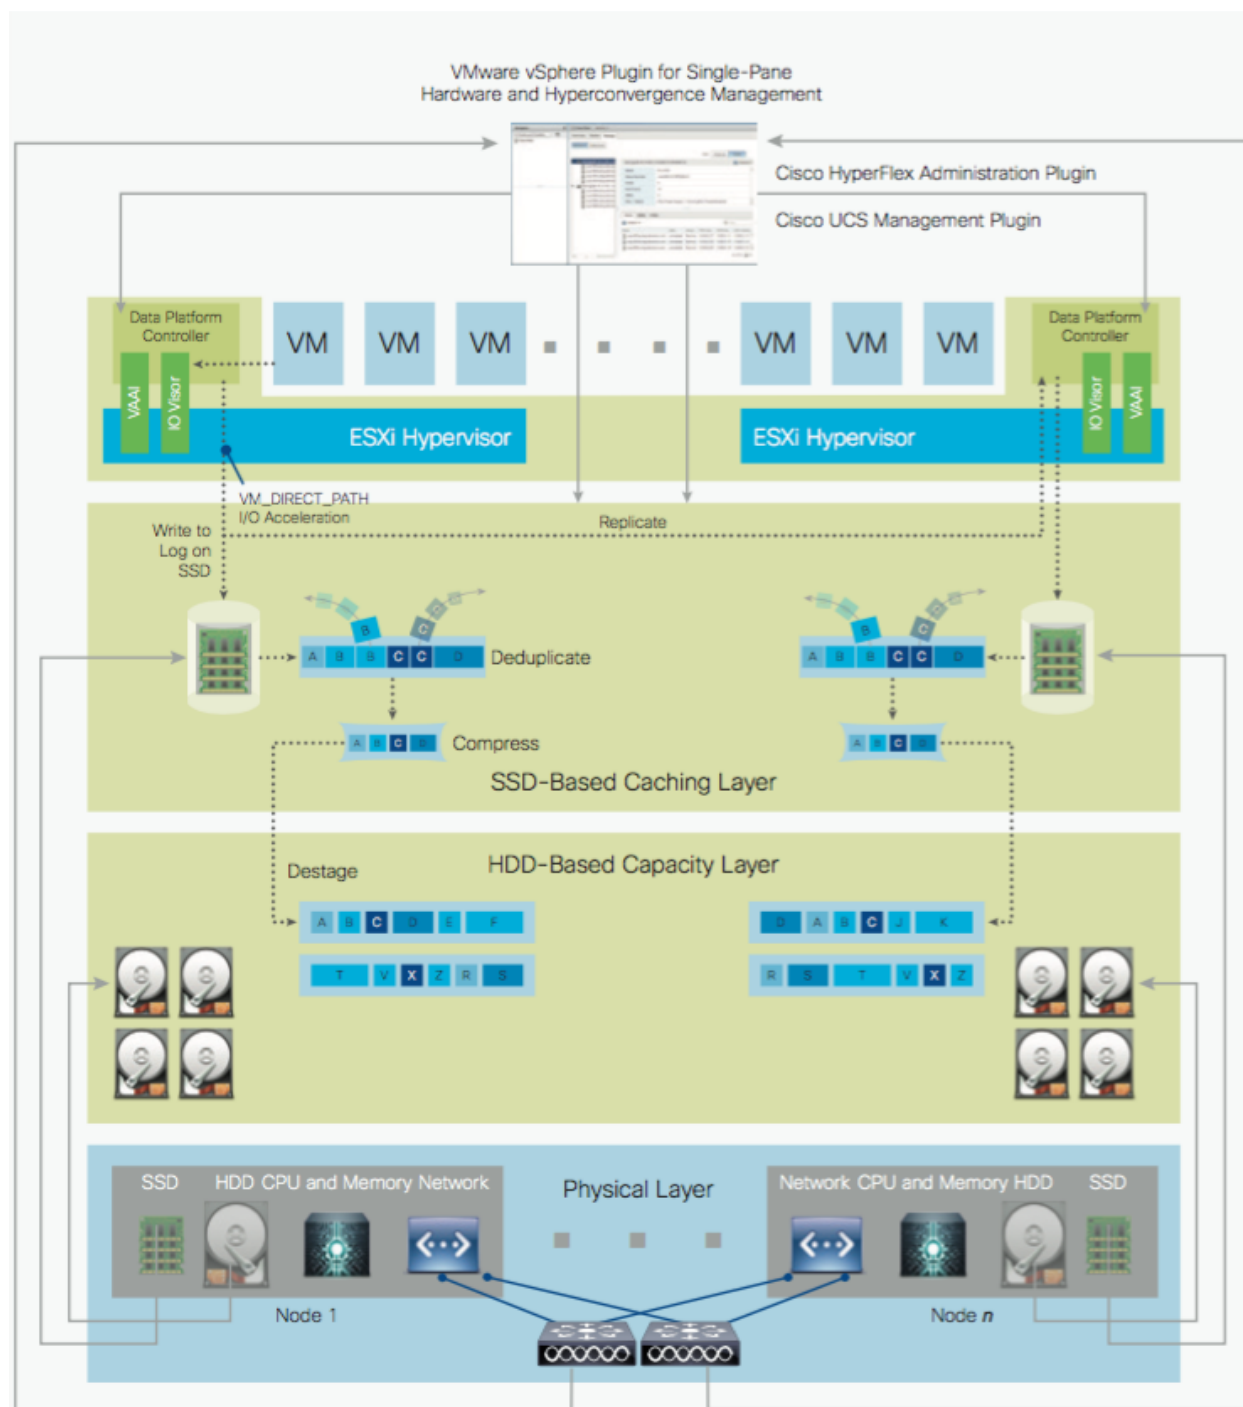
Data Write Operations

For each write operation, data is written to the local caching SSD on the node where the write originated, and replica copies of that write are written to the caching SSD of the remote nodes in the cluster, per the replication factor setting. For example, at RF=3 a write will be written locally where the VM originated the write, and two additional writes will be committed in parallel on two other nodes. The write operation will not be acknowledged until all three copies are written to the caching layer SSDs. Written data is also cached in a write log area resident in memory in the controller VM, along with the write log on the caching SSDs (Figure 10). This process speeds up read requests when reads are requested of data that has recently been written.

Data Destaging, Deduplication and Compression

The Cisco HyperFlex HX Data Platform constructs multiple write caching segments on the caching SSDs of each node in the distributed cluster. As write cache segments become full, and based on policies accounting for I/O load and access patterns, those write cache segments are locked and new writes roll over to a new write cache segment. The data in the now locked cache segment is destaged to the HDD capacity layer of the node. During the destaging process, data is deduplicated and compressed before being written to the HDD capacity layer. The resulting data after deduplication and compression can now be written in a single sequential operation to the HDDs of the server, avoiding disk head seek thrashing and accomplishing the task in the minimal amount of time (Figure 10). Since the data is already deduplicated and compressed before being written, the platform avoids additional I/O overhead often seen on competing systems, which must later do a read/dedupe/compress/write cycle.

Figure 10 HyperFlex HX Data Platform Data Movement



Data Read Operations

For data read operations, data may be read from multiple locations. For data that was very recently written, the data is likely to still exist in the write log of the local platform controller memory, or the write log of the local caching SSD. If local write logs do not contain the data, the distributed filesystem metadata will be queried to see if the data is cached elsewhere, either in write logs of remote nodes, or in the dedicated read cache area of the local and remote SSDs. Finally, if the data has not been accessed in a significant amount of time, the filesystem will retrieve the data requested from the HDD capacity layer. As requests for reads are made to the distributed filesystem and the data is retrieved from the HDD capacity layer, the caching SSDs populate their dedicated read cache area to speed up subsequent requests for the same data. This multi-tiered distributed system with several layers of

caching techniques, ensures that data is served at the highest possible speed, leveraging the caching SSDs of the nodes fully and equally.

Veeam Availability Suite

Veeam is an industry leader within the data protection market. In the era of Digital Transformation, Veeam recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise by helping organizations meet today's service-level objectives, enabling recovery of any IT service and related applications and data within seconds and minutes. Veeam consistently pushes the envelope in bringing sophisticated backup and disaster recovery functionality to enterprises and cloud providers.

Backup

Veeam Backup & Replication operates at the virtualization layer and uses an image-based approach for VM backup. To retrieve VM data, no agent software needs to be installed inside the guest OS. Instead, Veeam Backup & Replication leverages vSphere snapshot capabilities and Application Aware Processing. When a new backup session starts, a snapshot is taken to create a cohesive point-in-time copy of a VM including its configuration, OS, applications, associated data, system state and so on. Veeam Backup & Replication uses this point-in-time copy to retrieve VM data. Image-based backups can be used for different types of recovery, including full VM recovery, VM file recovery, Instant VM Recovery, file-level recovery and application item recovery.


Use of the image-based approach allows Veeam Backup & Replication to overcome shortfalls and limitations of traditional backup. It also helps streamline recovery verification and the restore process — to recover a single VM, there is no need to perform multiple restore operations. Veeam Backup & Replication uses a cohesive VM image from the backup to restore a VM to the required state without the necessity for manual reconfiguration and adjustment. In Veeam Backup & Replication, backup is a job-driven process where one backup job can be used to process one or more VMs. A job is a configuration unit of the backup activity. Essentially, the job defines when, what, how and where to back up. It indicates what VMs should be processed, what components should be used for retrieving and processing VM data, what backup options should be enabled and where to save the resulting backup file. Jobs can be started manually by the user or scheduled to run automatically. The resulting backup file stores compressed and deduplicated VM data. Compression and Deduplication is done by the Veeam Proxy server.

Regardless of the Backup method you use, the first run of a job creates a full backup of VM image. Subsequent job runs are incremental — Veeam Backup & Replication copies only those data blocks that have changed since the last backup job run. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches, including VMware's Changed Block Tracking (CBT) technology.

Changed Block Tracking

To perform incremental backup, Veeam Backup & Replication needs to know which data blocks have changed since the previous job run.

Figure 11 Change Block Tracking



For VMware VMs with hardware version 7 or later, Veeam Backup & Replication employs VMware vSphere Changed Block Tracking (CBT) — a native VMware feature. Instead of scanning VMFS, Veeam Backup &

Replication queries CBT on vSphere through VADP and gets the list of blocks that have changed since the last run of this particular job. Use of CBT increases the speed and efficiency of block-level incremental backups. CBT is enabled by default; if necessary, you can disable it in the settings of a specific backup job.

Restore

Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios:

- **Veeam Explorer** enables you to restore Single Application specific items
- **Instant VM Recovery** enables you to instantly start a VM directly from a backup file
- **Full VM recovery** enables you to recover a VM from a backup file to its original or another location
- **VM file recovery** enables you to recover separate VM files (virtual disks, configuration files and so on)
- **Virtual drive restore** enables you to recover a specific hard drive of a VM from the backup file, and attach it to the original VM or to a new VM
- **Windows file-level recovery** enables you to recover individual Windows guest OS files (from FAT, NTFS and ReFS file systems)
- **Multi-OS file-level recovery** enables you to recover files from 15 different guest OS file systems

Veeam Backup & Replication uses the same image-level backup for all data recovery operations. You can restore VMs, VM files and drives, application objects and individual guest OS files to the most recent state or to any available restore point.

Veeam Explorer

Veeam Explorers are tools included in all editions of Veeam Backup & Replication. As of v9 and restore application items directly from VM backups and replicas. It provides fast and effortless Active Directory, Exchange, SharePoint, SQL Server and Oracle recovery without needing to provision extra storage, deploy agents, restore an entire virtual machine (VM) for granular recovery or spin anything up in an isolated network. This includes powerful, easy-to-use and affordable eDiscovery and granular recovery for the following:

- **Microsoft Active Directory:** Search and restore all Active Directory object types (e.g., users, groups, computer accounts, contacts, expiring links), Group Policy Objects (GPOs), Active Directory-integrated Microsoft DNS records and Configuration Partition objects.
- **Microsoft Exchange:** Get instant visibility into Exchange 2010, 2013 and 2016 backups, advanced search capabilities and quick recovery of individual Exchange items (e.g., emails, contacts, notes, etc.), Online Archive mailboxes, Purges folder support and hard-deleted (i.e., permanently deleted) items; eDiscovery features include detailed export reports and export size estimation based on query search criteria.
- **Microsoft SharePoint:** Get instant visibility into SharePoint 2010, 2013 and 2016 backups, search for and quickly restore full SharePoint sites, item permissions and specific files. Export recovered items directly to their original SharePoint server or send them as an email attachment.
- **Microsoft SQL Server:** Get fast transaction- and table-level recovery of SQL databases, including agentless transaction log backup and replay, so you can restore your SQL databases to a precise point in time and achieve low RTPO.
- **Oracle:** Get transaction-level recovery of Oracle databases including agentless transaction log backup, so you can restore your Oracle databases to a precise point in time, self-service restore and restore via PowerShell.



Each Explorer has a corresponding User guide.

Instant VM Recovery

With instant VM recovery, you can immediately restore a VM into your production environment by running it directly from the backup file. Instant VM recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production VMs. It is like having a "temporary spare" for a VM; users remain productive while you can troubleshoot an issue with the failed VM.

When instant VM recovery is performed, Veeam Backup & Replication uses the Veeam vPower technology to mount a VM image to an ESX(i) host directly from a compressed and deduplicated backup file. Since there is no need to extract the VM from the backup file and copy it to production storage, you can restart a VM from any restore point (incremental or full) in a matter of minutes.

After the VM is back online you can use VMware storage vMotion to migrate the VM back to production storage.

VM Object Recovery

Veeam Backup & Replication can help you to restore specific VM files (.vmdk, .vmx and others) if any of these files are deleted or the datastore is corrupted. This option provides a great alternative to full VM restore, for example, when your VM configuration file is missing and you need to restore it. Instead of restoring the whole VM image to the production storage, you can restore the specific VM file only. Another data recovery option provided by Veeam Backup & Replication is restore of a specific hard drive of a VM. If a VM hard drive becomes corrupted for some reason (for example, with a virus), you can restore it from the image-based backup to any good-to-know point in time.

Replication

To ensure efficient and reliable data protection in your virtual environment, Veeam Backup & Replication complements image-based backup with image-based replication. Replication is the process of copying a VM from its primary location (source host) to a destination location (redundant target host). Veeam Backup & Replication creates an exact copy of the VM (replica), registers it on the target host and maintains it in sync with the original VM.

Replication provides the best recovery time objective (RTO) and recovery point objective (RPO) values, as you actually have a copy of your VM in a ready-to-start state. That is why replication is commonly recommended for the most critical VMs that need minimum RTOs. Veeam Backup & Replication provides means to perform both onsite replication for high availability (HA) scenarios and remote (offsite) replication for disaster recovery (DR) scenarios. To facilitate replication over WAN or slow connections, Veeam Backup & Replication optimizes traffic transmission — it filters out unnecessary data blocks (such as, duplicate data blocks, zero data blocks or blocks of swap files) and compresses replica traffic. Veeam Backup & Replication also allows you to apply network-throttling rules to prevent replication jobs from consuming the entire bandwidth available in your environment.

Replication is a job-driven process with one replication job used to process one or more VMs. You can start the job manually every time you need to copy VM data or, if you want to run replication unattended, create a schedule to start the job automatically. Scheduling options for replication jobs are similar to those for backup jobs.

WAN Acceleration

WAN accelerators are optional components in the replication infrastructure. You can use WAN accelerators if you replicate VMs over a slow connection or over the WAN.

In the replication process, WAN accelerators are responsible for global data caching and deduplication. To use WAN acceleration, you must deploy two WAN accelerators in the following manner:

- The **source WAN accelerator** must be deployed in the source side, close to the backup proxy running the source-side Data Mover Service.
- The **target WAN accelerator** must be deployed in the target side, close to the backup proxy running the target-side Data Mover Service.

Deployment Types

Veeam Backup & Replication supports a number of replication scenarios that depend on the location of the target host and the data transport path.

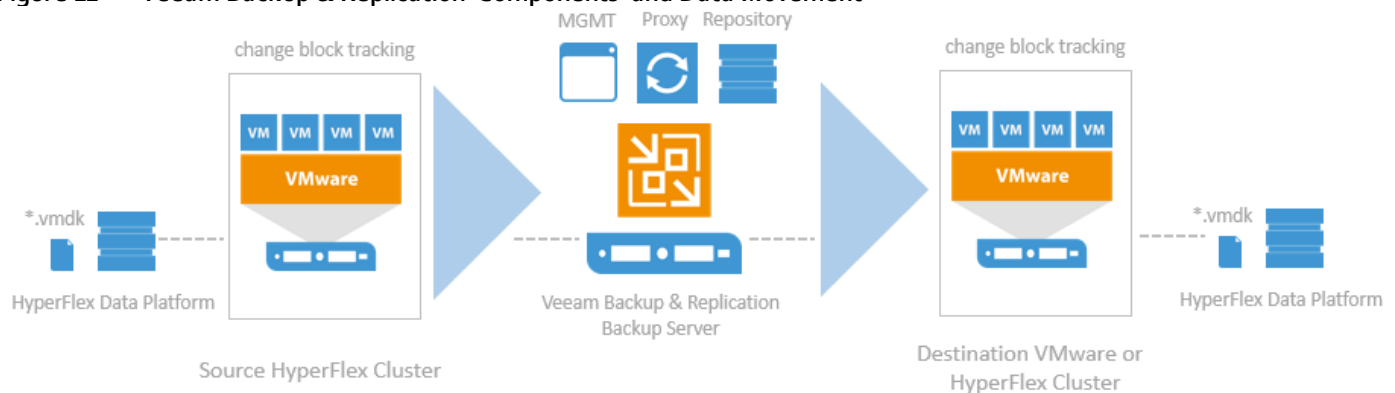
Onsite Replication

If the source host and the target host are located in the same site, you can perform onsite replication.

Onsite replication requires the following replication infrastructure components:

- **Backup proxy.** In the onsite replication scenario, the source-side Data Mover Service and the target-side Data Mover Service are started on the same backup proxy. The backup proxy must have access to the backup server, source host, target host and backup repository holding replica metadata.
- **Backup repository** for storing replica metadata

Figure 12 Veeam Backup & Replication Components and Data Movement



In the onsite replication scenario, Veeam Backup & Replication does not perform data compression. Replication traffic is transferred uncompressed between the two Data Mover Services started on the same backup proxy.

Offsite Replication

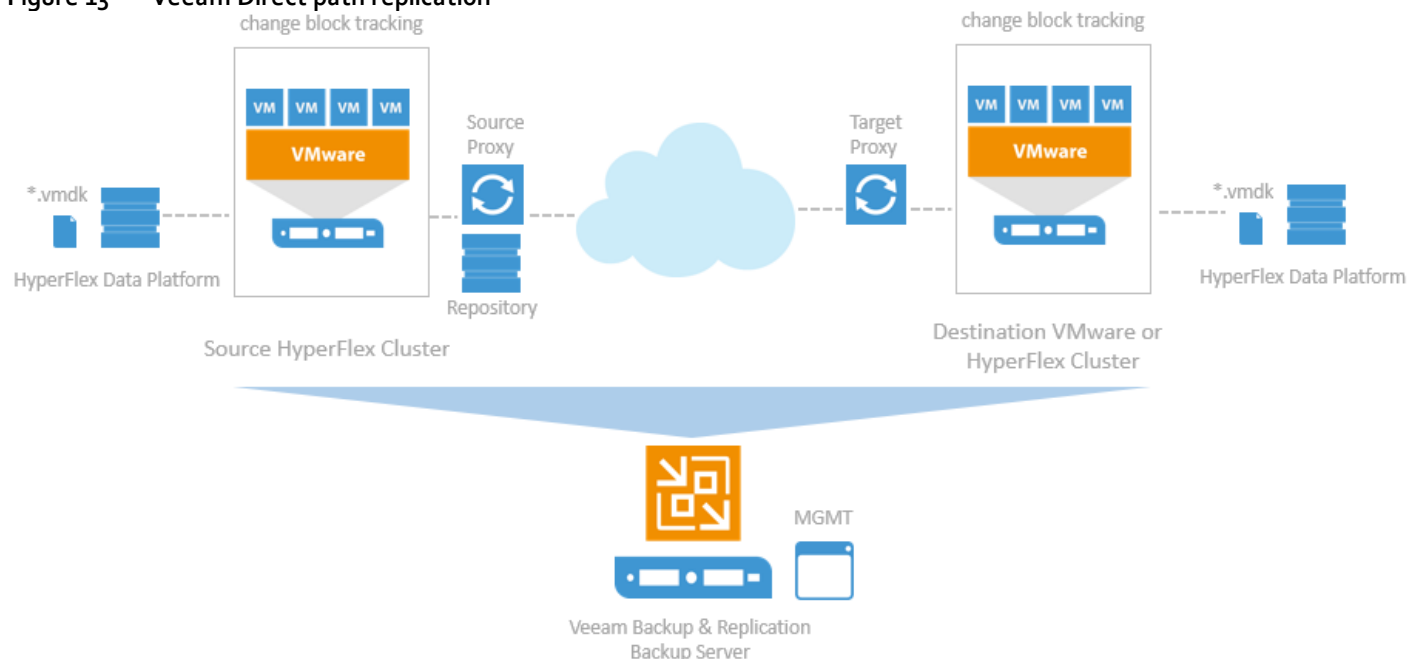
If the source host is located in the primary site and the target host is located in the DR site, you can perform offsite replication.

Offsite replication can run over two data paths:

- Direct data path
- Via a pair of WAN accelerators

Direct Data Path

The common requirement for offsite replication is that one Data Mover Service runs in the production site, closer to the source host, and another Data Mover Service runs in the remote DR site, closer to the target host. During backup, the Data Mover Services maintain a stable connection, which allows for uninterrupted operation over the WAN or slow links. For more information, see [Resume on WAN Disconnect](#).

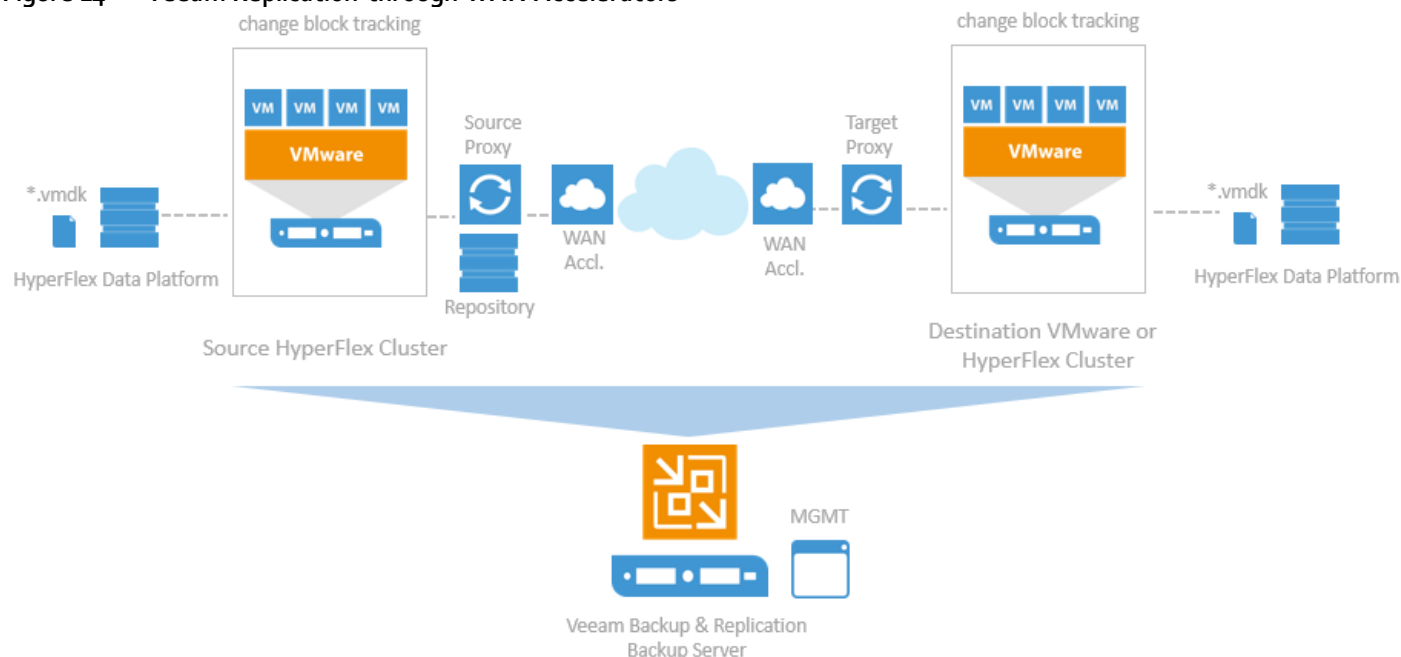
Figure 13 Veeam Direct path replication

Via WAN Accelerators

If you have a high latency WAN link, you can replicate VM data via a pair of WAN accelerators. WAN accelerators provide advanced technologies to optimize VM data transfer:

- Global data caching and deduplication
- Resume on disconnect for uninterrupted data transfer

WAN accelerators add a new layer in the backup infrastructure — a layer between the source-side Data Mover Service and the target-side Data Mover Service. The data flow goes from the source backup proxy via a pair of WAN accelerators to the target backup proxy that, finally, transports VM data to the target host.

Figure 14 Veeam Replication through WAN Accelerators

Failover and Failback

In case of software or hardware malfunction, you can quickly recover a corrupted VM by failing over to its replica. When you perform failover, a replicated VM takes over the role of the original VM. You can fail over to the latest state of a replica or to any of its good known restore points.

In Veeam Backup & Replication, failover is a temporary intermediate step that should be further finalized. Veeam Backup & Replication offers the following options for different disaster recovery scenarios:

- You can perform **permanent failover** to leave the workload on the target host and let the replica VM act as the original VM. Permanent failover is suitable if the source and target hosts are nearly equal in terms of resources and are located on the same HA site.
- You can perform **failback** to recover the original VM on the source host or in a new location. Failback is used in case you failed over to a DR site that is not intended for continuous operations and would like to move the operations back to the production site when the consequences of a disaster are eliminated.

Veeam Backup & Replication supports failover and failback operations for one VM and for several VMs. In case one or several hosts fail, you can use failover plans to restore operations with minimum downtime.

Failover-Plans

If you have a number of VMs running interdependent applications, you need to failover them one by one, as a group. To do this automatically, you can prepare a failover plan.

In a failover plan, you set the order in which VMs must be processed and time delays for VMs. The time delay is an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. It helps to ensure that some VMs, such as a DNS server, are already running at the time the dependent VMs start. The failover plan must be created in advance. In case the primary VM group goes offline, you can start the corresponding failover plan manually. When you start the procedure, you can choose to fail over to the latest state of a VM replica or to any of its good known restore points.

Planned Failover

If you know that your primary VMs are about to go offline, you can proactively switch the workload to their replicas. A planned failover is smooth manual switching from a primary VM to its replica with minimum interrupting in operation. You can use the planned failover, for example, if you plan to perform datacenter migration, maintenance or software upgrade of the primary VMs. You can also perform planned failover if you have an advance notice of a disaster approaching that will require taking the primary servers offline.

Failback

If you want to resume operation of a production VM, you can fail back to it from a VM replica. When you perform failback, you get back from the VM replica to the original VM, shift your I/O and processes from the target host to the production host and return to the normal operation mode.

If you managed to restore operation of the source host, you can switch from the VM replica to the original VM on the source host. If the source host is not available, you can restore the original VM to a new location and switch back to it.


Backup Server

Components

Veeam Availability Suite combines the backup, restore and replication capabilities of Veeam Backup & Replication™ with the advanced monitoring, reporting and capacity planning functionality of Veeam ONE™. Veeam Availability Suite delivers everything you need to reliably protect and manage your Cisco HyperFlex VMware environment. Veeam Backup & Replication is a modular solution that lets you build a scalable backup infrastructure for environments of different sizes and configuration. The installation package of Veeam Backup & Replication includes a set of components that you can use to configure the backup infrastructure. Some components are mandatory and provide core functionality; some components are optional and can be installed to provide additional functionality for your business and deployment needs. You can co-install all Veeam Backup & Replication components on the same machine, physical or virtual, or you can set them up separately for a more scalable approach.

Figure 15 shows an overview on the main Veeam components.

Figure 15 Veeam Backup & Replication Components




Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure that fills the role of the “configuration and control center”. The backup server performs all types of administrative activities:

- **Coordinates** backup, replication, recovery verification and restore tasks
- **Controls** job scheduling and resource allocation
- **Manages** all Proxy and Repository servers and other components of the backup infrastructure

The backup server is used to set up and manage backup infrastructure components as well as specify global settings for the backup infrastructure.

Figure 16 Veeam Backup Server Management



In addition to its primary functions, a newly deployed backup server also performs the roles of the default backup proxy and the backup repository.

The backup server uses the following services and components:

- **Veeam Backup Service** is a Windows service that coordinates all operations performed by Veeam Backup & Replication such as backup, replication, recovery verification and restore tasks. The Veeam

Backup Service runs under the Local System account or account that has the Local Administrator permissions on the backup server.


- **Veeam Backup Shell** provides the application user interface and allows user access to the application's functionality.
- **Veeam Guest Catalog Service** is a Windows service that manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files. Index data is stored in the Veeam Backup Catalog — a folder on the backup server. The Veeam Guest Catalog Service running on the backup server works in conjunction with search components installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft Search Server.
- **Veeam Backup SQL Database** is used by Veeam Backup Service, Veeam Backup Shell and Veeam Guest Catalog Service to store data about the backup infrastructure, jobs, sessions and so on. The database instance can be located on a SQL Server installed either locally (on the same machine where the backup server is running) or remotely.
- **Veeam Backup PowerShell Snap-In** is an extension for Microsoft Windows PowerShell 2.0. Veeam Backup PowerShell adds a set of cmdlets to allow users to perform backup, replication and recovery tasks through the command-line interface of PowerShell or run custom scripts to fully automate operation of Veeam Backup & Replication.
- **Backup Proxy Services.** In addition to dedicated services, the backup server runs a set of data mover services.

Backup Proxy

The backup proxy is an architecture component that sits between data source and target and is used to process jobs and deliver backup traffic. In particular, the backup proxy tasks include retrieving VM data from the production storage, compressing, deduplication and sending it to the backup repository (for example, if you run a backup job) or another backup proxy (for example, if you run a replication job). As the data handling task is assigned to the backup proxy, the backup server becomes the “point of control” for dispatching jobs to proxy servers.

The role of a backup proxy can be assigned to a dedicated Windows server (physical or virtual) in your environment. You can deploy backup proxies both in the primary site and in remote sites. To optimize performance of several concurrent jobs, you can use a number of backup proxies. In this case, Veeam Backup & Replication will distribute the backup workload between available backup proxies.

Figure 17 Veeam distributed Proxy Server Deployment



Use of backup proxies lets you easily scale your backup infrastructure up and down based on your demands. Backup proxies run lightweight services that take a few seconds to deploy. The primary role of the backup proxy is to provide an optimal route for backup traffic and enable efficient data transfer.

The backup proxy uses the following services and components:

- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyses the system, installs and upgrades necessary components and services depending on the role selected for the server.
- **Veeam Data Mover Service** is responsible for deploying and coordinating executable modules that act as "data movers" and perform main job activities on behalf of Veeam Backup & Replication, such as communicating with VMware Tools, copying VM files, performing data deduplication and compression and so on.

Backup Repository

A backup repository is a location used by Veeam Backup & Replication jobs to store backup files, copies of VMs and metadata for replicated VMs. By assigning different repositories to jobs and limiting the number of parallel jobs for each one, you can balance the load across your backup infrastructure.

You can configure one of the following types of backup repositories:

- **Microsoft Windows server with local or directly attached storage.** The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), or iSCSI/FC SAN LUN in case the server


is connected into the SAN fabric.

- **Linux server with local, directly attached storage, or mounted NFS.** The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), NFS share, or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.
- **CIFS (SMB) share.** SMB share cannot host Veeam Data Mover Services. For this reason, data to the SMB share is written from the gateway server. By default, this role performed by a backup proxy that is used by the job for data transport.
- **Deduplicating storage appliance.** Veeam Backup & Replication supports different deduplicating storage appliances.

Backup & Replication Console

The Veeam Backup & Replication console is a separate client-side component that provides access to the backup server. The console is installed locally on the backup server by default. You can also use it in a standalone mode — install the console on a dedicated machine and access Veeam Backup & Replication remotely over the network. The console lets you log into Veeam Backup & Replication and perform all kinds of data protection and disaster recovery operations as if you are working on the backup server.

Figure 18 Veeam Backup & Replication Console



You can install as many remote consoles as you need so that multiple users can access Veeam Backup & Replication simultaneously. Veeam Backup & Replication prevents concurrent modifications on the backup server.

Backup Proxy

Transport Modes

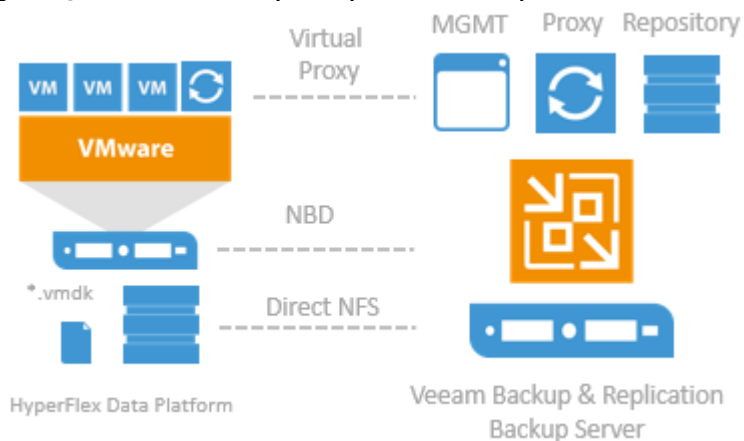
Job efficiency and time required for job completion greatly depends on the transport mode. The transport mode is a method that is used by the Veeam Data Mover Service to retrieve VM data from the source and write VM data to the target.

For data retrieval, Veeam Backup & Replication offers the following modes:

- Direct storage access

- Virtual appliance
- Network (NBD)

Figure 19 Veeam Backup & Replication Transport Modes



In the **Direct storage access** mode, Veeam Backup & Replication reads/writes data directly from/to the storage system where VM data or backups are located. With the Direct NFS access mode for Cisco HyperFlex, Veeam Backup & Replication bypasses the ESX(i) host and reads/writes data directly from/to NFS datastores. To do this, Veeam Backup & Replication deploys its native NFS client on the backup proxy and uses it for VM data transport. VM data still travels over the LAN but there is no load on the ESX(i) host.

The **Virtual appliance mode** is recommended if the role of a backup proxy is assigned to a VM. In the Virtual appliance mode, Veeam Backup & Replication uses the VMware SCSI HotAdd capability that allows attaching devices to a VM while the VM is running. During backup, replication or restore disks of the processed VM are attached to the backup proxy. VM data is retrieved or written directly from/to the datastore, instead of going through the network.

The **Network** mode can be used with any infrastructure configuration. In this mode, data is retrieved via the ESX(i) host over the LAN using the Network Block Device protocol (NBD). The Network mode is the recommended data transport mode to be used with Cisco HyperFlex in combination with Native HX Snapshots. To take the full advantage of the mode a 10Gbit/s Ethernet is mandatory.

Veeam Repository Sizing

When estimating the amount of required disk space, you should know the following:

- Number of backup VMs, total size of VMs and the data change rate
- Frequency of backups
- Retention period for backups

In addition, when testing is not possible beforehand, you should make assumptions on compression and deduplication ratios, change rates, and other factors. The following figures are typical for most deployments; however, it is important to understand the specific environment to find out possible exceptions:

- Data reduction thanks to Compression and Deduplication is usually 2:1 or more; it is common to see 3:1 or better, but you should always be conservative when estimating required space.
- Typical daily change rate is between 2 and 5% in a mid-size or enterprise environment; this can greatly vary among servers; some servers show much higher values. If possible, run monitoring tools like Veeam ONE to have a better understanding of the real change rate values.

- Include additional space for one-off full backups.
- Include additional space for backup chain transformation (forward forever incremental, reverse incremental) – at least the size of a full backup multiplied by 1.25x.
- Using the numbers above, you can estimate required disk space for any job. Besides, always leave plenty of extra headroom for future growth, additional full backups, moving VMs, restoring VMs from tape.



A repository sizing tool that can be used for estimation is available at <http://vee.am/rps>. Note that this tool is not officially supported by Veeam, and it should be used "as is", but it is nonetheless heavily used by Veeam Architects and regularly updated.

Solution Design

Data Protection for Cisco HyperFlex with Veeam Availability Suite is designed to deliver reliable backup and recovery solution with low recovery time objectives (RTOs) and recovery point objectives (RPOs) for all applications and data residing in virtual machines within the HyperFlex environment.

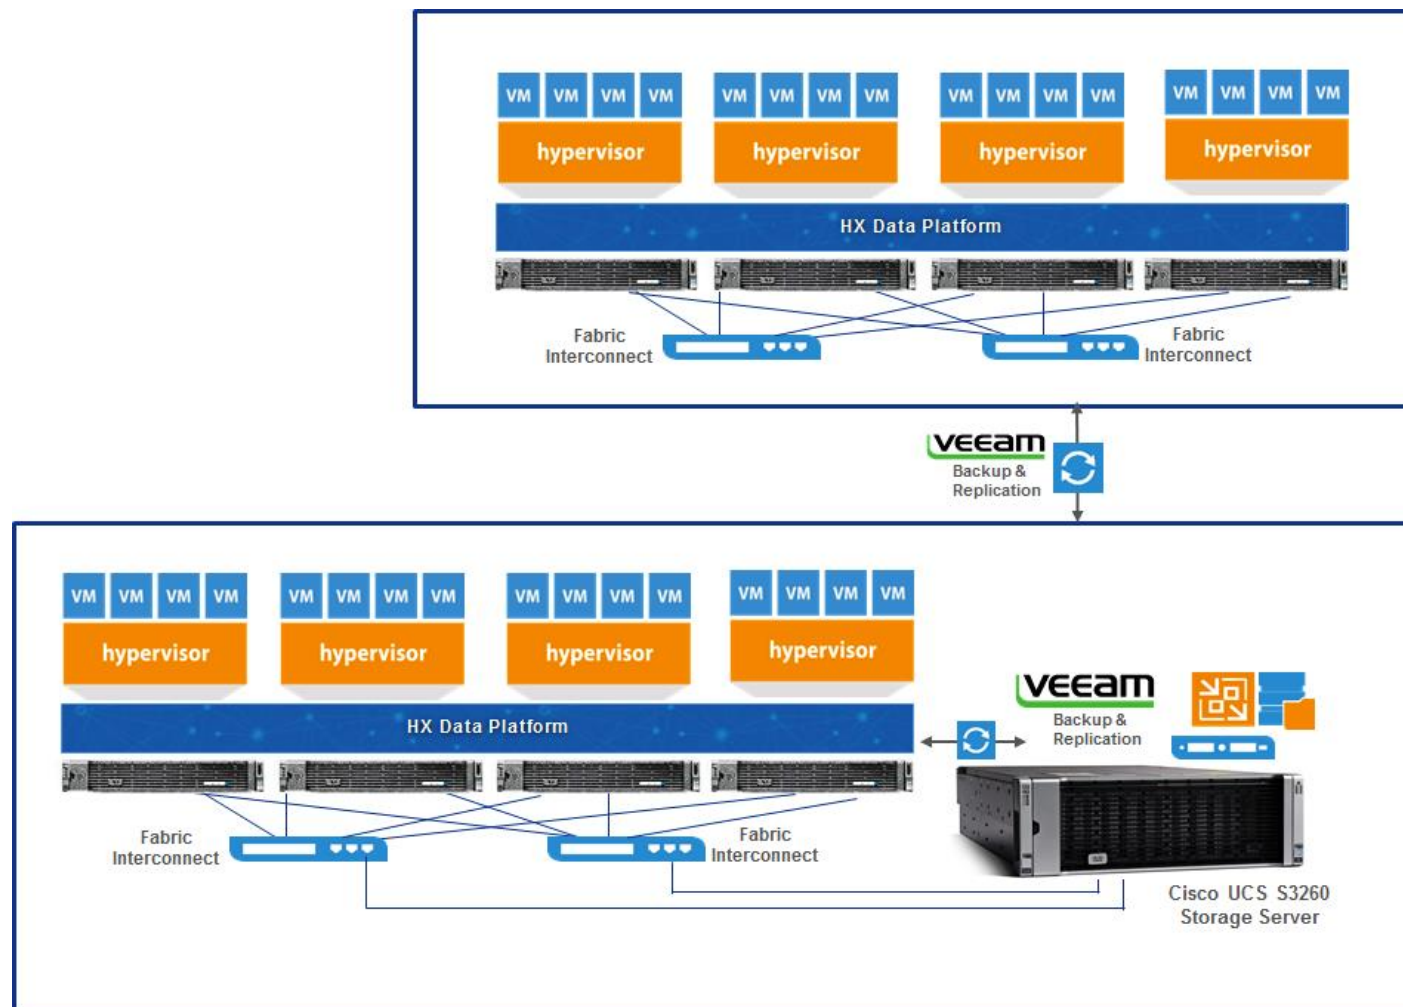
In addition to reliable backup and recovery of application data and VMs the solution provides

- Granular recovery of virtual machines and files
- Ability to automatically verify every backup, VM and replica
- Instant VM recovery of failed VM in less than two minutes
- Multiple backup end points such as tape drives, on cloud or on local repository
- SureBackup and SureReplica for Backup and Replication verification

This section elaborates on the deployment architecture and design considerations to protect application data through Veeam Availability Suite. The application VMs reside within multiple HyperFlex Clusters within same Data Center.

0 illustrates the end-to-end deployment scenarios for Cisco HyperFlex with Veeam Availability Suite.

Figure 20 Deployment Overview: Cisco HyperFlex with Veeam Availability Suite




Single Site Backup and Replication for HyperFlex provides protection for applications, data, and VMs within the same deployment site or data center. The key features of Single Site Backup and Replication are as follows:

- Veeam Availability Suite which include Veeam Backup Server, Veeam Proxy and Veeam Repository resides on Cisco UCS S3260 storage server.
- Backup of Primary HyperFlex Cluster VMs and application data to Veeam Repository
- Replication of HyperFlex Cluster VMs to either a standalone VMWare ESXi Cluster or to another HyperFlex cluster through common Veeam Proxy Server
- Backup of secondary HyperFlex Cluster VMs located in the same Data Center to common Veeam Repository and Proxy deployed on Cisco UCS S3260 Storage Server.

Figure 21 elaborates on the use case for Cisco Hyper Flex Single Site Backup and Replication

Figure 21 Cisco HyperFlex Single Site Backup and Replication




Deployment Architecture

Veeam Availability Suite (Veeam AS) provides, resilient protection of VMs and Application Data on Cisco HyperFlex Cluster, deployed either in same Data Center, or in Remote Office Branch Office (ROBO), or different HyperFlex Cluster deployed across Data Center in different geographical locations.

This section describes the reference architecture for Veeam AS with Cisco HyperFlex deployed in same campus or data center. Figure 22 details the physical topology of the present deployment.

Figure 22 Deployment Architecture: Cisco HyperFlex Single Site Backup and Replication



The solution detailed in the figure above includes Primary and secondary HyperFlex Clusters and provides Backup and Replication of VMs and application data through Veeam Availability Suite deployed on Cisco UCS S3260 Storage server. The details on the Primary and Secondary HX Clusters are as follows:

- Primary HyperFlex Cluster
 - Veeam Application Suite 9.5 deployed on Cisco UCS S3260 Storage Server. This includes Veeam Repository, Veeam Backup Server and the Veeam Proxy Server.
 - Cisco UCS S3260 Storage Server is directly attached to a pair of Cisco UCS 6200 Series Fabric Interconnects (FI)
 - Each of the 40Gbps ports on S3260 are connect to 4 X10Gbps ports on each of the Fabric Interconnect cluster through a QSFP to four SFP+ active optical breakout cable

- Cisco HX Cluster with Cisco HX Data Platform 1.8a and ESXi Hypervisor 6.0 update 2. This is the primary HX Cluster wherein all the application VMs reside.
 - Cisco HX Cluster and Cisco UCS S3260 Storage Server are connected to the same pair of Fabric Interconnects
 - The backup of the application VMs are created on S3260 Storage Server repository through Veeam AS
 - Application VMs are replicated either to standalone ESXi Cluster or the Secondary HX Cluster residing in the same Data Center or Campus
 - Cisco Nexus 9300 switches
- Secondary HyperFlex Cluster
 - Cisco HX Cluster with Cisco HX Data Platform 1.8a and ESXi Hypervisor 6.0 update 2. This is the secondary cluster wherein, either the primary VM replica reside or the actual application VMs are executed
 - Backup of the application VM on secondary HX Cluster are created on Cisco UCS S3260 Storage Server, residing in the primary HX Cluster domain
 - Cisco UCS 6200 Series Fabric Interconnects
 - Cisco Nexus 9300 switches

The primary and Secondary HX Cluster domains, reside in the same data center or campus and can be connected through either 1Gbe or 10 Gbe data links. In the present architecture, both the cluster domains are connected through a 1 Gbe data link.

Deployment Hardware and Software

Table 1 details the software revisions used throughout this document.

Software Versions

Table 1 Software Revisions


	Components	Software Version	Comments
Compute & Storage	Cisco UCS S3260 M4 Storage Server	3.1(2b)	Directly managed through Fabric Interconnect. Veeam AS is installed on the same. Provides Storage Veeam Repository
	Cisco HX220c M4		Hyper Converged node for HX Cluster
	Cisco HX240c M4		Hyper Converged Node for HX Cluster
Management	Cisco UCS Manager	3.1(2b)	UCS Management for all servers directly attached to Fabric Interconnects
Backup and Replication	Veeam Availability Suite	9.5	Pre-configured with Veeam Backup Server, Veeam Proxy , Veeam Repository
	Operating System	Windows 2016 DataCenter Edition	
Hyper Converged Software	Cisco HX Data Platform	HX Data Platform Release 1.8a	
Virtualization	VMWare VSphere	6.0 U2	
	VMWare vCenter	6.0 U2	
Network	Cisco Nexus 9372PX (N9k-9372PX)	6.1(2)I3(4b)	Cisco Platform Switch for ToR, MoR, EoR deployments; Provides connectivity to users and other networks and deployed in NX-OS Standalone mode
	Cisco UCS 6248UP FI	3.1(2b)	Fabric Interconnect with embedded UCS Manager

Configuration Guidelines

This CVD provides the details to configure a Cisco UCS S3260 Storage Server setup with Veeam 9.5, thus providing backup, restore and replication of a VM deployed in a Cisco HyperFlex infrastructure.

Figure 23 illustrates the high-level procedures and steps for this installation

Figure 23 Installation Workflow



Cisco Nexus A and Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning.

To indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the `network port vlan create` command:

Usage:

```

network port vlan create ?
  [-node] <nodename>           Node
  { [-vlan-name] {<netport>|<ifgrp>} VLAN Name
  | -port {<netport>|<ifgrp>} Associated Network Port
  [-vlan-id] <integer> }       Network Switch VLAN Identifier
  
```

Example:

```
network port vlan -node <node01> -vlan-name a0a-<vlan id>
```



This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses.

Table 2 describes the VLANs necessary for deployment as outlined in this guide.

Error! Reference source not found. lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	VLAN ID Used in Validating This Document
hx-inband-mgmt	ESXi host management interfaces HX Storage Controller VM management interfaces HX Storage Cluster roaming management interface Veeam Network	3175
hx-storage-data	ESXi host storage vmkernel interfaces HX Storage Controller storage network interfaces HX Storage Cluster roaming storage interface	3172
hx-vm-data	Guest VM network interfaces	3174
hx-vmotion	ESXi host vMotion vmkernel interfaces	3173
Native-VLAN	VLAN to which untagged frames are assigned	2

Table 3 Configuration Variables

Variable	Description
<<var_password>>	Global default administrative password
<<var_nexus_A_hostname>>	Cisco Nexus A host name
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway
<<var_nexus_B_hostname>>	Cisco Nexus B host name
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway

Network Switch Configuration

This section provides the detailed steps to configure the Cisco Nexus 9000s to use in a Veeam environment.



Follow these steps precisely because failure to do so could result in an improper configuration.

For detailed configuration details, refer to: [Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide](#)

Physical Connectivity

Follow the physical connectivity guidelines as covered in section [Deployment Architecture](#).

Cisco Nexus Base

This section describes how to configure the Cisco Nexus switches to use in a base Veeam environment. This procedure assumes that you are using Cisco Nexus 9000 7.0(3)I4(2).



The following procedure includes setting up the NTP distribution on the In-Band Management VLAN. The interface-vlan feature and NTP commands are used in this set up. This procedure also assumes the default VRF will be used to route the In-Band Management VLAN.

Set Up Initial Configuration

Cisco Nexus 9372PX A

To set up the initial configuration for the Cisco Nexus A switch on <<var_nexus_A_hostname>>, complete the following steps:

Configure the Switch



On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes

Do you want to enforce secure password standard (yes/no): yes

Enter the password for "admin": <<var_password>>

Confirm the password for "admin": <<var_password>>

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name: <<var_nexus_A_hostname>>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address: <<var_nexus_A_mgmt0_ip>>

Mgmt0 IPv4 netmask: <<var_nexus_A_mgmt0_netmask>>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway: <<var_nexus_A_mgmt0_gw>>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter

Configure the ntp server? (yes/no) [n]: y

NTP server IPv4 address: <<var_global_ntp_server_ip>>

Configure default interface layer (L3/L2) [L2]: Enter

Configure default switchport interface state (shut/noshut) [noshut]: shut

Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter

Would you like to edit the configuration? (yes/no) [n]: Enter

1. Review the configuration summary before enabling the configuration.

Use this configuration and save it? (yes/no) [y]: Enter

Cisco Nexus 9372PX B

To set up the initial configuration for the Cisco Nexus B switch on <<var_nexus_B_hostname>>, complete the following steps:

Configure the Switch



On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes

Do you want to enforce secure password standard (yes/no): yes

Enter the password for "admin": <<var_password>>

Confirm the password for "admin": <<var_password>>

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name: <<var_nexus_B_hostname>>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address: <<var_nexus_B_mgmt0_ip>>

Mgmt0 IPv4 netmask: <<var_nexus_B_mgmt0_netmask>>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway: <<var_nexus_B_mgmt0_gw>>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter

Configure the ntp server? (yes/no) [n]: y

NTP server IPv4 address: <<var_global_ntp_server_ip>>

Configure default interface layer (L3/L2) [L2]: Enter

Configure default switchport interface state (shut/noshut) [noshut]: shut

Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter

Would you like to edit the configuration? (yes/no) [n]: Enter

2. Review the configuration summary before enabling the configuration.

Use this configuration and save it? (yes/no) [y]: Enter

Cisco Nexus Switch Configuration

The steps elaborated below, are for reference. In Case, HyperFlex and Cisco Nexus 9000 is already configured, then the below steps are not required.

Enable Licenses

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To license the Cisco Nexus switches, complete the following steps:

1. Log in as admin.
2. Run the following commands:

```
config t
feature interface-vlan
feature lacp
feature vpc
```

```
feature lldp
```

Set Global Configurations

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To set global configurations, complete the following step on both switches:

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <<var_global_ntp_server_ip>> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <<var_ib-mgmt_vlan_gateway>>
copy run start
```

Create VLANs

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To create the necessary virtual local area networks (VLANs), complete the following step on both switches:

1. From the global configuration mode, run the following commands:

```
vlan <<var_veeam-network>>
name hx-inband-mgmt
exit
vlan <<var_native_vlan_id>>
name Native-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus 9372PX A

1. From the global configuration mode, run the following commands:

```
ntp source <<var_switch_a_ntp_ip>>
interface Vlan <<var_ib-mgmt_vlan_id>>
ip address <<var_switch_a_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>
no shutdown
exit
```

Cisco Nexus 9372PX B

1. From the global configuration mode, run the following commands:

```
ntp source <<var_switch_b_ntp_ip>>
interface Vlan<<var_ib-mgmt_vlan_id>>
ip address <<var_switch_b_ntp_ip>>/<<var_ib-mgmt_vlan_netmask_length>>
no shutdown
exit
```

Create Port Channels

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To create the necessary port channels between devices, complete the following step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
exit

interface Eth1/51-52
channel-group 10 mode active
no shutdown
exit

interface Po13
description <<var_ucs_6248_clustername>>-a
exit

interface Eth1/25
channel-group 13 mode active
no shutdown
exit

interface Po14
description <<var_ucs_6248_clustername>>-b
exit

interface Eth1/26
channel-group 14 mode active
no shutdown
exit
```

copy run start

Configure Port Channel Parameters



Cisco virtual PortChannels (vPC) allows a device to connect to two different physical Cisco Nexus switches using a single logical Cisco PortChannel interface.

Cisco Nexus 9372PX A and Cisco Nexus 9372PX B

To configure port channel parameters, complete the following step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan 1
switchport trunk allowed vlan << var_veeam-network >>
spanning-tree port type network
exit
interface Po13
switchport mode trunk
switchport trunk native vlan 1
switchport trunk allowed vlan << var_veeam-network
spanning-tree port type edge trunk
mtu 9216
exit
interface Po14
switchport mode trunk
switchport trunk native vlan 1
switchport trunk allowed vlan << var_veeam-network
spanning-tree port type edge trunk
mtu 9216
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus 9372PX A

To configure virtual port channels (vPCs) for switch A, complete the following step:

1. From the global configuration mode, run the following commands:

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 10
peer-keepalive destination <<var_nexus_B_mgmt0_ip>> source <<var_nexus_A_mgmt0_ip>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
exit
interface Po10
vpc peer-link
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Cisco Nexus 9372PX B

To configure vPCs for switch B, complete the following step:

1. From the global configuration mode, run the following commands.

```
vpc domain <<var_nexus_vpc_domain_id>>
role priority 20
peer-keepalive destination <<var_nexus_A_mgmt0_ip>> source <<var_nexus_B_mgmt0_ip>>
peer-switch
peer-gateway
auto-recovery
delay restore 150
exit
interface Po10
vpc peer-link
```



```
exit
interface Po13
vpc 13
exit
interface Po14
vpc 14
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the environment. If an existing Cisco Nexus environment is present, it is recommended to use vPCs to uplink the Cisco Nexus 9372PX switches included in the present environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Cisco UCS S3260 Storage Server and HyperFlex Configuration

Cisco UCS Base Configuration

This deployment details the configuration steps for the Cisco UCS 6248UP Fabric Interconnects (FI) in a design that supports both HyperFlex (HX) Cluster and Cisco UCS S3260 Storage Server. The base configuration of Cisco UCS will remain similar for both Cisco HX and Cisco UCS S3260 Storage Server.

Perform Initial Setup of Cisco UCS 6248UP Fabric Interconnects

This section describes the steps to configure the Cisco Unified Computing System (Cisco UCS) to use in a HyperFlex environment. These steps are necessary to provision the Cisco HyperFlex and Cisco UCS S3260 Storage Server and should be followed precisely to avoid improper configuration.

Cisco UCS 6248UP Fabric Interconnect A

To configure Fabric Interconnect A, complete the following steps:

1. Make sure the Fabric Interconnect cabling is properly connected, including the L1 and L2 cluster links, and power the Fabric Interconnects on by inserting the power cords.
2. Connect to the console port on the first Fabric Interconnect, which will be designated as the A fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.
3. Start your terminal emulator software.
4. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, 1 stop bit.
5. Open the connection just created. You may have to press ENTER to see the first prompt.
6. Configure the first Fabric Interconnect, using the following example as a guideline:

```
---- Basic System Configuration Dialog ----
```

```
This setup utility will guide you through the basic configuration of
the system. Only minimal configuration including IP connectivity to
the Fabric interconnect and its clustering mode is performed through these steps.
```

```
Type Ctrl-C at any time to abort configuration and reboot system.
To back track or make modifications to already entered values,
complete input till end of section and answer no when prompted
to apply configuration.
```

```
Enter the configuration method. (console/gui) ? console
```

```
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
```

```
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
```

```
Enforce strong password? (y/n) [y]: y
```

```
Enter the password for "admin":
```

```
Confirm the password for "admin":
```

```
Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes
```

```

Enter the switch fabric (A/B) []: A

Enter the system name:  FI-HX1

Physical Switch Mgmt0 IP address : 10.29.149.98

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.29.149.1

Cluster IPv4 address : 10.29.149.100

Configure the DNS Server IP address? (yes/no) [n]: yes

    DNS IP address : 171.70.168.183

Configure the default domain name? (yes/no) [n]: yes

    Default domain name : hxl.lab.cisco.com

Join centralized management environment (UCS Central)? (yes/no) [n]: no

Following configurations will be applied:

Switch Fabric=A
System Name=FI-HX1
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=10.29.149.98
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.29.149.1
Ipv6 value=0
DNS Server=171.70.168.183
Domain Name=hxl.lab.cisco.com

Cluster Enabled=yes
Cluster IP Address=10.29.149.1
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.

Configuration file - Ok

```

Cisco UCS 6248UP Fabric Interconnect B

To configure Fabric Interconnect B, complete the following steps:

1. Connect to the console port on the first Fabric Interconnect, which will be designated as the B fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.
2. Start your terminal emulator software.
3. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, 1 stop bit.
4. Open the connection just created. You may have to press ENTER to see the first prompt.
5. Configure the second Fabric Interconnect, using the following example as a guideline:

```

---- Basic System Configuration Dialog ----

```

This setup utility will guide you through the basic configuration of the system. Only minimal configuration including IP connectivity to the Fabric interconnect and its clustering mode is performed through these steps.

Type Ctrl-C at any time to abort configuration and reboot system.
To back track or make modifications to already entered values, complete input till end of section and answer no when prompted to apply configuration.

Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? y

Enter the admin password of the peer Fabric interconnect:

```
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.149.98
Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
Cluster IPv4 address           : 10.29.149.100
```

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address
Physical Switch Mgmt0 IP address : 10.29.149.99

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.

Configuration file - Ok

Cisco UCS Setup

Log in to Cisco UCS Manager

To log in to the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS fabric interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 3.1(2b)

This document assumes you are using Cisco UCS 3.1(2b). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.1(2b), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To create anonymous reporting, complete the following step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products:

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.

If you decide to enable this feature in future, you can do so from the "Anonymous Reporting" in the Call Home settings under the Admin tab.

[View Sample Data](#)

Do you authorize the disclosure of this information to Cisco Smart CallHome?

☐ Yes ☐ No

☐ Don't show this message again.

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for in band server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

Create Block of IPv4 Addresses

From : 10.29.149.80 Size : 16

Subnet Mask : 255.255.255.0 Default Gateway : 10.29.149.1

Primary DNS : 0.0.0.0 Secondary DNS : 0.0.0.0

OK Cancel

- Click OK to create.
- Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

- In Cisco UCS Manager, click the Admin tab in the navigation pane.
- Select All > Timezone Management.

All

Collection Policy Chassis
Collection Policy Flex
Collection Policy Host
Collection Policy Port
Collection Policy Server

▼ fabric

All / Time Zone Management / Timezone

General Events

Actions

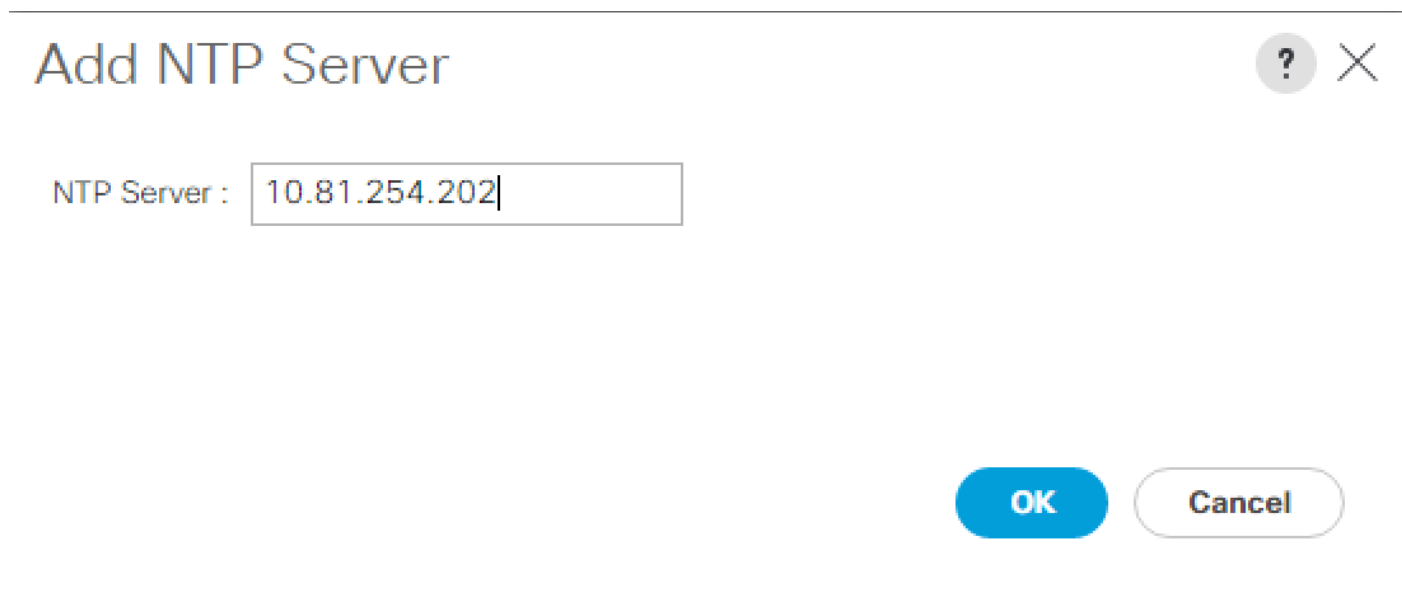
Add NTP Server

Properties

Time Zone : America/Los_Angeles (Pacific Time)

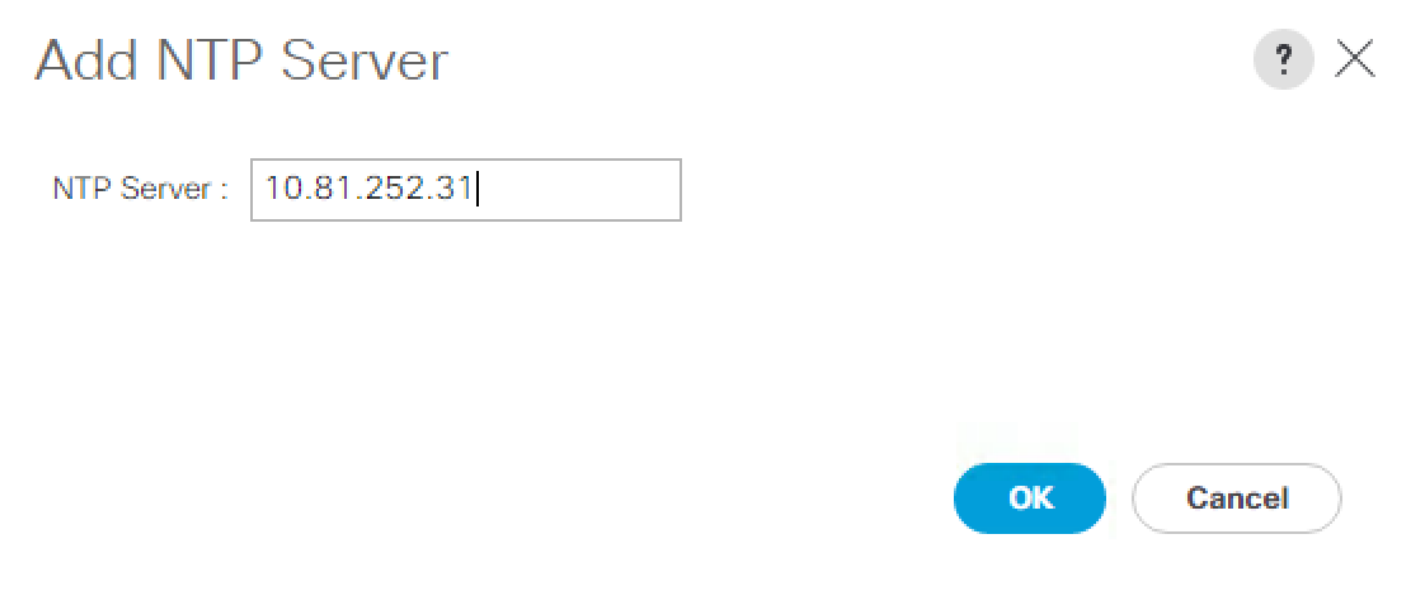
NTP Server : America/Los_Angeles (Pacific Time)

- In the Properties pane, select the appropriate time zone in the Timezone menu.
- Click Save Changes and then click OK.
- Click Add NTP Server.
- Enter <<var_switch_a_ntp_ip>> and click OK.



A dialog box titled "Add NTP Server" with a question mark icon and a close button (X) in the top right corner. The main content area contains the text "NTP Server :" followed by a text input field containing the IP address "10.81.254.202". At the bottom right, there are two buttons: a blue "OK" button and a white "Cancel" button with a grey border.

7. Click Add NTP Server.
8. Enter <<var_switch_b_ntp_ip>> and click OK.



A dialog box titled "Add NTP Server" with a question mark icon and a close button (X) in the top right corner. The main content area contains the text "NTP Server :" followed by a text input field containing the IP address "10.81.252.31". At the bottom right, there are two buttons: a blue "OK" button and a white "Cancel" button with a grey border.

9. Click OK.

Uplink Ports

The Ethernet ports of a Cisco UCS 6248UP Fabric Interconnect are all capable of performing several functions, such as network uplinks or server ports, and more. By default, all ports are unconfigured, and their function must be defined by the administrator. To define the specified ports to be used as network uplinks to the upstream network, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Fabric Interconnects > Fabric Interconnect A > Fixed Module or Expansion Module 2 as appropriate > Ethernet Ports

3. Select the ports that are to be uplink ports, right click them, and click Configure as Uplink Port.
4. Click Yes to confirm the configuration and click OK.
5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module 2 as appropriate > Ethernet Ports
6. Select the ports that are to be uplink ports, right-click them, and click Configure as Uplink Port.
7. Click Yes to confirm the configuration and click OK.
8. Verify all the necessary ports are now configured as uplink ports.

Equipment / Fabric Interconnects **Fabric Interconnect A (subordinate)** /

Fabric Interconnects						IO Modules	Thermal	Power	Fans	Installed Firmware	Faults	Events	Performance
+ - Advanced Filter Export Print													
Name	Address	IF Role	IF Type	Overall Status	Admin State								
Port 19	8C:60:4F:BF:0D:7A	Server	Physical	Sfp Not Present	Enabled								
Port 20	8C:60:4F:BF:0D:7B	Server	Physical	Sfp Not Present	Enabled								
Port 21	8C:60:4F:BF:0D:7C	Server	Physical	Sfp Not Present	Enabled								
Port 22	8C:60:4F:BF:0D:7D	Server	Physical	Sfp Not Present	Enabled								
Port 23	8C:60:4F:BF:0D:7E	Server	Physical	Sfp Not Present	Enabled								
Port 24	8C:60:4F:BF:0D:7F	Server	Physical	Sfp Not Present	Enabled								
Port 25	8C:60:4F:BF:0D:80	Network	Physical	Up	Enabled								
Port 26	8C:60:4F:BF:0D:81	Network	Physical	Up	Enabled								
Port 27	8C:60:4F:BF:0D:82	Unconfigured	Physical	Sfp Not Present	Disabled								
Port 28	8C:60:4F:BF:0D:83	Unconfigured	Physical	Sfp Not Present	Disabled								
Port 29	8C:60:4F:BF:0D:84	Unconfigured	Physical	Sfp Not Present	Disabled								
Port 30	8C:60:4F:BF:0D:85	Unconfigured	Physical	Sfp Not Present	Disabled								
Port 31	8C:60:4F:BF:0D:86	Unconfigured	Physical	Sfp Not Present	Disabled								
Port 32	8C:60:4F:BF:0D:87	Unconfigured	Physical	Sfp Not Present	Disabled								
FC Ports													
Fabric Interconnect B (p...													

Uplink Port Channels

If the Cisco UCS uplinks from one Fabric Interconnect are to be combined into a port channel or vPC, you must separately configure the port channels using the previously configured uplink ports. To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Under LAN > LAN Cloud, click to expand the Fabric A tree.
3. Right-click Port Channels underneath Fabric A and select Create Port Channel.
4. Enter the port channel ID number as the unique ID of the port channel.
5. Enter the name of the port channel.

6. Click Next.
7. Click on each port from Fabric Interconnect A that will participate in the port channel, and click the >> button to add them to the port channel.
8. Click Finish.
9. Click OK.
10. Under LAN > LAN Cloud, click to expand the Fabric B tree.
11. Right-click Port Channels underneath Fabric B and select Create Port Channel.
12. Enter the port channel ID number as the unique ID of the port channel.
13. Enter the name of the port channel.
14. Click Next.
15. Click on each port from Fabric Interconnect B that will participate in the port channel, and click the >> button to add them to the port channel.
16. Click Finish.
17. Click OK.
18. Verify the necessary port channels have been created. It can take a few minutes for the newly formed port channels to converge and come online.

The screenshot displays the Cisco UCS Management Center interface. On the left, a navigation pane shows the hierarchy: LAN > LAN Cloud > Fabric A > Port Channels. The selected item is 'Port-Channel 11 NX9K-vPC-A'. The main content area shows the configuration for this port channel.

General Tab:

- Status:** Overall Status is **Up** (green arrow icon). Additional Info is empty.
- Actions:** Enable Port Channel, Disable Port Channel, Add Ports.
- Properties:**
 - ID: 11
 - Fabric ID: A
 - Port Type: Aggregation
 - Transport Type: Ether
 - Name: NX9K-vPC-A
 - Description: (empty field)
 - Flow Control Policy: default
 - LACP Policy: default
 - Note: Changing LACP policy may flap the port-channel if the suspend-individual value change
 - Admin Speed: 1 Gbps, 10 Gbps (selected), 40 Gbps
 - Operational Speed(Gbps): 20

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and of additional fabric extenders for further C-Series connectivity. To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and select Equipment in the list on the left.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.

5. Click Save Changes.
6. Click OK.

Server Ports

The Ethernet ports of a Cisco UCS Fabric Interconnect connected to the rack-mount servers, or to the blade chassis or to Cisco UCS S3260 Storage Server must be defined as server ports. When a server port is activated, the connected server or chassis will begin the discovery process shortly afterwards. Rack-mount servers, blade chassis, and S3260 chassis are automatically numbered in the order that they are first discovered. For this reason, it is important to configure the server ports sequentially in the order you wish the physical servers and/or chassis to appear within UCS manager. For example, if you installed your servers in a cabinet or rack with server #1 on the bottom, counting up as you go higher in the cabinet or rack, then you need to enable the server ports to the bottom-most server first, and enable them one-by-one as you move upward. You must wait until the server appears in the Equipment tab of Cisco UCS Manager before configuring the ports for the next server. The same numbering procedure applies to blade server chassis.

To define the specified ports to be used as server ports, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane.
2. Select Fabric Interconnects > Fabric Interconnect A > Fixed Module or Expansion Module 2 as appropriate > Ethernet Ports.
3. Select the first port that is to be a server port, right-click it, and click Configure as Server Port.
4. Click Yes to confirm the configuration and click OK.
5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module 2 as appropriate > Ethernet Ports.
6. Select the matching four ports as chosen for Fabric Interconnect A that will be configured as Server Port.



Cisco UCS S3260 storage server has two 40Gbps port per server node whereas Cisco UCS FI6248UP has 10 Gbps ports. Therefore, each of the 40 Gb ports on S3260 is connected to four 10Gbps ports on each of Fabric Interconnect 6248UP, through a QSFP to four SFP+ active optical breakout cable.

7. Click Yes to confirm the configuration, and click OK.
8. Repeat step 6-8 for Fabric Interconnect B
9. Wait for a brief period, until the rack-mount server appears in the Equipment tab underneath Equipment > Rack Mounts > Servers, or the chassis appears underneath Equipment > Chassis.

Name	Slot	Port ID	MAC	If Role	If Type	Overall Sta...	Admin State
Port 1	1	1	8C:60:4F:...	Server	Physical	Up	Enabled
Port 2	1	2	8C:60:4F:...	Server	Physical	Up	Enabled
Port 3	1	3	8C:60:4F:...	Server	Physical	Up	Enabled
Port 4	1	4	8C:60:4F:...	Server	Physical	Up	Enabled
Port 5	1	5	8C:60:4F:...	Server	Physical	Admin Do...	Disabled
Port 6	1	6	8C:60:4F:...	Server	Physical	Admin Do...	Disabled
Port 7	1	7	8C:60:4F:...	Server	Physical	Admin Do...	Disabled
Port 8	1	8	8C:60:4F:...	Server	Physical	Admin Do...	Disabled

Server Discovery


As previously described, when the server ports of the Fabric Interconnects are configured and active, the servers connected to those ports will begin a discovery process. During discovery, the servers' internal hardware inventories are collected, along with their current firmware revisions. Before continuing with the HyperFlex and Cisco UCS S3260 storage server installation processes, wait for all of the servers to finish their discovery process and show as unassociated servers that are powered off, with no errors. To view the servers' discovery status, complete the following steps:

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane, and click Equipment in the top of the navigation tree on the left.
2. In the properties pane, click the Servers tab.
3. Click the Chassis > Chassis1 Tab and view the Chassis status in the Overall Status column.
4. When the chassis is discovered, the S3260 storage server is displayed as shown below:

The screenshot displays the Cisco UCS Manager interface for Chassis 1. The left-hand navigation pane shows the hierarchy: Equipment > Chassis > Chassis 1. The main content area is divided into several sections:

- Equipment / Chassis / Chassis 1**: The breadcrumb path at the top.
- General**: The active tab, showing:
 - Fault Summary**: Four status icons (all green) with a count of 0 for each.
 - Status**: Overall Status is **Operable** (green arrow icon).
 - Status Details**:
 - Configuration State: **OK** (green arrow icon)
 - Operability: **N/A** (red X icon)
 - Power: **OK** (green arrow icon)
 - Thermal: **N/A** (red X icon)
 - Actions**:
 - Associate Chassis Profile
 - Acknowledge Chassis
 - Decommission Chassis
- Physical Display**: A photograph of the Cisco UCSC C3X60 server rack.
- Properties**:
 - ID: 1
 - Product Name: Cisco UCSC C3X60
 - Vendor: Cisco Systems Inc
 - PID: UCSC-C3X60
 - Revision: 0
 - Serial: FOX2033G4MA
 - Chassis Profile:

5. Click the Rack-Mount Servers or Storage Server sub-tab as appropriate, and view the servers' status in the Overall Status column. Below are HX Servers for four node Hyper Flex Cluster:



HyperFlex Installation

The Cisco HyperFlex software is distributed as a deployable virtual machine, contained in an Open Virtual Appliance (OVA) file format. The HyperFlex OVA file is available for download at www.cisco.com.

When the OVA is installed, the HyperFlex installer is accessed via a webpage using your local computer and a web browser. The HyperFlex Installer configures Cisco UCS and deploys the HyperFlex Data Platform on a Cisco UCS Cluster. To configure Cisco UCS for HyperFlex and then install HyperFlex Data Platform, refer to: [Cisco HyperFlex Deployment Guide](#).

As shown in Figure 24, the present setup is deployed with four node HyperFlex Cluster with HX220C-M4 nodes.

Figure 24 Cisco UCS Manager Summary for a Four Node HX Cluster

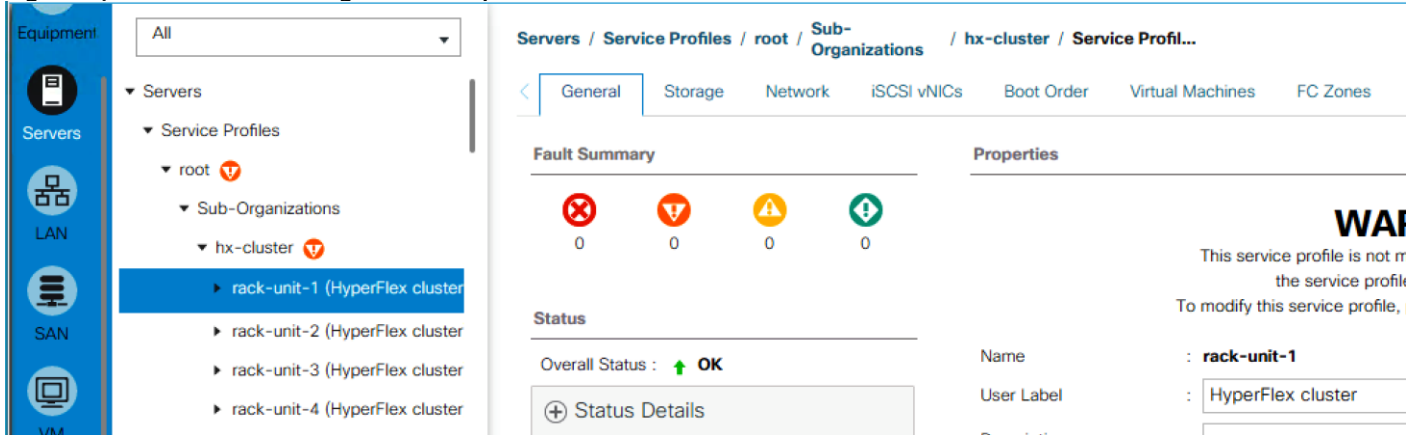
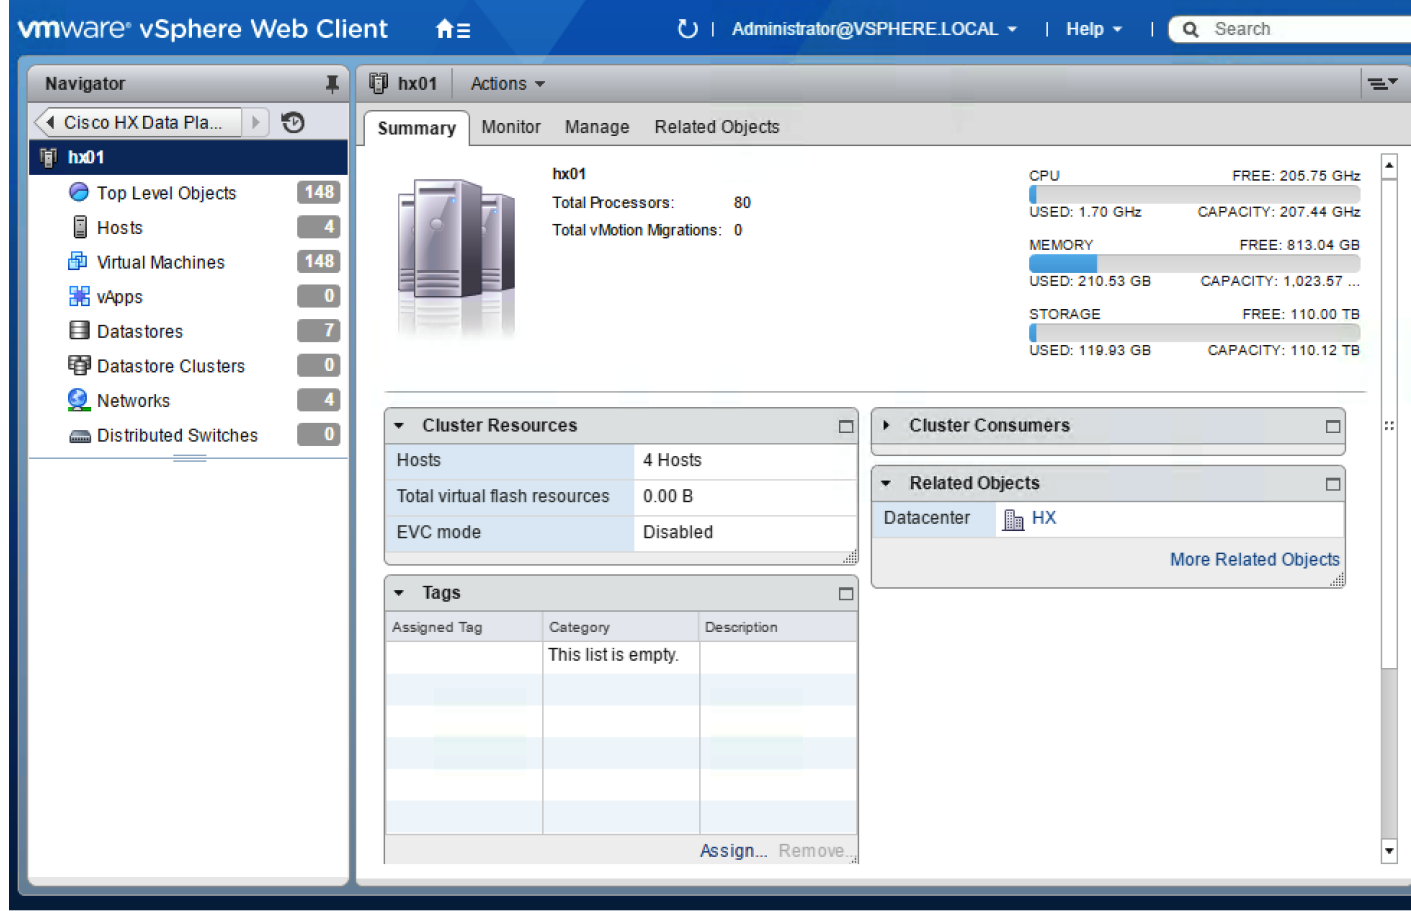


Figure 25 details the HyperFlex Cluster summary through the HX vCenter Plugin.

Figure 25 HX Cluster Summary



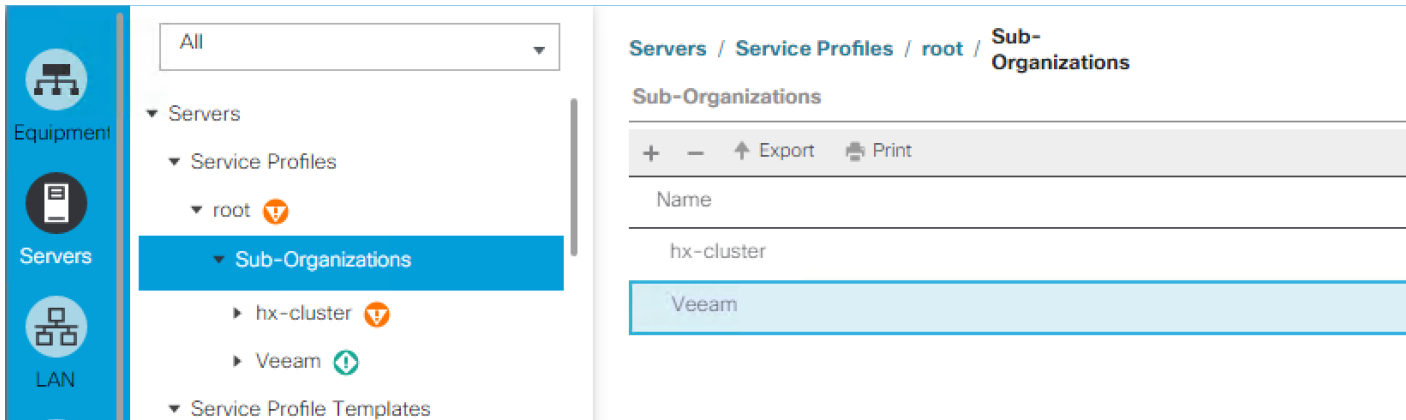
Cisco UCS S3260 Configuration

This section details the steps for the Cisco UCS configuration for the S3260 Storage Server. These steps are independent of the Cisco UCS configuration for HX Cluster.

Create Sub-Organization

In this setup, there are two sub-organizations under the root, each for HX and Veeam Infrastructure. Sub-organizations help to restrict user access to logical pools and objects in order to facilitate security and to provide easier user interaction. For Veeam Backup infrastructure, create a sub-organization as "Veeam." To create a sub-organization, complete the following steps:

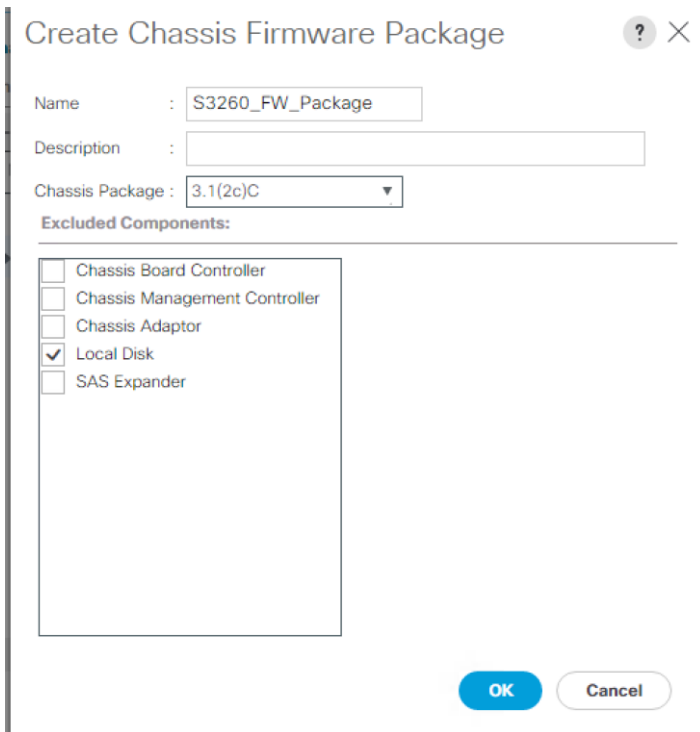
1. In the Navigation pane, click the Servers tab.
2. In the Servers tab, expand Service Profiles > root. You can also access the Sub-Organizations node under the Policies or Pools nodes.
3. Right-click Sub-Organizations and choose Create Organization.



Create Chassis Firmware Packages

To create S3260 Chassis Firmware packages, complete the following steps:

1. In the Navigation pane, click the Chassis tab.
2. In the Chassis tab, expand Policies > root > sub-Organizations > Veeam.
3. Right-click Chassis Firmware Packages and select Create Chassis Firmware Packages.
4. Enter S3260_FW_Package as the Package name.
5. Select 3.1(2c)C from the Chassis Package drop-down.
6. Click OK.



Create Disk Zoning Policy

You can assign disk drives to the server nodes using disk zoning. Disk zoning can be performed on the controllers in the same server or on the controllers on different servers.

To create S3260 Disk Zoning Policy, complete the following steps:

1. In the Navigation pane, click Chassis.
2. Expand Policies > root > Sub-Organizations > Veeam.
3. Right-click Disk Zoning Policies and choose Create Disk Zoning Policy.

Create Disk Zoning Policy

Name : S3260_DiskZone

Description :

Preserve Config : ☐

Disk Zoning Information

+ - Advanced Filter ↑ Export Print

Name	Slot Number	Ownership	Assigned to Ser...	Assigned to Con...	Controller Type
No data available					

+ Add

Delete

Modify

OK

Cancel

4. Enter S3260_DiskZone as the Disk Zone Name
5. In the Disk Zoning Information Area, Click Add
6. Select Ownership as Dedicated
7. Select Server as 1 (Disk would be assigned to node 1 of S3260 Storage server)
8. Select Controller as 1
9. Slot range as 1-28 (in the present setup, there are 28 X6 TB SAS drives).

Add Slots to Policy



Ownership : ☐ Unassigned ☒ Dedicated ☐ Shared ☐ Chassis Global Hot Spare

Server :

Controller :

Controller Type : **SAS**

Slot Range :

OK

Cancel

10. Click OK
11. Click OK Again to complete the Disk Zoning Configuration Policy.

Create Disk Zoning Policy

Name

S3260_DiskZone

Description

Preserve Config

☐

Disk Zoning Information

+

-

Advanced Filter

Export

Print

Name	Slot Number	Ownership	Assigned to Ser...	Assigned to Con...	Controller Type
▶ disk-slot-1	1	Dedicated			
▶ disk-slot-2	2	Dedicated			
▶ disk-slot-3	3	Dedicated			
▶ disk-slot-4	4	Dedicated			
▶ disk-slot-5	5	Dedicated			
▶ disk-slot-6	6	Dedicated			

+

Add

Delete

Modify

OK

Cancel

Setting S3260 Disk to Unconfigured Good

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:



To prepare all disks from the S3260 Storage Servers for storage profiles, the disks have to be converted from JBOD to Unconfigured Good. To convert the disks, complete the steps below.

1. Select the Equipment tab in the left pane of the Cisco UCS Manager GUI.
2. Go to Equipment > Chassis > Chassis 1 > Storage Enclosures > Enclosure1
3. Select disks and right-click Set JBOD to Unconfigured Good.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > Sub-organizations > Veeam.



In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.

4. Select Create MAC Pool to create the MAC address pool.
5. Enter `MAC_Pool_A` as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select **Sequential** as the option for Assignment Order.

Create MAC Pool

Name :

Description :

Assignment Order : ☐ Default ☒ Sequential

< Prev **Next >** Finish Cancel

8. Click Next.
9. Click Add.
10. Specify a starting MAC address.



For the present solution, the recommendation is to place `A0` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.

Create a Block of MAC Addresses

First MAC Address : 00:25:B5:55:A0:00

Size : 32

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

OK

Cancel

12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.
15. Right-click MAC Pools under the root organization.
16. Select Create MAC Pool to create the MAC address pool.
17. Enter `MAC_Pool_B` as the name of the MAC pool.
18. Optional: Enter a description for the MAC pool.

Create MAC Pool

1 Define Name and Description

2 Add MAC Addresses

Name : MAC_Pool_B

Description :

Assignment Order : ☐ Default ☒ Sequential

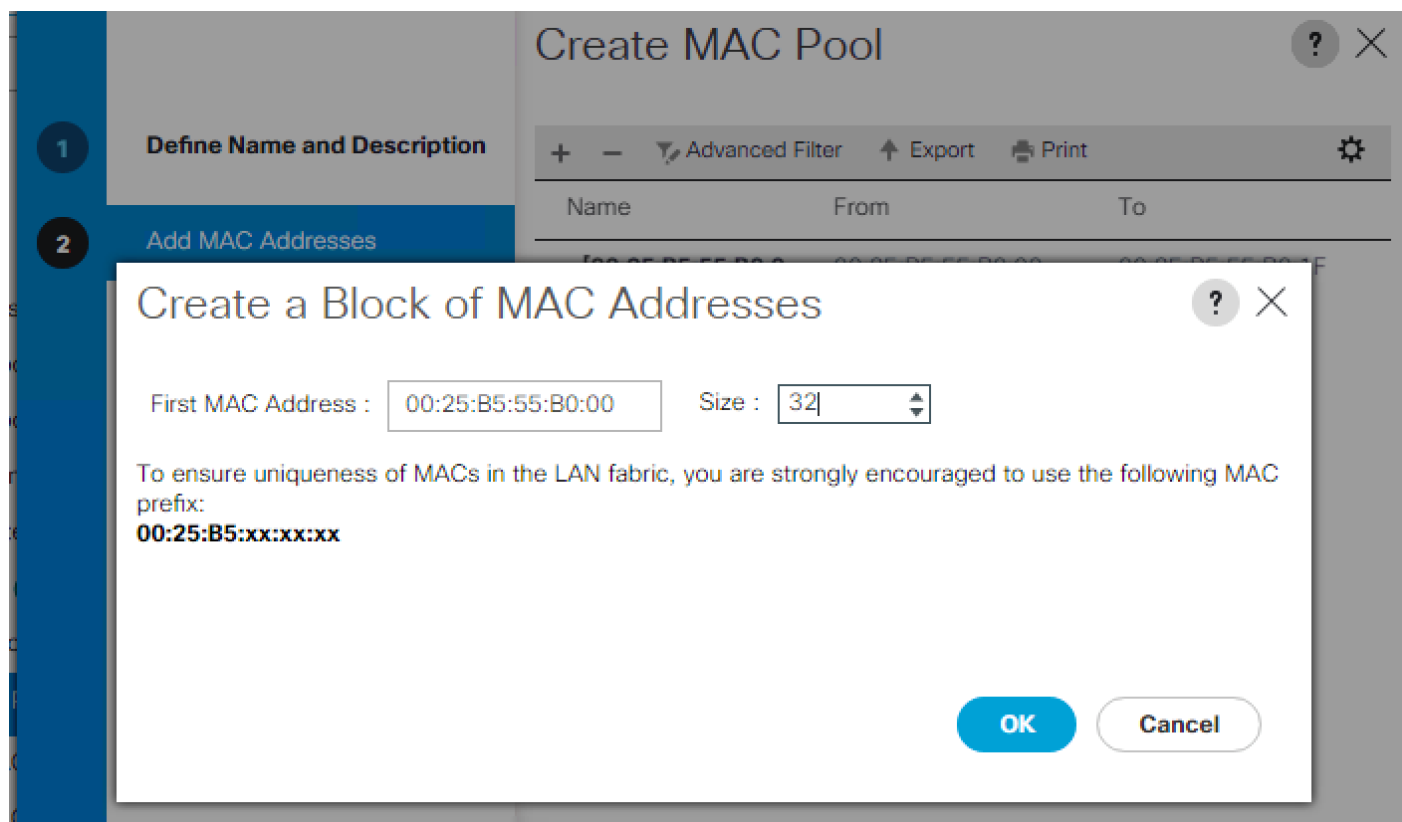
< Prev Next > **Finish** Cancel

19. Click Next.
20. Click Add.
21. Specify a starting MAC address.



For the present solution, it is recommended to place B0 in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses.

22. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.



23. Click OK.
24. Click Finish.
25. In the confirmation message, click OK.

Create UUID Suffix Pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root > Sub-Organizations > Veeam.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID_Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.
9. Click Next.

10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available server resources.

Create a Block of UUID Suffixes

From :

Size :

OK

Cancel

13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter Infra_Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select S3260 server node to be used for and click >> to add them to the Infra_Pool server pool.

Properties for: Server Pool S3260-Pool

General

Servers

Faults

Events

Actions

Delete

Add Servers

Show Pool Usage

Name : **S3260-Pool**

Description :

Size : **1**

Assigned : **0**

Pool Policies

Advanced Filter

Export

Print

Name

Description

Qualification

No data available

9. Click Finish.

10. Click OK.

Create VLANs



If HyperFlex is already configured, this step is not required. HX Management and Veeam network are on the same VLAN.

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter hx-inband-mgmt as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Keep the Sharing Type as None.
8. Click OK and then click OK again.

Create Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organizations > Veeam.
3. Expand Host Firmware Packages.
4. Right-click and Select Create Host Firmware Package.
5. Select the version 3.1(2b)B for Blade and 3.1(2b)C for Rack Packages.

Create Host Firmware Package

Name : S3260Firmware

Description :

How would you like to configure the Host Firmware Package?

☒ Simple ☐ Advanced

Blade Package : 3.1(2b)B

Rack Package : 3.1(2b)C

Excluded Components:

- ☐ Adapter
- ☐ Host NIC Option ROM
- ☐ CIMC
- ☐ Board Controller
- ☐ Flex Flash Controller
- ☐ BIOS
- ☐ PSU
- ☐ SAS Expander
- ☐ Storage Controller Onboard Device
- ☐ Storage Device Bridge
- ☐ GPUs
- ☐ FC Adapters

- Click OK to add the host firmware package.

QoS Policy

To enable quality of service and create QoS policy in the Cisco UCS fabric, complete the following steps:

- In Cisco UCS Manager, click the LAN tab in the navigation pane.
- Select LAN > LAN Cloud > QoS System Class.
- In the right pane, click the General tab.
- Make sure the Class of Service is configured as shown below. This is already created through the HyperFlex installer.

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU
Platinum	<input checked="" type="checkbox"/>	5	<input type="checkbox"/>	4	25	9216
Gold	<input checked="" type="checkbox"/>	4	<input checked="" type="checkbox"/>	4	25	normal
Silver	<input checked="" type="checkbox"/>	2	<input checked="" type="checkbox"/>	best-effort	6	normal
Bronze	<input checked="" type="checkbox"/>	1	<input checked="" type="checkbox"/>	best-effort	6	9216
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	best-effort	6	normal
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	32	fc



The QoS System Class describes are in alignment with the HyperFlex Infrastructure configuration.

- In Cisco UCS Manager, click the LAN tab in the navigation pane.
- Select LAN > Policies > root > Sub-Organizations> Veeam >QoS Policies.
- Right-click QoS Policies and select Create QoS Policy.
- Enter the name as 'Silver' and select Silver in the priority.
- Click OK.

Create QoS Policy

?
×

Name :

Egress

Priority :

Burst(Bytes) :

Rate(Kbps) :

Host Control : ☒ None ☐ Full

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root >Sub-Organization > Veeam.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Veeam_NCP as the policy name.
6. For CDP, select the Enabled option.
7. Click OK to create the network control policy.

Create Network Control Policy

?

×

Name : Veeam_NCP

Description :

CDP : ☐ Disabled ☒ Enabled

MAC Register Mode : ☒ Only Native Vlan ☐ All Host Vlans

Action on Uplink Fail : ☒ Link Down ☐ Warning

MAC Security

Forge : ☒ Allow ☐ Deny

LLDP

OK

Cancel

- Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

- In Cisco UCS Manager, click the Servers tab in the navigation pane.
- Select Policies > root >Sub-Organizations >Veeam.
- Right-click Power Control Policies.
- Select Create Power Control Policy.
- Enter No-Power-Cap as the power control policy name.
- Change the power capping setting to No Cap.
- Click OK to create the power control policy.
- Click OK.

Create Power Control Policy

Name

:

No-Power-Cap

Description

:

Fan Speed Policy :

Any

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

☒ No Cap
 ☐ cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK

Cancel

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > sub-Organizations >Veeam.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter S3260-BIOS as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.

7. Change Consistent Device Naming to enabled.
8. Click Finish to create the BIOS policy.

Create BIOS Policy

Name : S3260-BIOS

Description :

Reboot on BIOS Settings Change : ☐

Quiet Boot : ☒ disabled ☐ enabled ☐ Platform Default

Post Error Pause : ☐ disabled ☐ enabled ☒ Platform Default

Resume Ac On Power Loss : ☐ stay-off ☐ last-state ☐ reset ☒ Platform Default

Front Panel Lockout : ☐ disabled ☐ enabled ☒ Platform Default

Consistent Device Naming : ☐ disabled ☒ enabled ☐ Platform Default

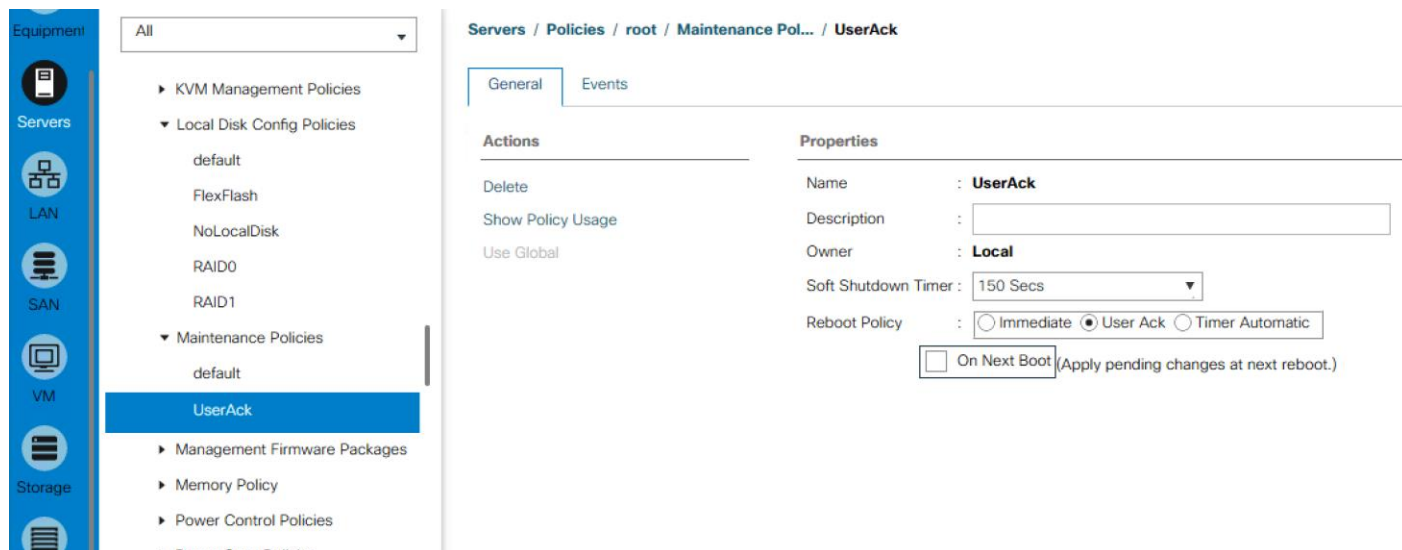
< Prev Next > Finish

9. Click OK.

Create Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Maintenance Policies and Select Create Maintenance Policy.
4. Change the Reboot Policy to User Ack.
5. (Optional: Click “On Next Boot” to delegate maintenance windows to server owners)



6. Click Save Changes.
7. Click OK to accept the change.

Create Adaptor Policy

To create adaptor policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Sub-Organizations > Veeam.
3. Right Click on Adaptor Policies and Select Ethernet Adaptor Policy.
4. Enter name as veeam_adaptorpol.
5. Enter Transmit Queues = Receive Queues = 8, Ring Size = 4096.
6. Enter Completion Queues = 16 and Interrupts = 32.
7. Under Options, ensure Receive Side Scaling (RSS) is enabled.
8. Click OK.

Create Ethernet Adapter Policy



Resources

Transmit Queues	:	<input type="text" value="8"/>	[1-1000]
Ring Size	:	<input type="text" value="4096"/>	[64-4096]

Receive Queues	:	<input type="text" value="8"/>	[1-1000]
Ring Size	:	<input type="text" value="4096"/>	[64-4096]

Completion Queues	:	<input type="text" value="16"/>	[1-2000]
Interrupts	:	<input type="text" value="32"/>	[1-1024]

Options

Transmit Checksum Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Receive Checksum Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
TCP Segmentation Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
TCP Large Receive Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Receive Side Scaling (RSS)	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Accelerated Receive Flow Steering	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled

OK

Cancel



To enable maximum throughput, it is recommended to change the default size of Rx and Tx Queues. RSS should be enabled, since it allows the distribution of network receive processing across multiple CPUs in a multiprocessor system.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 2 vNIC Templates will be created.



The Infra VLANs were used on S3260 deployed in this environment. The HyperFlex Infra creates its own vNIC templates through HX Installer

Create Data vNICs

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Policies > root > Sub-Organizations > Veeam.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.

5. Enter vNIC_Template_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Updating Template as the Template Type.
9. Select Redundancy Type as No Redundancy
10. Under VLANs, select the checkbox for hx-inband-mgmt VLAN.

Create vNIC Template

Name

:

vNIC_Template_A

Description

:

Fabric ID

:

☒ Fabric A
 ☐ Fabric B
 ☐ Enable Failover

Redundancy

Redundancy Type

:

☒ No Redundancy
 ☐ Primary Template
 ☐ Secondary Template

Target

☒ Adapter
 ☐ VM

Warning

If **VM** is selected, a port profile by the same name will be created.

If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type

:

☐ Initial Template
 ☒ Updating Template

VLANs

Advanced Filter

Export

Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	hx-inband-mgmt	<input checked="" type="radio"/>

OK

Cancel

11. Set hx-inband-mgmt as the native VLAN.
12. For MTU, enter 1500.
13. In the MAC Pool list, select MAC_Pool_A.

14. In the Network Control Policy list, select Veeam_NCP.
15. Select QoS Policy as Veeam_QoS.

Create vNIC Template

☒ **hx-inband-mgmt** ☒

☐ **hx-storage-data** ☐

☐ **hx-vmotion** ☐

☐ **Infra-Mgmt** ☐

☐ **Lab-Network-149** ☐

CDN Source : ☒ vNIC Name ☐ User Defined

MTU :

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

16. Click OK to create the vNIC template.

17. Click OK.

Repeat these similar steps for vNIC_Template_B:

1. In the navigation pane, select the LAN tab.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template

5. Enter vNIC_Template_B as the vNIC template name.
6. Select Fabric B.
7. Do not select the Enable Failover checkbox.
8. Under Target, make sure the VM checkbox is not selected.
9. Select Redundancy Type as No Redundancy
10. Select Updating Template as the template type.
11. Under VLANs, select the checkboxes for VLAN-149, VLAN.

Create vNIC Template

Name

:

vNIC_Template_B

Description

:

Fabric ID

:

☐ Fabric A ☒ Fabric B ☐ Enable Failover

Redundancy

Redundancy Type

:

☒ No Redundancy ☐ Primary Template ☐ Secondary Template

Target

☒ Adapter

☐ VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type

:

☐ Initial Template ☒ Updating Template

VLANs

Advanced Filter ↑ Export Print ⚙		
Select	Name	Native VLAN
<input checked="" type="checkbox"/>	hx-inband-mgmt	<input checked="" type="radio"/>

OK

Cancel

12. Set VLAN-149 as the native VLAN.
13. Select vNIC Name for the CDN Source.

14. For MTU, enter 1500.
15. Select QoS Policy as Veeam_QoS.
16. In the MAC Pool list, select MAC_Pool_B.
17. In the Network Control Policy list, select Veeam_NCP.

Create vNIC Template

<input checked="" type="checkbox"/>	hx-inband-mgmt	<input checked="" type="radio"/>
<input type="checkbox"/>	hx-storage-data	<input type="radio"/>
<input type="checkbox"/>	hx-vmotion	<input type="radio"/>
<input type="checkbox"/>	Infra-Mgmt	<input type="radio"/>
<input type="checkbox"/>	Lab-Network-149	<input type="radio"/>

CDN Source : ☒ vNIC Name ☐ User Defined

MTU : 1500

Warning

Make sure that the MTU has the same value in the QoS System Class corresponding to the Egress priority of the selected QoS Policy.

MAC Pool : MAC_Pool_B(31/32) ▼

QoS Policy : Veeam-QoS ▼

Network Control Policy : Veeam_NCP ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

OK

Cancel

18. Click OK to create the vNIC template.
19. Click OK.

Create Disk Group Policy

A storage profile encapsulates the storage requirements for one or more service profiles. LUNs configured in a storage profile can be used as boot LUNs or data LUNs, and can be dedicated to a specific server. You can also specify a local LUN as a boot device. The introduction of storage profiles allows you to do the following:

- Configure multiple virtual drives and select the physical drives that are used by a virtual drive. You can also configure the storage capacity of a virtual drive.
- Configure the number, type and role of disks in a disk group.
- Associate a storage profile with a service profile

Cisco UCS Manager's Storage Profile and Disk Group Policies are utilized to define storage disks, disk allocation and management in the Cisco UCS S3260 system. We would create two disk Group Policies as

- RAID 1 from two Rear SSDs for OS Boot
- RAID6 from 28 HDD as defined under Disk Zoning Policy

To create Disk Group Policy, complete the following steps:

1. In Cisco UCS Manager, click the Storage tab in the navigation pane.
2. Select Storage Policies > root > Sub-Organizations > Veeam > Disk Group Policies.
3. Right Click on Disk Group Policy and Select Create Disk Group Policy.
4. Enter name as RAID1_OS.
5. Select RAID Level as RAID1 Mirrored.
6. Number of drives as 2 and Drive Type as SSD.
7. Click OK.

Create Disk Group Policy

Name :

Description :

RAID Level :

☒ Disk Group Configuration (Automatic)
 ☐ Disk Group Configuration (Manual)

Disk Group Configuration (Automatic)

Number of drives : [0-60]

Drive Type : ☐ Unspecified ☐ HDD ☒ SSD

Number of Dedicated Hot Spares : [0-60]

Number of Global Hot Spares : [0-60]

Min Drive Size (GB) : [0-10240]

Use Remaining Disks : ☐

Virtual Drive Configuration

Strip Size (KB) :

Access Policy :

8. Create second Disk Group Policy with RAID6.
9. Repeat Step 1 to 3.
10. Enter name as S3260_RAID6.
11. Select RAID6, Number of drives as 26 and Dedicated Hot Spares as 2
12. Select Drive type as HDD.
13. Click OK.

Create Disk Group Policy

Name :

Description :

RAID Level :

☒ Disk Group Configuration (Automatic) ☐ Disk Group Configuration (Manual)

Disk Group Configuration (Automatic)

Number of drives : [0-60]

Drive Type : ☐ Unspecified ☒ HDD ☐ SSD

Number of Dedicated Hot Spares : [0-60]

Number of Global Hot Spares : [0-60]

Min Drive Size (GB) : [0-10240]

Use Remaining Disks : ☐

Virtual Drive Configuration

Strip Size (KB) :

Access Policy :

Read Policy : ☒ Platform Default ☐ Read Ahead ☐ Normal

Write Cache Policy : ☒ Platform Default ☐ Write Through ☐ Write Back Good Bbu ☐ Always Write Back

Create Storage Profile

To create Storage Profile for S3260, complete the following steps:

1. In Cisco UCS Manager, click the Storage tab in the navigation pane.
2. Select Storage Policies > root >Sub-Organizations >Veeam.
3. Right Click and Select Create Storage Profile.
4. Enter name as S3260_Str_Prfl_1.
5. Under Local Lun Selection, click Add.

Create Storage Profile

?

×

Name : S3260_Str_Prfl_1

Description :

LUNs

Local LUNs

Controller Definitions

Advanced Filter

Export

Print

Name	Size (GB)	Order	Fractional Size (MB)
No data available			

+

 Add Delete

i

 Info

OK

Cancel

6. Enter Name as OS_Boot.

7. Check Expand to Available; this creates a single lun with maximum space available.

8. Select Disk Group Selection as 'RAID1_OS' and click OK.

Create Local LUN

☒ Create Local LUN ☐ Prepare Claim Local LUN

Name : OS_Boot

Size (GB) : 1 **[0-102400]**

Fractional Size (MB) : 0

Auto Deploy : ☒ Auto Deploy ☐ No Auto Deploy

Expand To Available : ☒

Select Disk Group Configuration : RAID1_OS [Create Disk Group Policy](#)

OK

9. Click Add under Local LUN.
10. Enter Name as Veeam_Rep; this is the LUN used by Veeam Repository.
11. Check Expand to Available and Select Disk Group Configuration as 'S3260_RAID6'.
12. Click OK.

Create Local LUN

☒ Create Local LUN

☐ Prepare Claim Local LUN

Name

:

Veeam_Rep

Size (GB)

:

1

[0-102400]

Fractional Size (MB)

:

0

Auto Deploy

:

☒ Auto Deploy

☐ No Auto Deploy

Expand To Available

:

☒

Select Disk Group Configuration :

S3260_RAID6

▼

Create Disk Group Policy

OK

-
13. Click OK on the Create Storage Profile.

?

×

Create Storage Profile

Name : S3260_Str_Prfl_1

Description :

LUNs

Local LUNs

Controller Definitions

Advanced Filter

Export

Print

⚙

Name	Size (GB)	Order	Fractional Size (MB)
Veeam_Rep	1	Not Applicable	0
OS_Boot	1	Not Applicable	0

+

Add

🗑

Delete

ℹ

Info

OK

Cancel

Create Chassis Profile Template

A chassis profile defines the storage, firmware and maintenance characteristics of a chassis. You can create a chassis profile for the Cisco UCS S3260 Storage Server. When a chassis profile is associated to a chassis, Cisco UCS Central automatically configures the chassis to match the configuration specified in the chassis profile.

A chassis profile includes four types of information:

- Chassis definition—Defines the specific chassis to which the profile is assigned.
- Maintenance policy—Includes the maintenance policy to be applied to the profile.
- Firmware specifications—Defines the chassis firmware package that can be applied to a chassis through this profile.
- Disk zoning policy—Includes the zoning policy to be applied to the storage disks.

To create Chassis Profile Template for Cisco UCS S3260 storage server, complete the following steps:

1. In Cisco UCS Manager, click the Chassis tab in the navigation pane.
2. Select Chassis Profile Templates > root > Sub-Organizations > Veeam.

3. Right-click and select Create Chassis Profile Template.
4. Enter name as Chassis_Template.
5. Select Type as Updating Template.

Create Chassis Profile Template

You must enter a name for the chassis profile template and specify the template type. You can also enter a description of the template.

Name :

The template will be created in the following organization. Its name must be unique within this organization.

Where : **org-root/org-Veeam**

Type : ☐ Initial Template ☒ Updating Template

Optionally enter a description for the template. The description can contain information about when and where the chassis profile template should be used.

< Prev Next > **Finish** Cancel

6. Select default as the Maintenance Policy and click Next.
7. Select Chassis Firmware Package as 'S3260_FW_Package'.
8. Select Disk Zoning Policy as 'S3260_DiskZone' and click Finish.

Create Service Profile Template

With a service profile template, you can quickly create several service profiles with the same basic parameters, such as the number of vNICs and vHBAs, and with identity information drawn from the same pools.



If you need only one service profile with similar values to an existing service profile, you can clone a service profile in the Cisco UCS Manager GUI.

For example, if you need several service profiles with similar values to configure servers to host database software, you can create a service profile template, either manually or from an existing service profile. You then use the template to create the service profiles.

Cisco UCS supports the following types of service profile templates:

- **Initial template:** Service profiles created from an initial template inherit all the properties of the template. However, after you create the profile, it is no longer connected to the template. If you need to make changes to one or more profiles created from this template, you must change each profile individually.

- **Updating template:** Service profiles created from an updating template inherit all the properties of the template and remain connected to the template. Any changes to the template automatically update the service profiles created from the template.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organizations > Veeam.
3. Right-click Veeam.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter S3260_SP_Template as the name of the service profile template.
6. Select the “Updating Template” option.
7. Under UUID, select UUID_Pool as the UUID pool.

Create Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to the template and enter a description.

Name :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-Veeam**

The template will be created in the following organization. Its name must be unique within this organization.
Type : ☐ Initial Template ☒ Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.
UUID

UUID Assignment:

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

8. Click Next.

Configure Storage Provisioning

1. Click Storage Profile Policy Tab and select S3260_Str_Prfl_1 (as created under Storage Profile section).

1

Identify Service Profile Template

2

Storage Provisioning

3

Networking

4

SAN Connectivity

5

Zoning

6

vNIC/vHBA Placement

7

vMedia Policy

8

Server Boot Order

9

Maintenance Policy

10

Server Assignment

11

Operational Policies

Create Service Profile Template

Optionally specify or create a Storage Profile, and select a local disk configuration policy.

Specific Storage Profile

Storage Profile Policy

Local Disk Configuration Policy

Storage Profile: S3260_Str_Prfl_1

Create Storage Profile

Name : S3260_Str_Prfl_1

Description :

LUNs

Local LUNs

Controller Definitions

Advanced Filter

Export

Print

Name	Size (GB)	Order	Fractional Size (MB)
OS_Boot	1	Not Applicable	0
Veeam_Rep	1	Not Applicable	0

< Prev

Next >

Finish

- Click Next.

Configure Networking Options

- Keep the default setting for Dynamic vNIC Connection Policy.
- Under ‘How would you like to configure LAN connectivity’, select Expert Mode.
- Select the “Use Connectivity Policy” option to configure the LAN connectivity.

1

Identify Service Profile Template

2

Storage Provisioning

3

Networking

4

SAN Connectivity

5

Zoning

6

vNIC/vHBA Placement

7

vMedia Policy

8

Server Boot Order

9

Maintenance Policy

10

Server Assignment

11

Operational Policies

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

Select a Policy to use (no Dynamic vNIC Policy by default) ▼

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

☐ Simple

☒ Expert

☐ No vNICs

☐ Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Nat
No data available			

Delete

+ Add

Modify

< Prev

Next >

4. Click Add.
5. Under Create vNIC option, enter name as eth0.
6. Select use vNIC Template and choose vNIC_Template_A.
7. Under Adaptor Policy Select 'veeam_adaptorpol' and click OK.

Create vNIC

Name :

eth0

Use vNIC Template :

☒

Redundancy Pair :

☐

Peer Name :

vNIC Template :

vNIC_Template_A ▼

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

veeam_adaptorpol ▼

Create Ethernet Adapter Policy

8. Repeat Step 1 to 7 and name the vNIC as eth1 and vNIC Template as vNIC_Template_B.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

Select a Policy to use (no Dynamic vNIC Policy by default) ▼

Create Dynamic vNIC Connection Policy

How would you like to configure LAN connectivity?

☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Na
vNIC eth1	Derived	derived	
vNIC eth0	Derived	derived	

Delete

Add

Modify

< Prev

Next >

9. Click Next.

96

Configure Storage Options

Skip the SAN Connectivity since you will use local storage for S3260 created through Storage Policy and Select No vHBAs

Configure Zoning Options

1. Set no Zoning options and click Next.

Configure vNIC/HBA Placement

1. In the "Select Placement" list, leave the placement policy as "Let System Perform Placement".
2. Click Next.

Configure vMedia Policy

1. From the vMedia Policy drop-down, select "ESXi-6.0U1b-HTTP".
2. Click Next.

Configure Server Boot Order

1. Choose Default Boot Policy.

Configure Maintenance Policy

1. Change the Maintenance Policy to userAck.

Create Service Profile Template

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to service profile.

Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be access

Maintenance Policy: **UserAck** [Create Maintenance Policy](#)

Name : **UserAck**
 Description :
 Soft Shutdown Timer : **150 Secs**
 Reboot Policy : **User Ack**

[< Prev](#) [Next >](#)

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select Assign Later.
2. Firmware Management at the bottom of the page select S3260Firmware as created in the previous section.

Create Service Profile Template

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: Assign Later ▼ [Create Server Pool](#)

Select the power state to be applied when this profile is associated with the server.

☒ Up ☐ Down

The service profile template is not automatically associated with a server. Either select a server from the list or associate manually later.

Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: S3260Firmware ▼ [Create Host Firmware Package](#)

[< Prev](#) [Next >](#)

3. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

1. In the BIOS Policy list, select S3260-BIOS.
2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

Create Service Profile Template

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy : S3260-BIOS ▼

External IPMI Management Configuration

Management IP Address

Monitoring Configuration (Thresholds)

Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : No-Power-Cap ▼ [Create Power Control Policy](#)

[< Prev](#) [Next >](#) [Finish](#)

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create Chassis Profile

To create chassis profile from the chassis profile template, complete the following steps:

1. Click the Chassis tab in the navigation pane.
2. Select Chassis Profile Templates > root > Sub-Organizations > Veeam > Chassis Profile Template Chassis_Template.
3. Right-click Chassis Profile Template Chassis_Template and Select Create Chassis Profiles from Template
4. Enter S3260_Chassis_SP as the Chassis profile prefix.
5. Enter 1 as "Name Suffix Starting Number" and 1 as Number of Instances.

Create Chassis Profiles From Template



Naming Prefix : S3260_ChassisSP

Name Suffix Starting Number : 1

Number of Instances : 1

OK Cancel

6. The screenshot below displays S3260_ChassisSP1 under Chassis > root > Sub_organizations > Veeam > Chassis Profile.

Chassis / Chassis Profi... / root / Sub-Organizations / Veeam / Chassis Pro...

General Policies Chassis FSM Faults Events

Fault Summary

0

0

0

1

Status

Overall Status :

Unassociated

Status Details

Actions

Rename Chassis Profile

Create a Clone

Create a Chassis Profile Template

Disassociate Chassis Profile

Change Chassis Profile Association

Properties

WARNING

This chassis profile is not modifiable because it is bound to the chassis profile template **Chassis_Template**. To modify this chassis profile, please unbind it from the template.

Name : S3260_ChassisSP1

User Label :

Description :

Owner : Local

Associated Chassis :

Chassis Profile Template : Chassis_Template

Template Instance : org-root/org-Veeam/cp-Chassis_Template

Assigned Chassis

Chassis Maintenance Policy

Associate Chassis Profile to S3260 Chassis

To Associate Chassis Profile to S3260 Chassis, complete the following steps:

1. Click the Chassis tab in the navigation pane.
2. Select Chassis Profiles > root > Sub-Organizations > Veeam.
3. Right-click 'S3260_Chassis_SP1' and select Change Chassis Profile Association.
4. In the Assignment tab, select Existing Chassis.
5. Thereafter select the existing chassis.

Associate Chassis Profile ? ×

Select a previously-discovered chassis by name, or manually specify a custom chassis by entering its chassis ID. If no chassis currently exists at that location, the system waits until one is discovered.

You can select an existing chassis you want to associate with this chassis profile.

Chassis Assignment: Select existing Chassis ▼

☒ Available Chassis
 ☐ All Chassis

Select	ID
<input checked="" type="radio"/>	1

Restrict Migration : ☐

OK
Cancel

6. Click OK.
7. Since you have selected User Ack for the Maintenance Policy, you need to acknowledge Chassis Reboot for Chassis Profile Association.
8. On the FSM tab, you will see the Association Status.

Chassis / Chassis Profi... / root / Sub-Organizations / Veeam / Chassis Pro...

GeneralPoliciesChassisFSMFaultsEvents

ChassisChassis Profile

FSM Status : In Progress

Description :

Current FSM Name : Associate

Completed at :

Progress Status :

86%

Remote Invocation Result :

Remote Invocation Error Code : None

Remote Invocation Description :

Step Sequence

Order	Name	Description	Status	Timestamp	Retried
1	Associate Download...	Download images(F...	Skip	2017-01-10T19:33:27	0
2	Associate Copy Re...	Copy images to peer...	Skip	2017-01-10T19:33:27	0
3	Associate Wait Befo...	Wait before installati...	Success	2017-01-10T19:33:27	0
4	Associate Update C...	Update CMC fw(FS...	Success	2017-01-10T19:33:27	0

9. When the Chassis is Associated you will see the assigned status as Assigned.

Equipment

Servers

LAN

SAN

VM

Storage

Chassis

Admin

All

Chassis

Chassis Profiles

root

Sub-Organizations

hx-cluster

Sub-Organizations

Veeam

S3260_ChassisSP1

Sub-Organizations

Chassis Profile Templates

root

Sub-Organizations

hx-cluster

Veeam

Chassis Profile Template Chas

Sub-Organizations

Policies

root

Chassis Firmware Packages

Chassis Maintenance Policies

Disk Zoning Policies

Chassis / Chassis Profi... / root / Sub-Organizations / Veeam / Chassis Pro...

GeneralPoliciesChassisFSMFaultsEvents

0000

Status

Overall Status : OK

Status Details

Assoc State : Associated

Assigned State : Assigned

Actions

Rename Chassis Profile

Create a Clone

Create a Chassis Profile Template

Disassociate Chassis Profile

Change Chassis Profile Association

Unbind from the Template

Bind to a Template

Change Chassis Maintenance Policy

WARNING

This chassis profile is not modifiable because it is bound to the chassis profile template **Chassis_Template**. To modify this chassis profile, please unbind it from the template.

Name : S3260_ChassisSP1

User Label :

Description :

Owner : Local

Associated Chassis : sys/chassis-1

Chassis Profile Template : Chassis_Template

Template Instance : org-root/org-Veeam/cp-Chassis_Template

Assigned Chassis

Chassis : sys/chassis-1

Restrict Migration : No

Chassis Maintenance Policy

Name : default

Maintenance Policy Instance : org-root/chassis-profile-maint-default

Description :

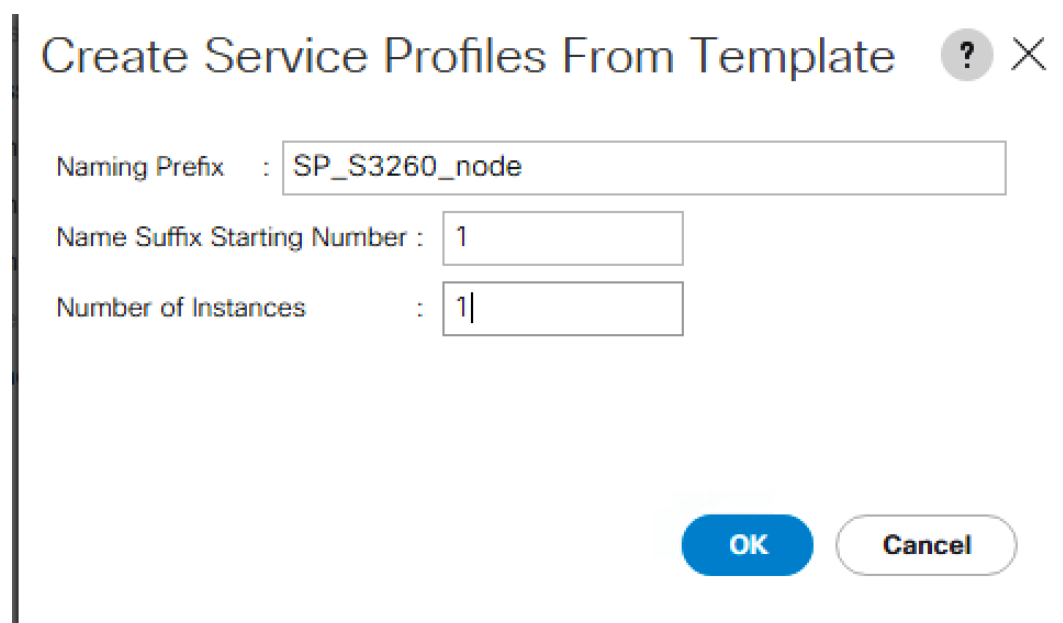
Reboot Policy : User Ack

Create Service Profiles

This section describes how to associate the Compute Node on S3260 Storage server to a Service Profile.

To create service profiles from the service profile template, complete the following steps:

1. On Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Sub-Organizations > Veeam > Service Template S3260_SP_Template.
3. Right-click S3260_SP_Template and select Create Service Profiles from Template.
4. Enter SP_S3260_node as the service profile prefix.
5. Enter 1 as “Name Suffix Starting Number.”
6. Enter 1 as the “Number of Instances.”
7. Click OK to create the service profile.



Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

OK Cancel

8. Click OK in the confirmation message.

Associate Service Profile to Server node of S3260 Chassis

To Associate Service Profile to server node of S3260 Chassis, complete the following steps:

1. Click the Chassis tab in the navigation pane.
2. Select Service Profiles > root > Sub-Organizations > Veeam.
3. Right-click 'S3260_Chassis_SP1' and select Change Service Profile Association.
4. In the Assignment tab, select Existing Chassis.
5. In the Associate Service Profile Window, select the M4 node for S3260 Chassis.
6. Click OK.

Associate Service Profile

?

×

Select an existing server pool or a previously-discovered server by name, or manually specify a custom server by entering its chassis and slot ID. If no server currently exists at that location, the system waits until one is discovered.

You can select an existing server or server pool, or specify the physical location of the server you want to associate with this service profile.

Server Assignment:

Select existing Server ▾

☒ Available Servers

☐ All Servers

Select	Chassis ID	Slot	Rack ID	PID	Procs	Memory	Adapters
<input checked="" type="radio"/>	1	1		UCSC-C3...	2	262144	1
<input type="radio"/>			11	UCSC-C2...	2	393216	1


Restrict Migration : ☐

OK

Cancel

7. A warning displays, click Yes.

Associate Service Profile

 Your changes:
Create: Server sys/chassis-1/blade-1 (org-root/org-Veeam/ls-SP_S3260_node1/pn)

Will cause the Immediate Reboot of:
Service Profile SP_S3260_node1 (org-root/org-Veeam/ls-SP_S3260_node1) [Server: sys/chassis-1/blade-1]

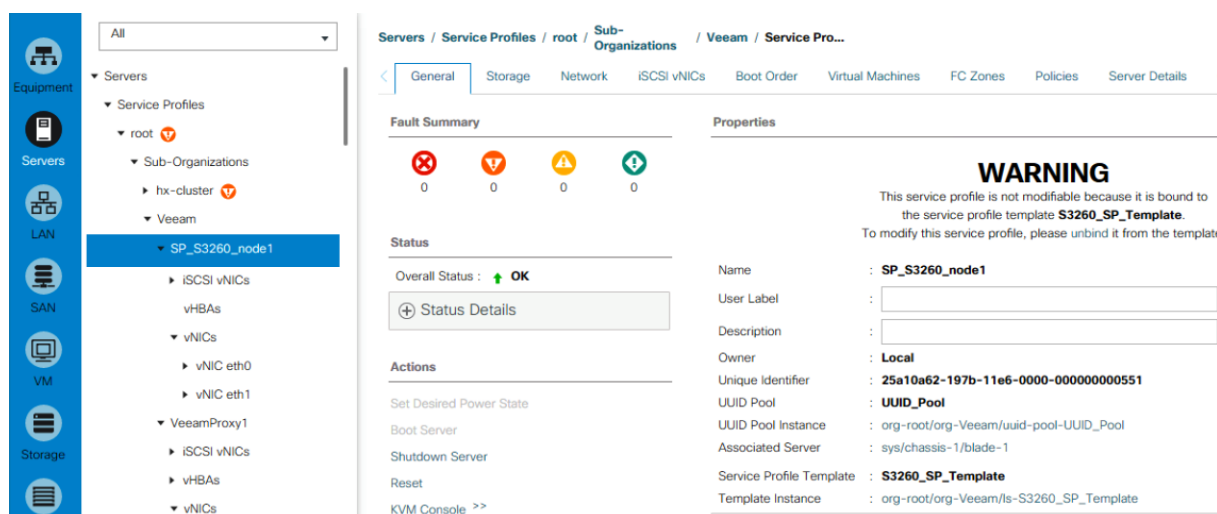
LUN Resource Selection Logs for Service Profile SP_S3260_node1, LUN OS_Boot:

Order	Description
1	Disk selection process started for local lun: org-root/org-Veeam/profile-S3260_Str_Prf_1/das-scsi-lun-OS_Boot
2	Try to find out an existing disk group for the new LUN
3	Cannot carve out of the existing disk groups. Trying to create a new disk group
4	Controller sys/chassis-1/blade-1/board/storage-PCH-1 does not support OOB
5	Select normal disk in slot: 201
6	Select normal disk in slot: 202

LUN Resource Selection Logs for Service Profile SP_S3260_node1, LUN Veeam-Rep:

Order	Description
1	Disk selection process started for local lun: org-root/org-Veeam/profile-S3260_Str_Prf_1/das-scsi-lun-Veeam-Rep
2	Try to find out an existing disk group for the new LUN
3	Cannot carve out of the existing disk groups. Trying to create a new disk group
4	Controller sys/chassis-1/blade-1/board/storage-PCH-1 does not support OOB
5	Select normal disk in slot: 1
6	Select normal disk in slot: 2
7	Select normal disk in slot: 3
8	Select normal disk in slot: 4
9	Select normal disk in slot: 5
10	Select normal disk in slot: 6

8. When Service Profile Association is complete, confirm that the overall status is OK.



9. Verify the Boot Lun and Veeam Repository LUN under Storage tab of Service Profile.

All

Equipment

Servers

LAN

SAN

VM

Storage

Chassis

Servers

Service Profiles

root

Sub-Organizations

hx-cluster

Veeam

SP_S3260_node1

iSCSI vNICs

vHBAs

vNICs

vNIC eth0

vNIC eth1

VeeamProxy1

iSCSI vNICs

vHBAs

vNICs

vNIC eth0

Sub-Organizations

Servers / Service Profiles / root / Sub-Organizations / Veeam / Service Pro...

General

Storage

Network

iSCSI vNICs

Boot Order

Virtual Machines

FC Zones

Policies

Server Details

CIMC

Storage Profiles

Local Disk Configuration Policy

vHBAs

vHBA Initiator Groups

Actions

Storage Profile Policy

Modify Storage Profile

Name : S3260_Str_Prfl_1

Description :

Storage Profile Instance : org-root/org-Veeam/profile-S3260_Str_Prfl_1

Local LUNs

Controller Definitions

Faults

Advanced Filter

Export

Print

Name	RAID Level	Size (MB)	Config State	Deploy Name	LUN ID	Drive State
OS_Boot	RAID 1 Mirrored	456809	Applying	OS_Boot-1	1000	unknown
Veeam-Rep	RAID 60 Striped ...	133514052	Applying	Veeam-Rep-1	1001	unknown

Add

Delete

Info

Veeam Availability Suite 9.5 Installation

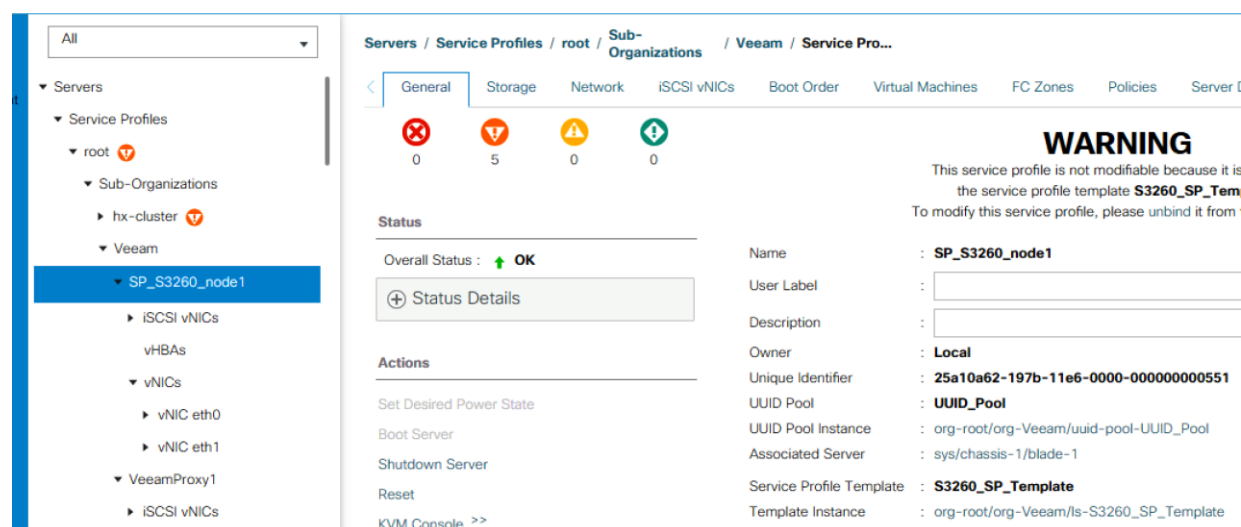
This section elaborates on the installation and configuration of Veeam Availability Suite 9.5 on Cisco UCS S3260 storage server node. The important high-level steps are as follows:

- Install and configure Windows 2016 on S3260 server node
- Install Veeam Availability Suite 9.5
- Configure Backup Infrastructure for Veeam Availability Suite

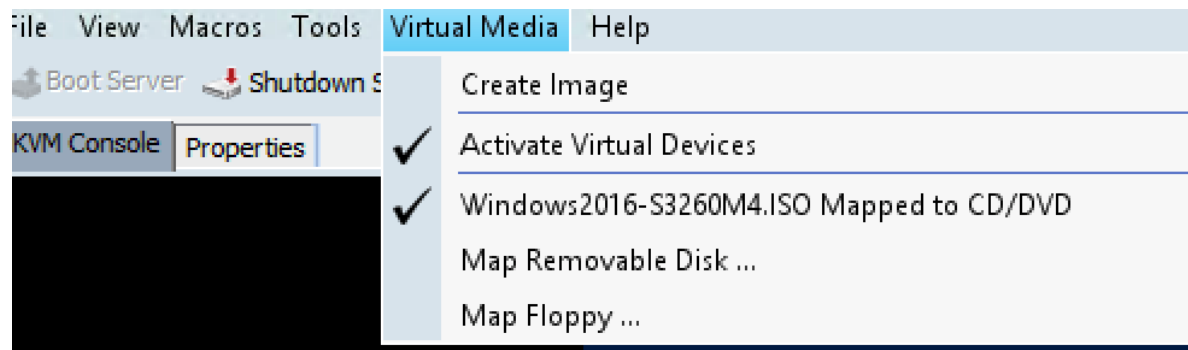
Install and Configure Windows 2016

To install and configure Windows 2016, complete the following steps:


1. In the Navigation pane, select Server tab.
2. In the Servers tab, expand Service Profiles > root > Sub-Organizations > Veeam > SP_S3260_node1.
3. Click KVM console and open the *.jnlp with java webstart.



4. In KVM Console, go to the Virtual Media tab and select Activate Virtual Devices.
5. On the Virtual Media tab, select MAP CD/DVD and browse to Windows 2016 Installer and Map Device.



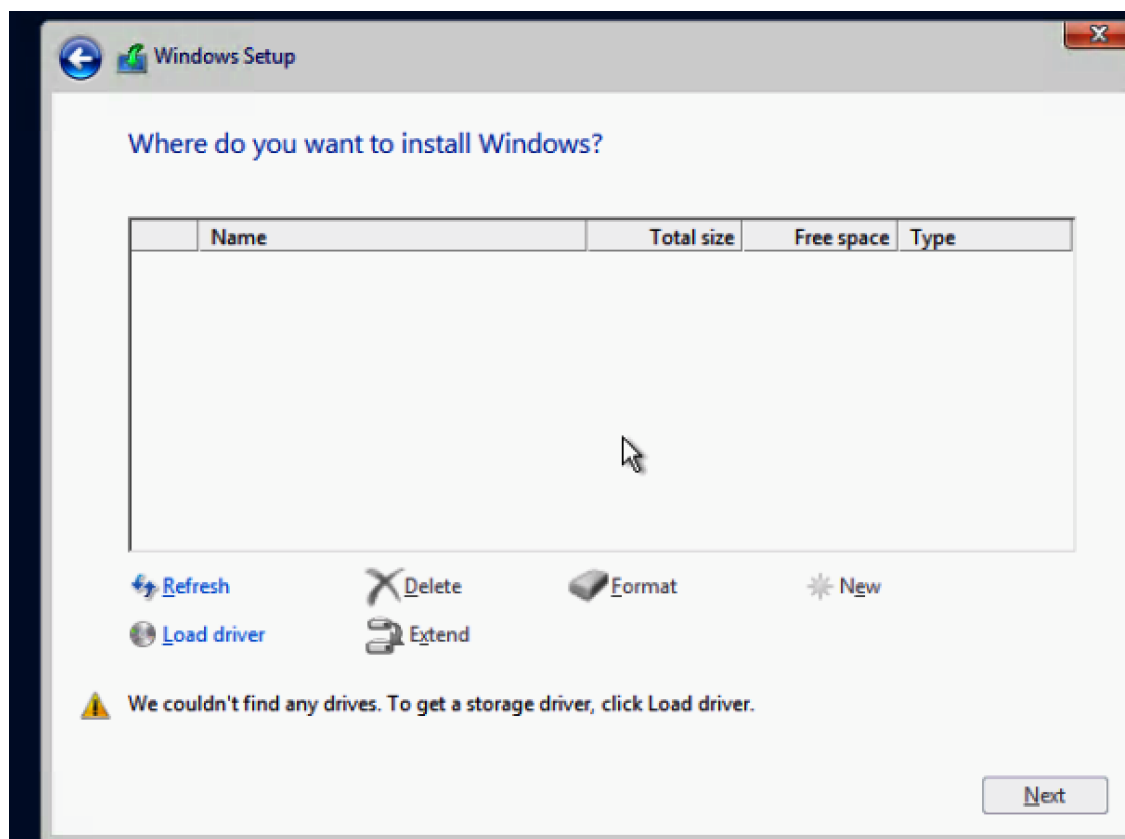
6. Reset the server and wait for the ISO image to load.
7. Install Windows 2016.



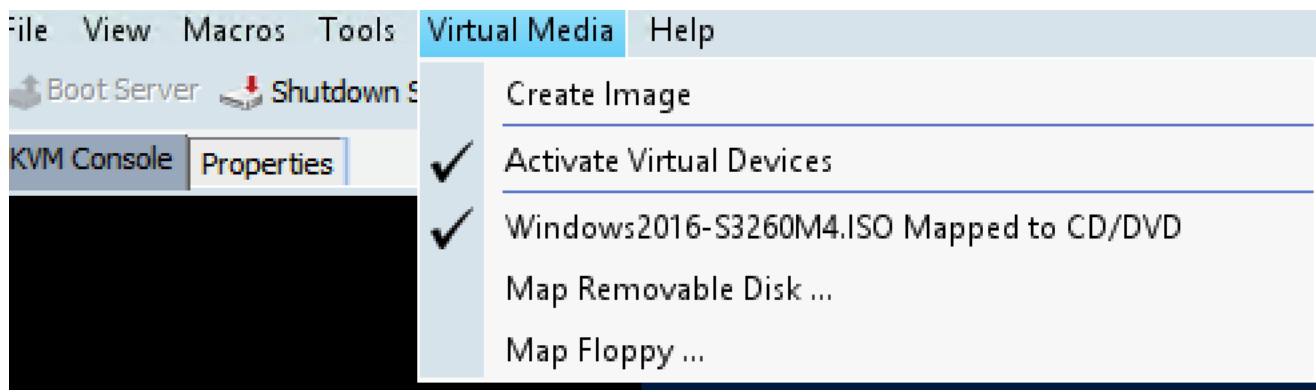
Load Driver for S3260 RAID Controller

To load the RAID Controller driver S3260, complete the following steps:


1. On the Screen 'Where do you want to install Windows?' click Load Driver.



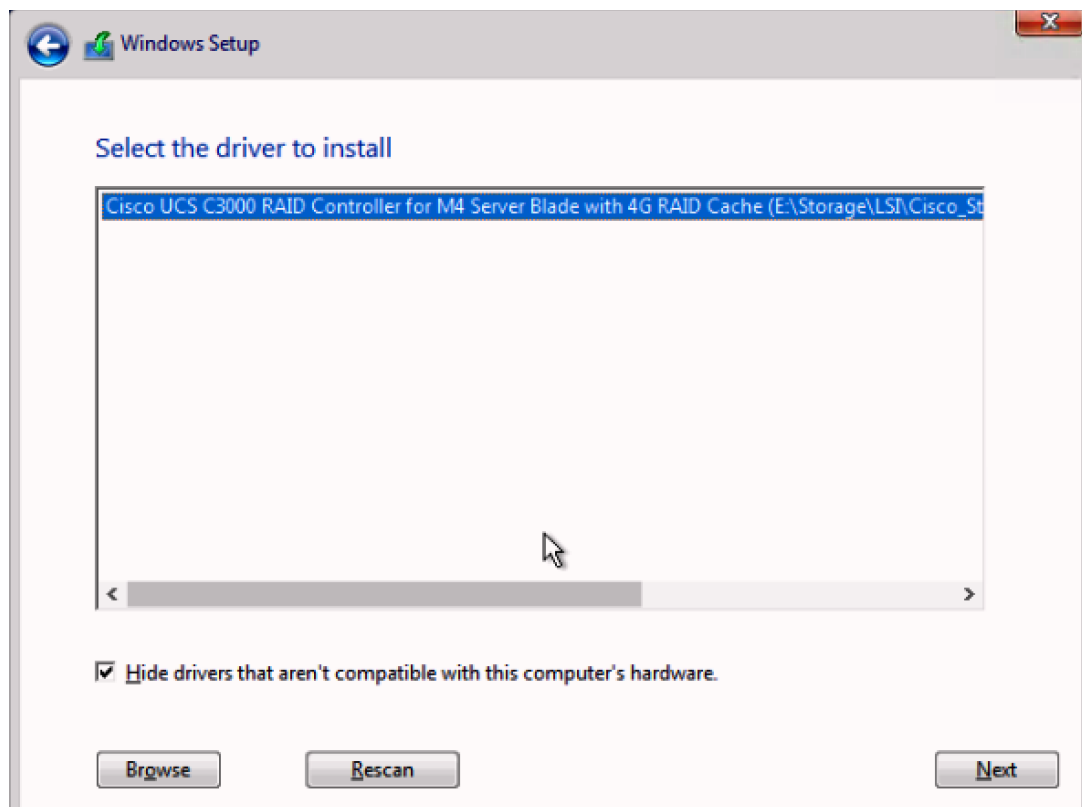
2. In Cisco UCS Manager, click the LAN tab in the navigation pane.
3. In Virtual Media, un-map the Windows installer ISO and map the S3260 drivers ISO. The drivers for S3260 can be downloaded from www.cisco.com at location > Downloads Home Products Servers > Unified Computing > UCS C-Series Rack-Mount UCS-Managed Server Software > Unified Computing System (UCS) Drivers-3.0(1a).




4. Click Browse and navigate to the location of the driver as shown in image below:



5. Click Rescan and view the correct RAID Controller driver in the "Select the driver to install" window.



6. Click Next to install the driver.
7. Return to 'Where do you want to install Windows' screen, uncheck the driver ISO image and re-map the Windows Installer Image.
8. Click Refresh.
9. Select the Drive2. This drive is RAID1 config created from the two SSD in the rear of S3260 chassis for OS installation through Storage Profile in the Cisco UCS Service Profile. Drive3 is the RAID6 configuration created from the top load SAS drives for Veeam Repository.
10. Click Next.




11. When the installation completes, proceed to the following section, Update Cisco VIC Driver for Windows 2016.

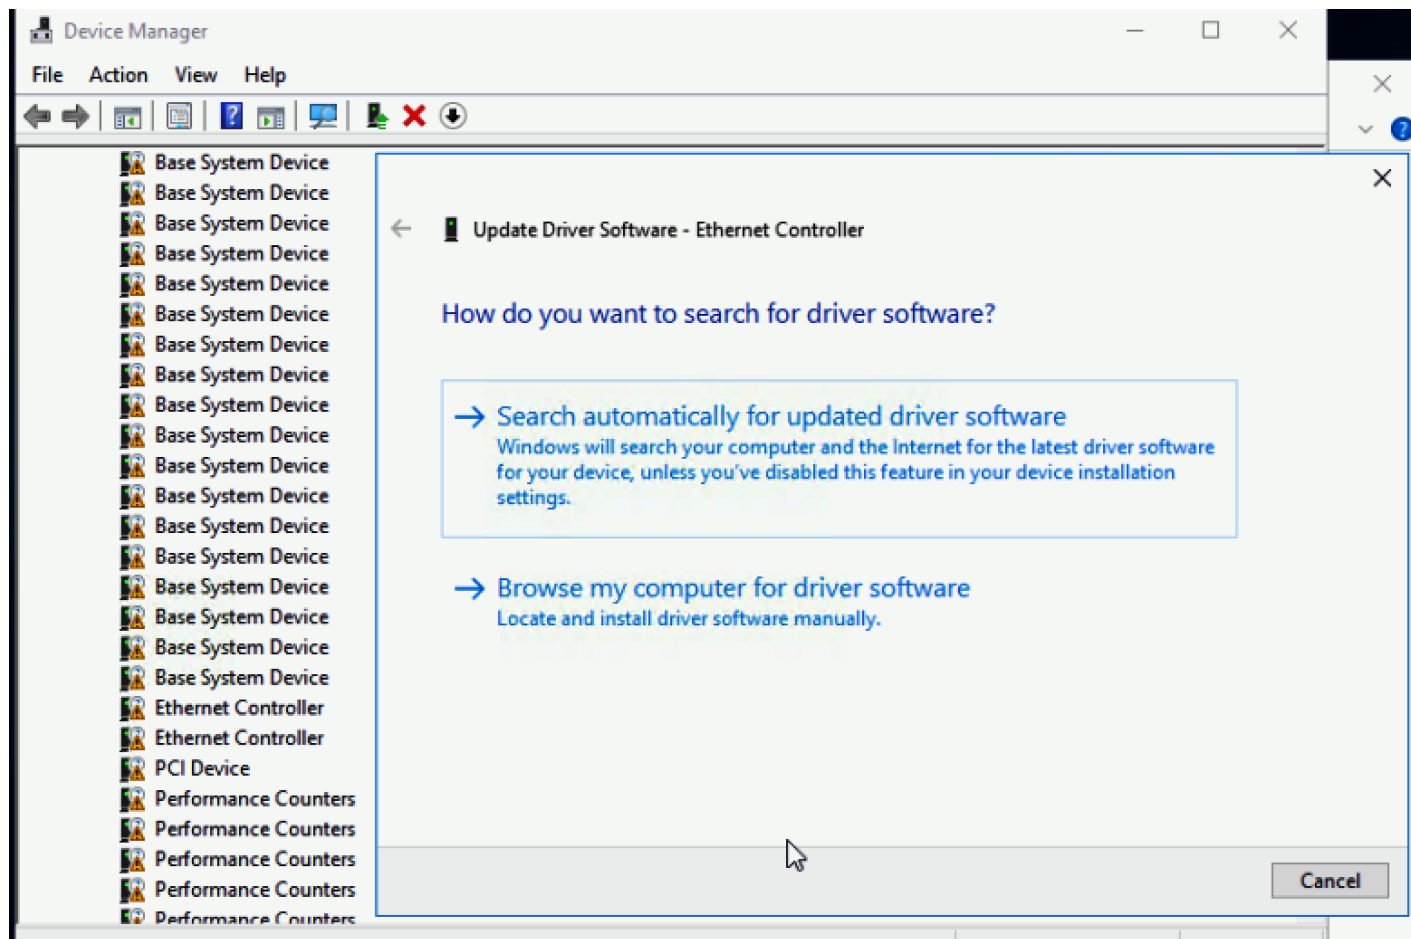
Update Cisco VIC Driver for Windows 2016

To update the Cisco VIC driver for Windows 2016, complete the following steps:

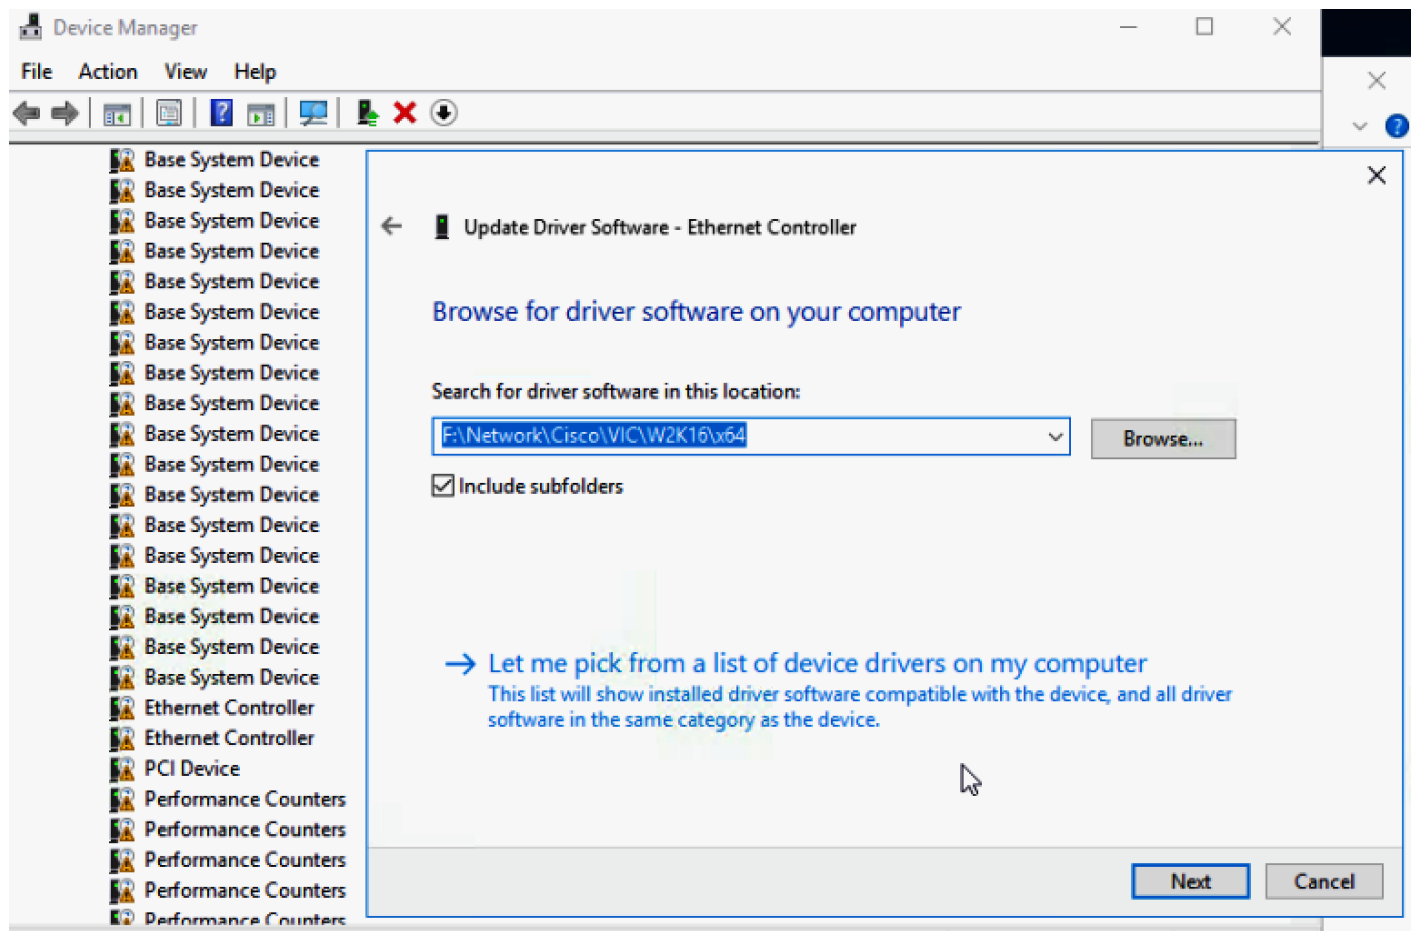
1. Open the UCS KVM Console and login to Windows 2016 installed on S3260 Storage Server.
2. Map the S3260 drivers through Map CD/DVD option under the Virtual Media tab in the KVM Console.
3. The S3260 drivers are located in the section 'Load Drivers for S3260 RAID Controller.'



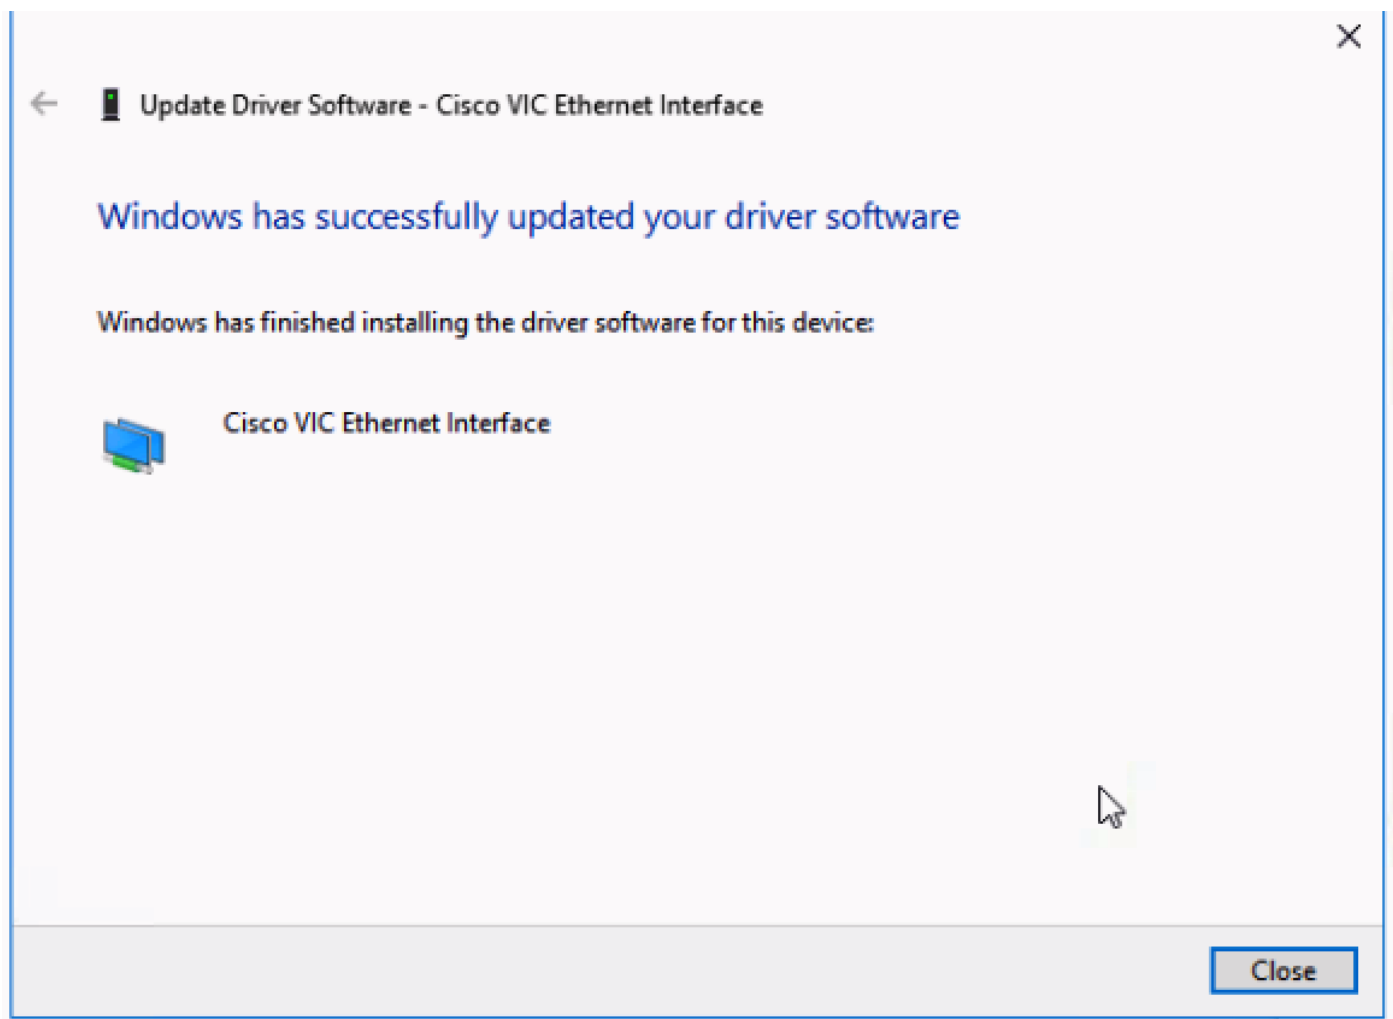
4. In Windows 2016, go to Control Panel > Device Manager.
5. Select Ethernet Controller and Select Update Driver.



6. Select for 'Browse for Driver' in My Computer.'
7. Select DVD Driver as mapped through Virtual Media in KVM Console and browse to \\Network\Cisco\\VIC\\W2K16\\x64.



8. Click Next to install the Cisco VIC Ethernet Interface driver.




9. Follow steps 5 to 8 to install the driver for the second VIC Ethernet interface.


Update Intel ChipSet Driver for Windows 2016

To update Intel ChipSet driver for Windows 2016, complete the following steps:

1. Update Driver for S3260 Intel Chip Set.
2. Under the S3260 Driver ISO mounted through Virtual Media of KVM Console, browse to \ChipSet\Intel\C3260\W2K16.



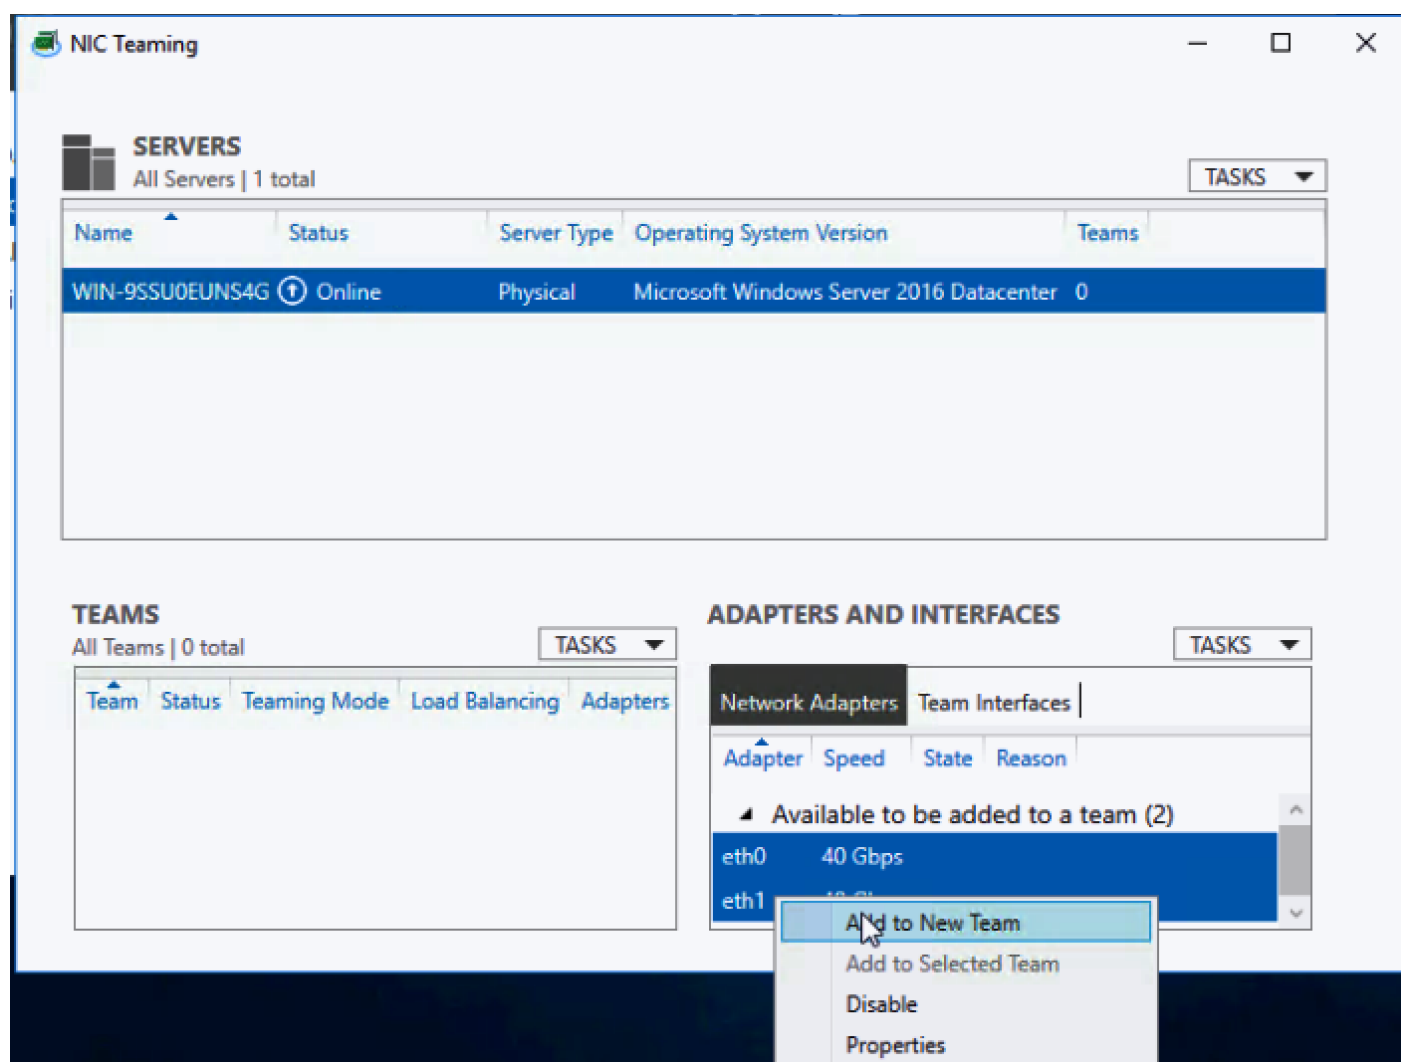
3. Execute SetupChipset.exe. When it is installed, restart the system.




Create NIC Teaming for Windows 2016

To create NIC Teaming for Windows 2016, complete the following steps:

1. Go to Server Manager > Local Server and click NIC Teaming option.
2. Right-click the two 40Gbps VIC interfaces and click Add to New Team.




- Enter NIC Team name as vic80gbps.
- Under Additional properties, select NIC teaming mode as 'Switch Independent.'
- Click OK. You will see that both the Ethernet interfaces are Green.



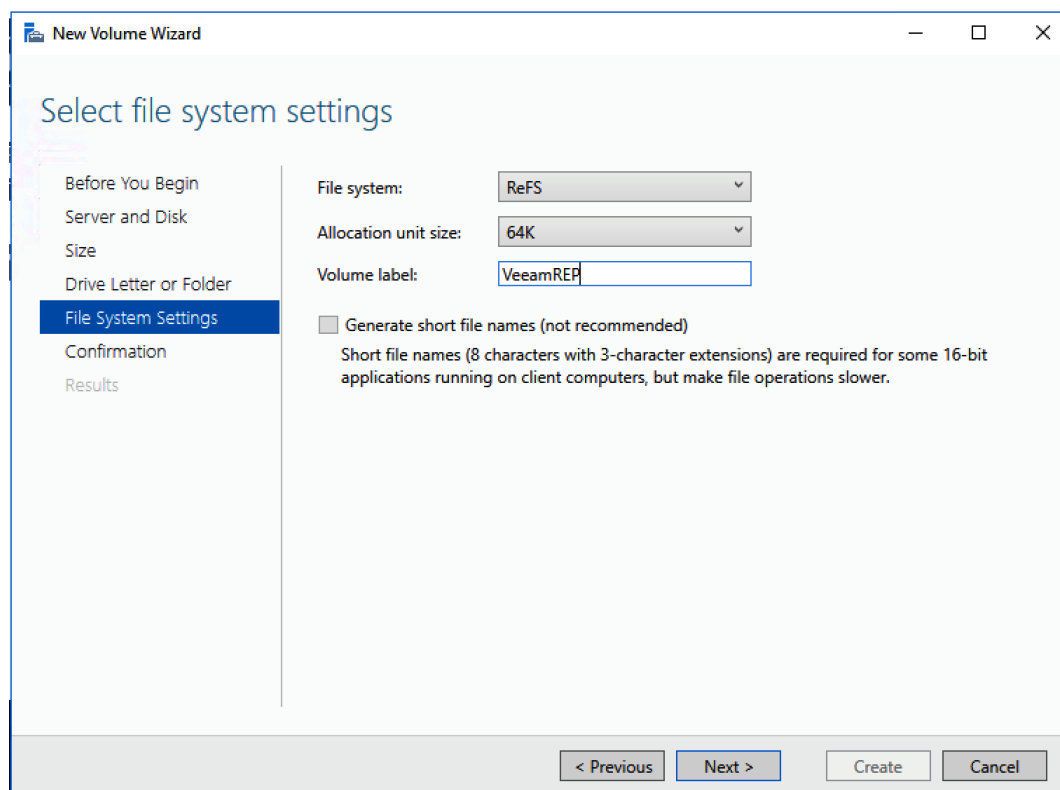
Create Disk Volume for Veeam Repository

To create disk volume for Veeam repository, complete the following steps:

1. Go to Server Manager > File and Storage Services.
2. Navigate to Volumes > Disks and select the volume with Partition type as 'Unknown.'
3. Create a New Volume.



4. Click Next until you reach 'Select File System settings' window.
5. Create a Volume Label, select File system to 'ReFS', Allocation unit size to '64k', Volume label to 'VeeamRep.'
6. Click Next.



7. Confirm the File System Settings and click Create.



ReFS volumes provide significantly faster synthetic full backup creation and transformation performance, as well as reduce storage requirements and improve reliability. Even more importantly, this functionality improves Availability of backup storage by significantly reducing its load — which results in improved backup and restore performance and enables customers to do much more with virtual labs running off of backup storage.

Install Veeam Availability Suite 9.5

To install Veeam Availability Suite 9.5, complete the following steps:



For detailed steps to install and configure Veeam 9.5, refer to the [Veeam Cisco UCS S3260 Configuration Guide](#).

- 1. Download the Veeam software from <https://www.veeam.com/data-center-availability-suite-vcp-download.html>. Download a free 30-day trial license key or obtain license key from Veeam.
- 2. Execute Veeam 9.5 setup file.

★ Quick access

Desktop

Downloads

Documents

Pictures

This PC

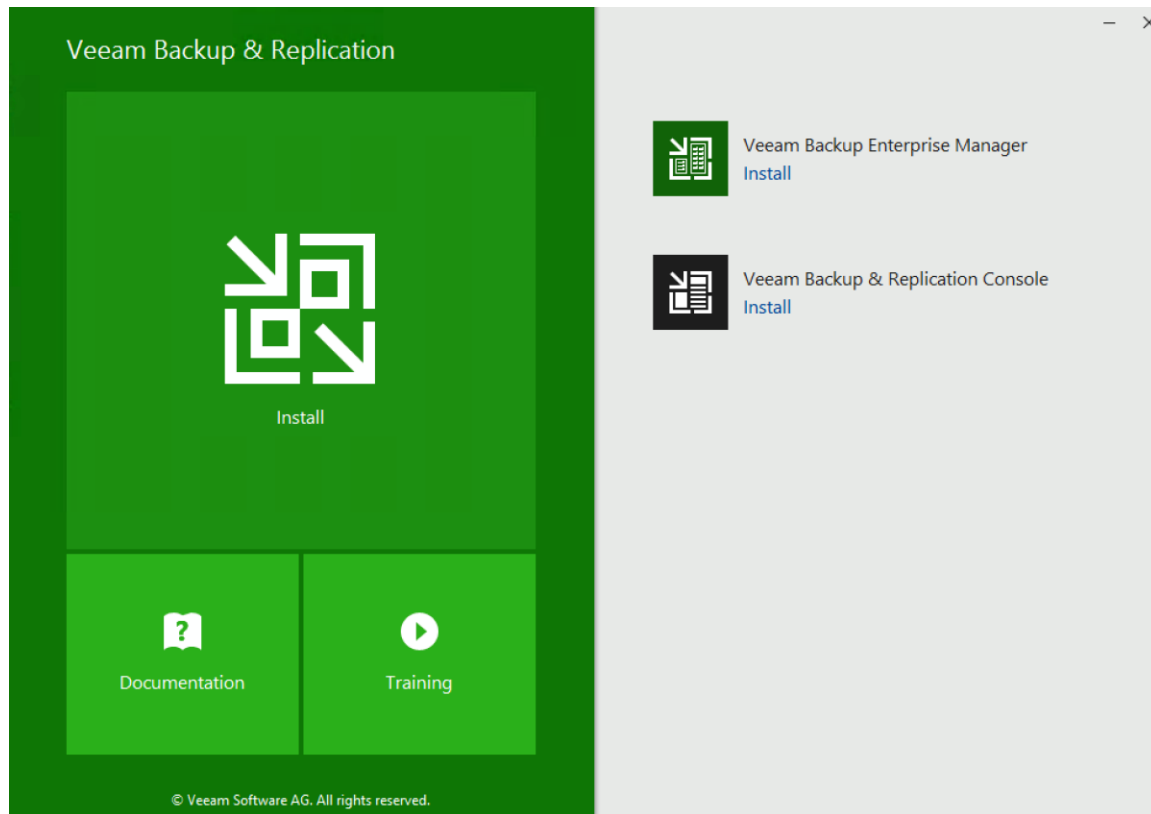
Network

Name	Date modified	Type
AIR	10/27/2016 9:03 PM	File folder
Backup	10/27/2016 9:03 PM	File folder
Catalog	10/27/2016 9:03 PM	File folder
Cloud Portal	1/12/2017 3:12 PM	File folder
EnterpriseManager	10/27/2016 9:03 PM	File folder
Explorers	10/27/2016 9:04 PM	File folder
Redistr	10/27/2016 9:04 PM	File folder
Search	10/27/2016 9:04 PM	File folder
Suite	10/27/2016 9:04 PM	File folder
autorun	10/27/2016 9:05 PM	Setup Info
Setup	10/27/2016 9:05 PM	Applicati

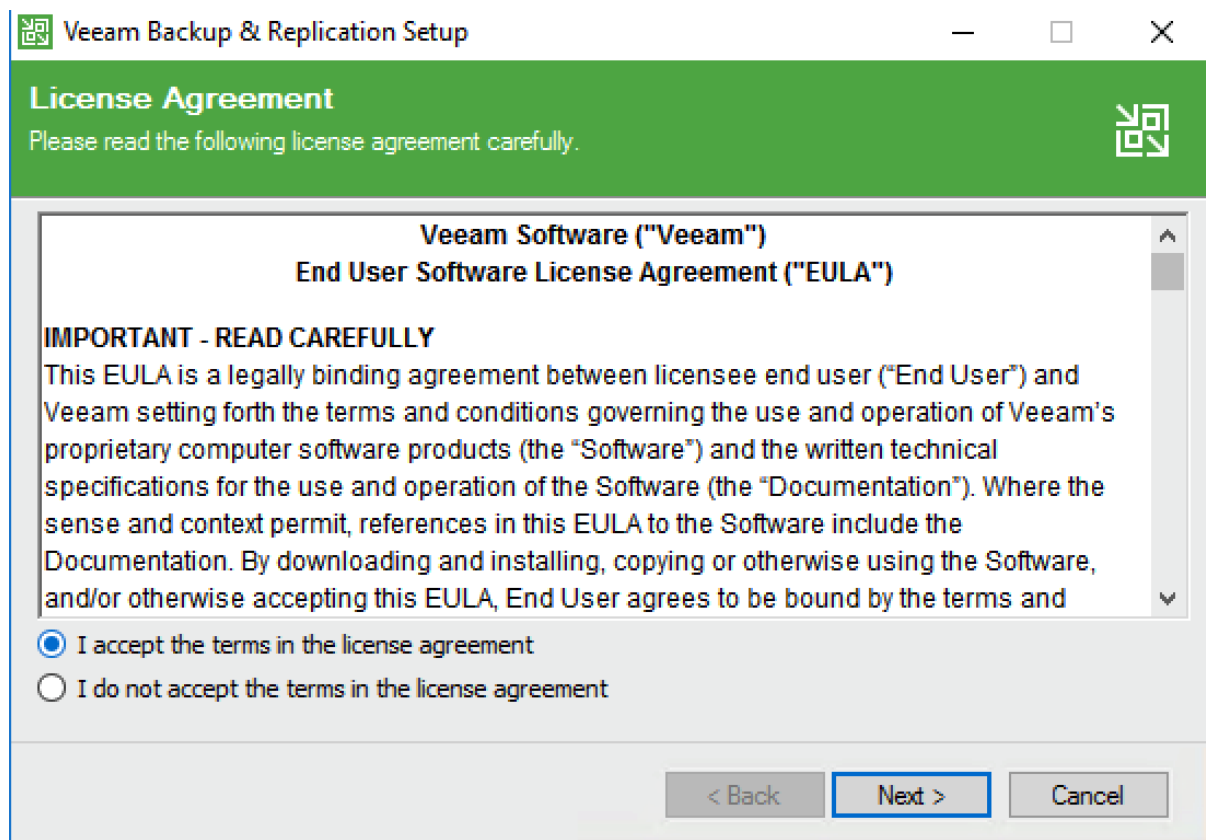
Open

Run as administrator

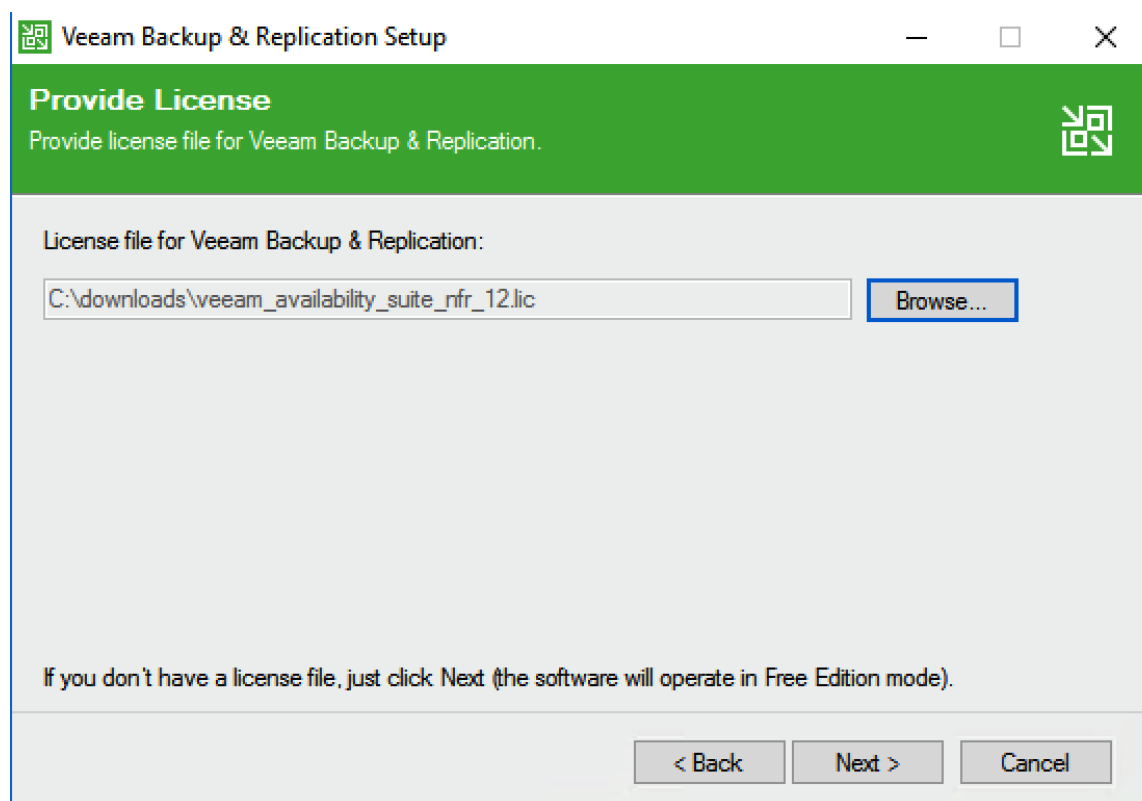
- 3. Click Install Link.



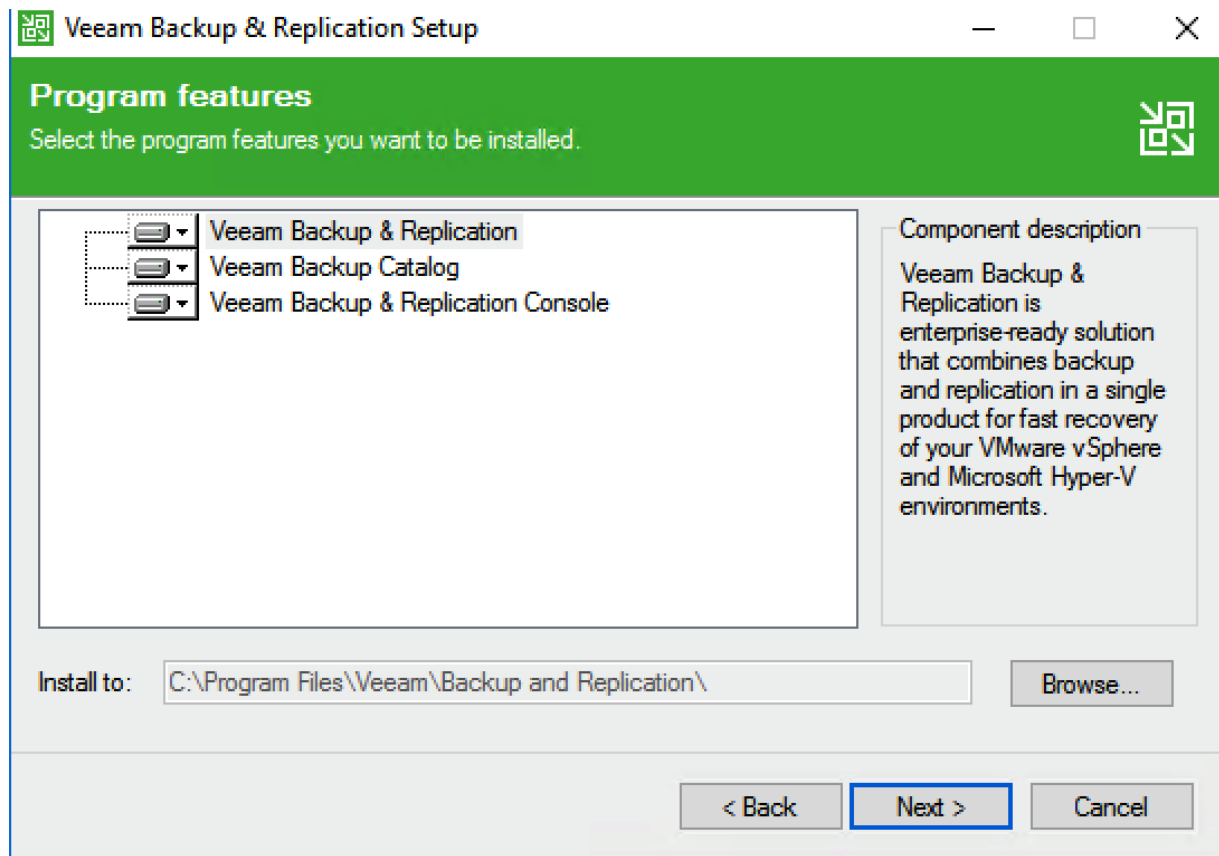
4. Accept the License Agreement.




5. Browse to the license file and click Next.



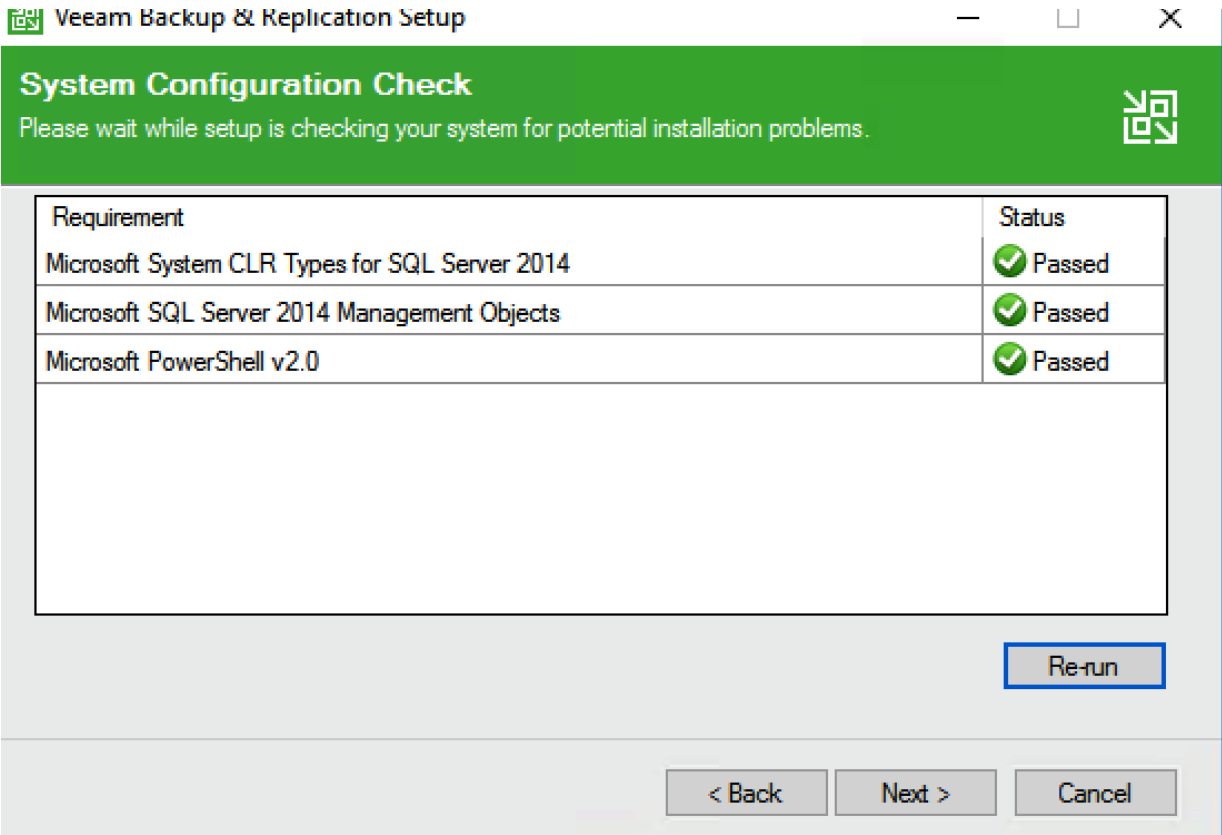
6. Click Next on Program features.



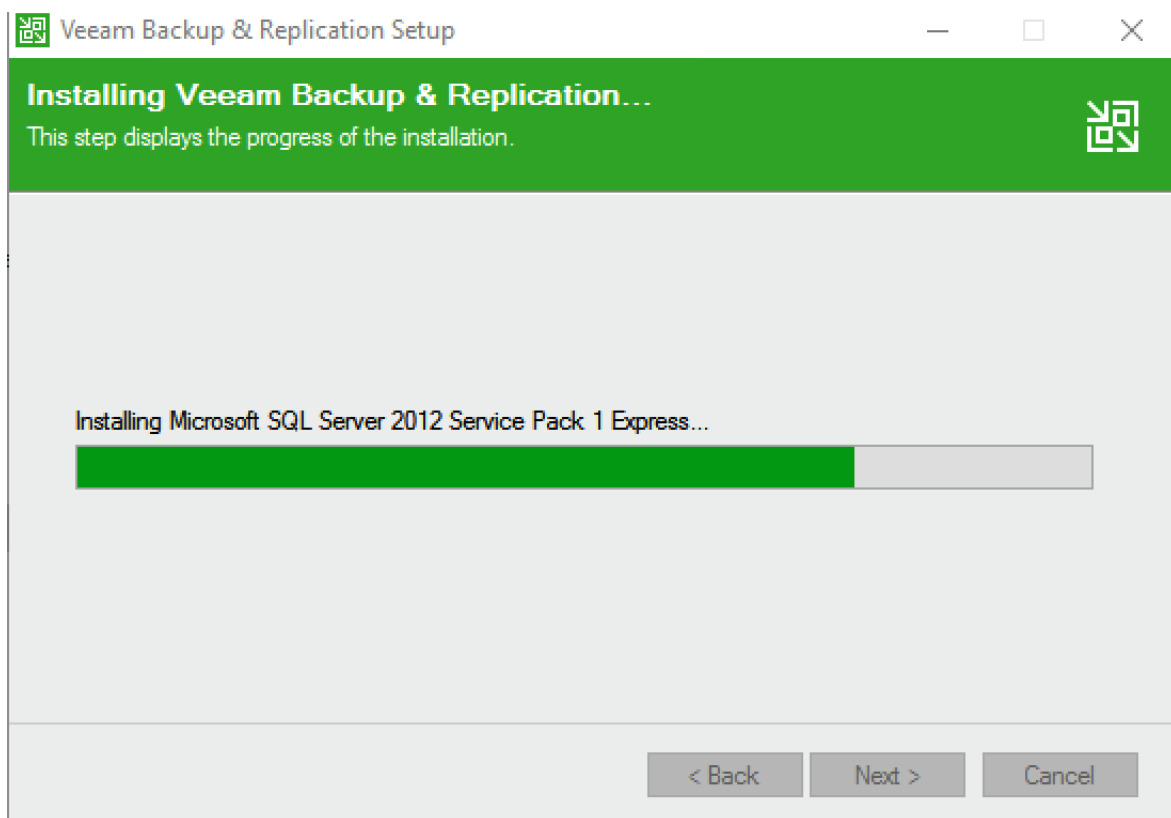
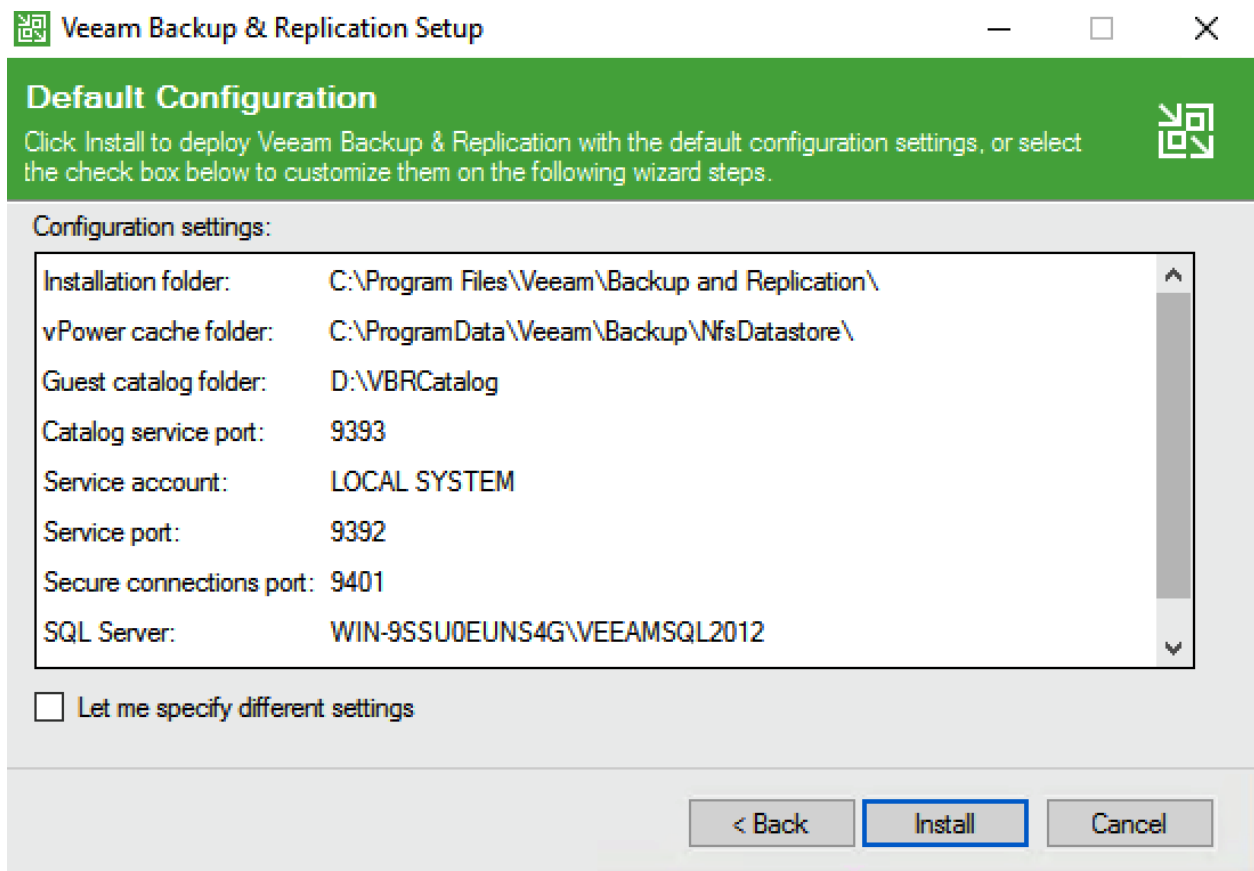
7. During System Check, Veeam verifies the SQL Server Installation and pre-requisite software components, click Install.



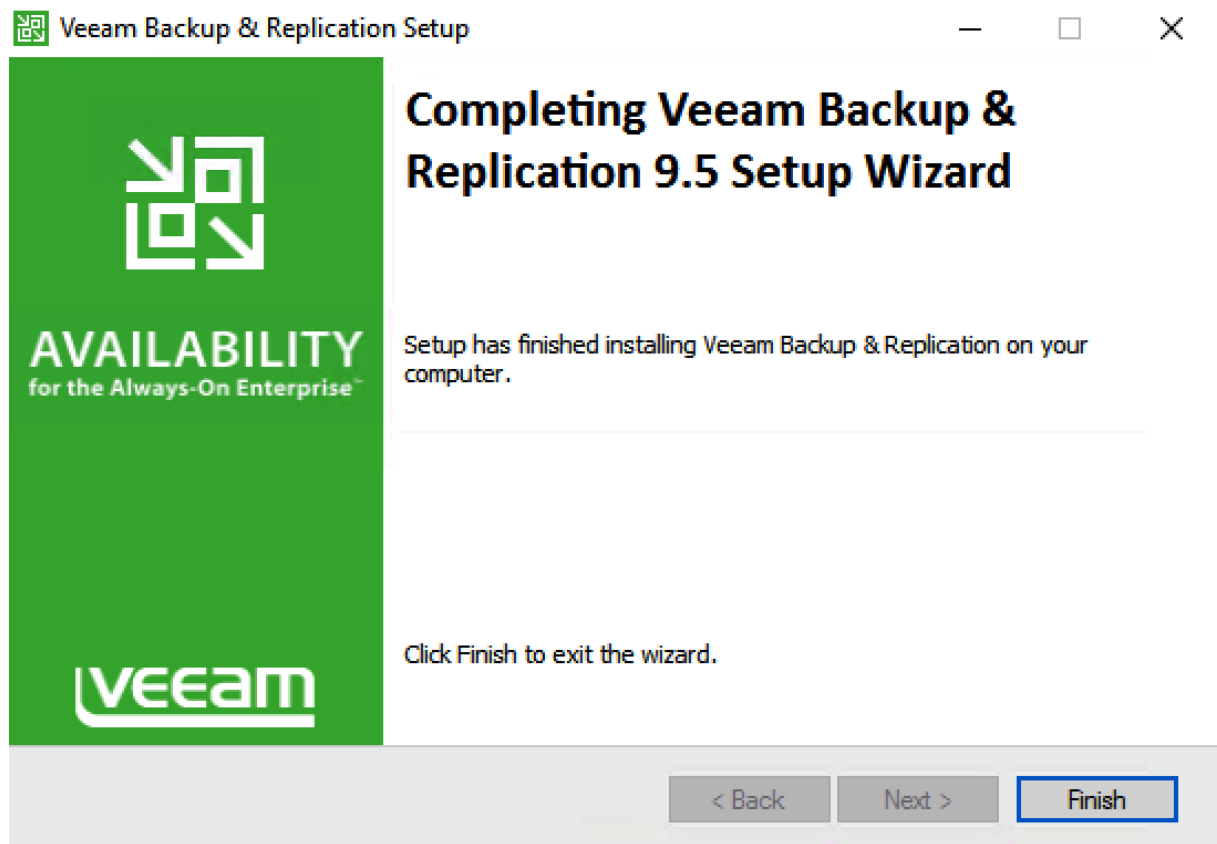
- This will install all the required dependencies. When the system Check passes, click Next.



9. Accept the default Installer locations and click Install.



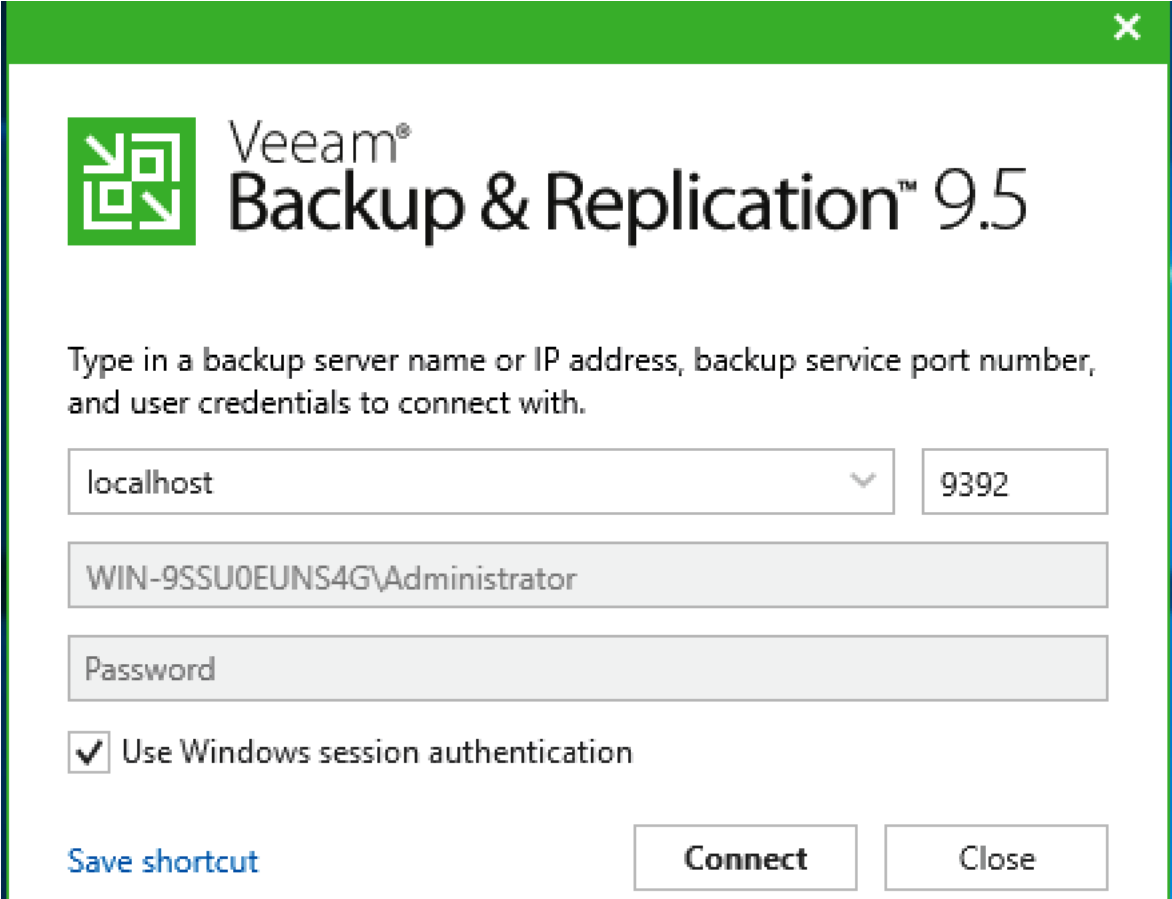
10. Click Finish, when the installation completes.



Configure Veeam Availability Suite 9.5

To configure Veeam Availability Suite 9.5, complete the following steps:

1. From the desktop, Open the Veeam Backup & Replication Console.
2. Click Connect on local host.

A screenshot of the Veeam Backup & Replication 9.5 connection dialog. The window has a green title bar with a close button. The Veeam logo and product name are at the top. Below is a text prompt asking for backup server details. There are three input fields: a dropdown menu showing 'localhost', a text box with '9392', and a text box with 'WIN-9SSU0EUNS4G\Administrator'. A password field is also present. A checkbox for 'Use Windows session authentication' is checked. At the bottom are three buttons: 'Save shortcut', 'Connect', and 'Close'.

Veeam®
Backup & Replication™ 9.5

Type in a backup server name or IP address, backup service port number, and user credentials to connect with.

localhost 9392

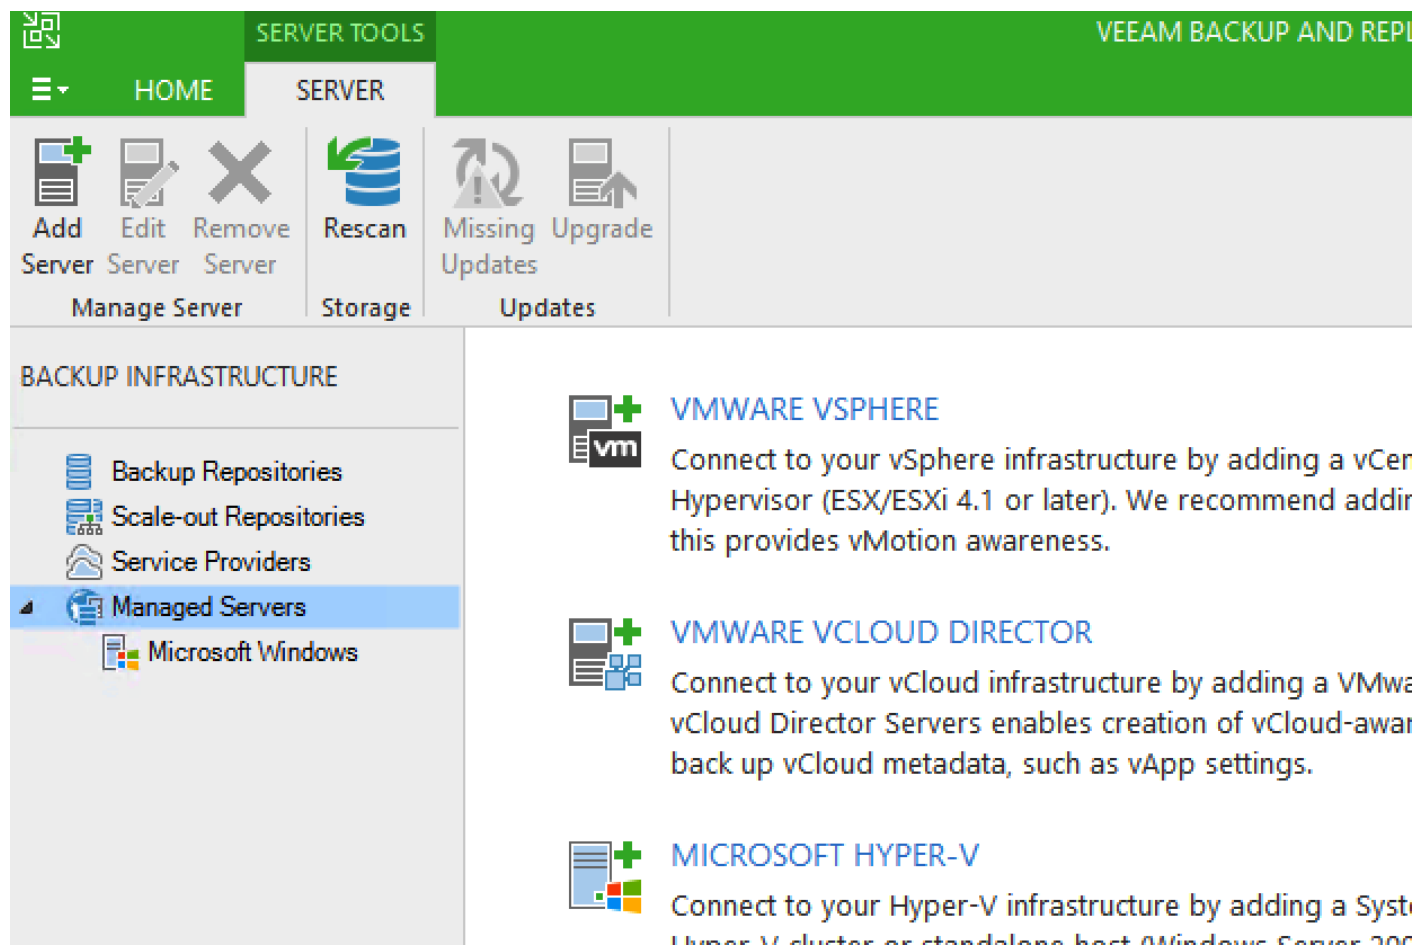
WIN-9SSU0EUNS4G\Administrator

Password

☒ Use Windows session authentication


[Save shortcut](#) **Connect** **Close**

3. By default, Veeam uses the D Drive as the Veeam Repository. This is the repository created through Disk Volume.
4. Right-click Managed Server and click 'Add Server.'
5. Select VMWare VSphere and add the vCenter URL of HyperFlex Cluster, click Next.



6. Enter the vCenter credentials and click Next.

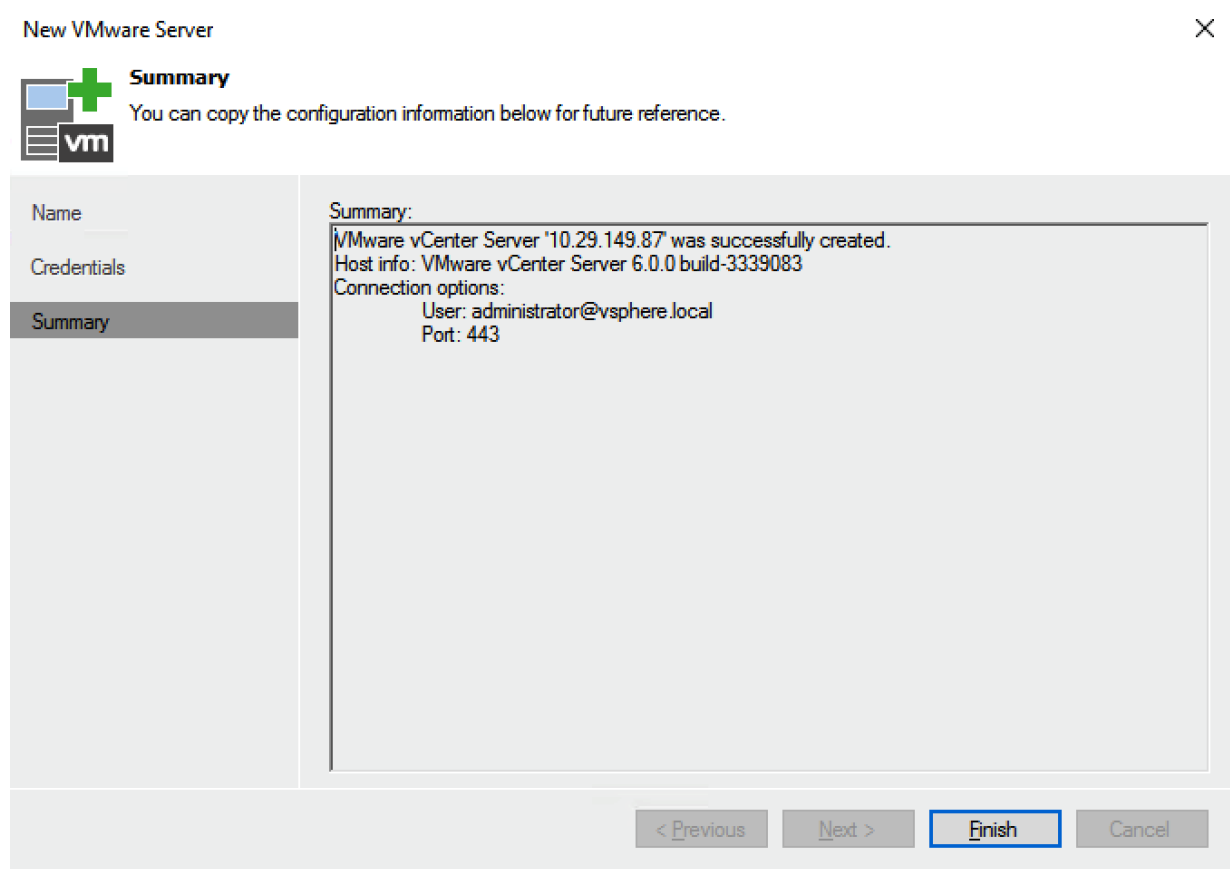
New VMware Server ✕

 **Name**
Specify DNS name or IP address of VMware server.

Name	DNS name or IP address: <input type="text" value="10.29.149.87"/>
Credentials	
SSH Connection	Description: <input type="text" value="Created by WIN-9SSU0EUNS4G\Administrator at 1/12/2017 3:40 PM."/>
Summary	


< Previous **Next >** Finish Cancel

- When Veeam Console collects all the deployment details from vCenter, click Finish.



8. Click Backup Proxies in the right navigation windows, choose VMware Backup Proxy and edit Properties.

Edit VMware Proxy ✕

 **Server**
Choose server for new backup proxy. You can only select between Microsoft Windows servers added to the managed servers which are not proxies already.

<p>Server</p> <p>Traffic Rules</p> <p>Summary</p>	<p>Choose server:</p> <p>WIN-9SSU0EUNS4G Add New...</p> <p>Proxy description:</p> <p>Created by Veeam Backup & Replication</p> <p>Transport mode:</p> <p>Automatic selection Choose...</p> <p>Connected datastores:</p> <p>Automatic detection (recommended) Choose...</p> <p>Max concurrent tasks:</p> <p>2 ✓</p>
--	--

< Previous
Next >
Finish
Cancel

9. Edit Max Concurrent Task to be equal to Number of physical cores in S3260 minus 2. In the present deployment, there is a dual 12-core Intel processor and therefore you can increase the Max Concurrent Task to 22
10. Under Transport Mode, click Choose and make sure that “failover to network mode, if primary mode fails, or is unavailable” is checked. This option is checked by default.

Transport Mode
✕

Backup proxy transport mode:

☒ **Automatic selection**
Data retrieval mode is selected automatically by analyzing backup proxy configuration and reachable VMFS and NFS datastores. Transport modes allowing for direct storage access will be used whenever possible.

☐ **Direct storage access**
Data is retrieved directly from shared storage, without impacting production hosts. For block storage, backup proxy server must be connected into SAN fabric via hardware or software HBA, and have VMFS volumes mounted.

☐ **Virtual appliance**
Data is retrieved directly from storage through hypervisor I/O stack by hot adding backed up virtual disks to a backup proxy VM. Datastores containing protected VMs must be connected to a host running backup proxy VM.

☐ **Network**
Data is retrieved from storage through hypervisor network stack using NBD protocol over host management interface. This mode has no special setup requirements. Recommended for 10 Gb Ethernet or faster.

Options

☒ Failover to network mode if primary mode fails, or is unavailable


☐ Enable host to proxy traffic encryption in Network mode (NBDSSL)


OK

Cancel

11. Click Finish.
12. Click Backup Repository in the right navigation window, select the Default Backup Repository and edit Properties.
13. Click Next until you reach the Repository Windows.
14. Increase the 'Limit Max Concurrent Task' to be 22 (Number of physical cores in S3260 minus 2).


Edit Backup Repository ✕

 **Repository**
Type in path to the folder where backup files should be stored, and set repository load control options.

Name	Location
Type	Path to folder: <input type="text" value="D:\Backup"/> Browse...
Server	 Capacity: ... Populate
Repository	Free space: ...
Mount Server	Load control
Review	Running too many concurrent tasks against the same repository may reduce overall performance, and cause I/O operations to timeout. Control storage device saturation with the following settings:
Apply	<input checked="" type="checkbox"/> Limit maximum concurrent tasks to: <input type="text" value="22"/> <input type="checkbox"/> Limit read and write data rates to: <input type="text"/> MB/s
Click Advanced to customize repository settings Advanced...	
< Previous Next > Finish Cancel	

15. Click Advanced and Enable “Use Per-VM backup File.”

Edit Backup Repository ✕

 **Repository**
Type in path to

Name	Storage Compatibility Settings ✕
Type	<input checked="" type="checkbox"/> Align backup file data blocks Allows to achieve better deduplication ratio on deduplicating storage devices leveraging constant block size deduplication. Increases the backup size when backing up to raw disk storage.
Server	<input type="checkbox"/> Decompress backup data blocks before storing VM data is compressed by backup proxy according to the backup job compression settings to minimize LAN traffic. Uncompressing the data before storing allows for achieving better deduplication ratio on most deduplicating storage appliances at the cost of backup performance.
Repository	<input type="checkbox"/> This repository is backed by rotated hard drives Backup jobs pointing to this repository will tolerate the disappearance of previous backup files by creating new full backup, clean up backup files no longer under retention on the newly inserted hard drives, and track backup repository location across unintended drive letter changes.
Mount Server	<input checked="" type="checkbox"/> Use per-VM backup files Per-VM backup files may improve performance with storage devices benefiting from multiple I/O streams. This is the recommended setting when backing up to deduplicating storage appliances.
Review	OK Cancel
Apply	Click Advanced to customize repository settings Advanced...
< Previous Next > Finish Cancel	

16. Click Finish.
17. Add the MaxSnapshotsPerDatastore parameter in Registry.



The default Number of Snapshots per datastore is 4, you may alter concurrent snapshots executed by Veeam on HX datastore and should consider the intensity of the IO workload on the HX Cluster during backup jobs. For instance, if the HX Cluster is not under heavy IO intensive transactions during Veeam backup jobs, then the Max Snapshots per Datastore can be increased.

18. Using the Registry Editor, go to HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication\; add a RegDWORD with the name 'MaxSnapshotsPerDatastore' value greater than 4.

Registry Editor

File Edit View Favorites Help

Computer

- HKEY_CLASSES_ROOT
- HKEY_CURRENT_USER
- HKEY_LOCAL_MACHINE
 - BCD00000000
 - HARDWARE
 - SAM
 - SECURITY
 - SOFTWARE
 - Classes
 - Clients
 - Cygwin
 - Description
 - Intel
 - Microsoft
 - ODBC
 - Partner
 - Policies
 - RegisteredApplications
 - Veeam
 - Veeam Backup and Replication
 - Veeam Backup Catalog
 - Veeam Mount Service
 - WOW6432Node
 - SYSTEM
 - HKEY_USERS
 - HKEY_CURRENT_CONFIG

Edit DWORD (32-bit) Value

Value name: MaxSnapshotsPerDatastore

Value data: 10

Base:

- ☒ Hexadecimal
- ☐ Decimal

OK Cancel

Name	Type	Value
IsComponentsU...	REG_DWORD	0x00000000 (0)
LoggingLevel	REG_DWORD	0x00000004 (4)
MaxLogCount	REG_DWORD	0x0000000a (10)
MaxLogSize	REG_DWORD	0x00002800 (10240)
SecureConnecti...	REG_DWORD	0x000024b9 (9401)
SqlDatabaseNa...	REG_SZ	VeeamBackup
SqlInstanceName	REG_SZ	VEEAMSQL2012
SqlLockInfo	REG_SZ	<CLockInfo xmlns:xsd="http://www.w3.org/2001/...
SqlLogin	REG_SZ	
SqlSecuredPass...	REG_SZ	
SqlServerName	REG_SZ	WIN-S448941PTOP
VddkReadBuffer...	REG_DWORD	0x00000000 (0)
VNXBlockNavi...	REG_SZ	C:\Program Files\Veeam\Backup and Replication\...
VNXeUemcliPath	REG_SZ	C:\Program Files\Veeam\Backup and Replication\...
WorkWithoutSQL	REG_DWORD	0x00000000 (0)
MaxSnapshotsP...	REG_DWORD	0x00000000 (0)

Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication

19. Restart the Windows 2016 Server for Registry Settings to be applied.

This completes the deployment of Veeam Availability Suite 9.5 on Cisco UCS S3260 Storage server with HyperFlex Cluster.

Validation

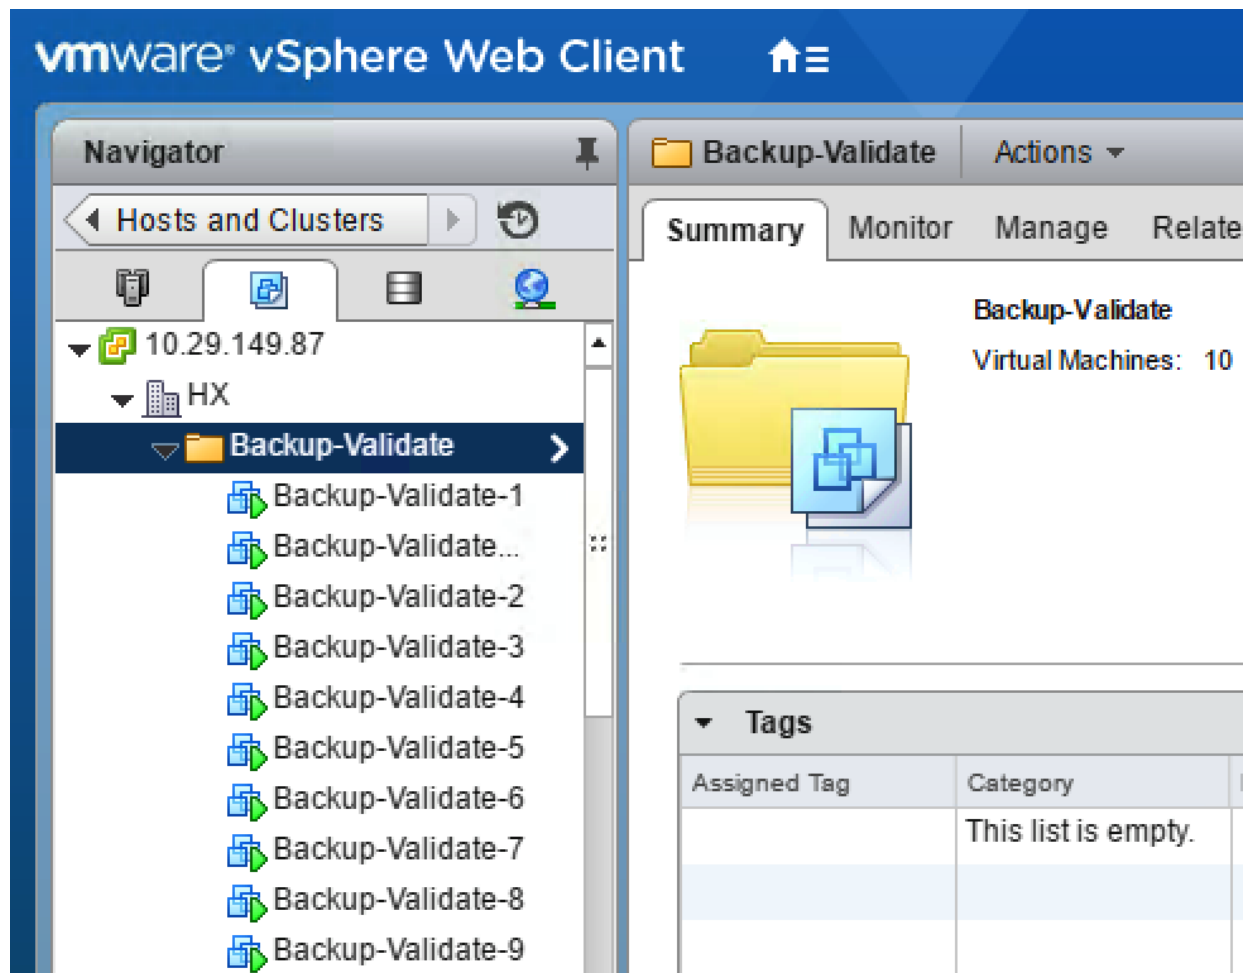
This section describes the test executed to validate the Veeam Backup and Recovery Solution on the Cisco HyperFlex platform. This solution and its scenarios are validated with high-level backup, Replication, Failover and Failback task between HyperFlex Cluster and Cisco UCS S3260 Storage Server.

Some of the important test executed are as follows:

- Single Site Backup & Replication
 - Create Backup of application VM on HX Cluster to Veeam Backup Repository
 - Restore File, Instant Recovery and Restore Entire VM to HX Cluster
 - Create Replica, Failback and Failover of application VM on HX Cluster to another HX Cluster

The validation below, displays a successful backup of several VM on HX cluster to Cisco UCS S3260 target repository.

1. Identify the source VM on HX Cluster for Backup. It is a good practice to have the backup VM moved to a VM Folder. This enables a single-click sentinel snapshot of all the VM through VM folder. Right-click the VM folder, go to Cisco HX Data platform > Snapshot Now. Create a HX Native/Sentinel Snapshot.



2. Go to Veeam Backup and Replication Console and Create Backup job.

- 137

SESSION TOOLS

VEEAM BACKUP AND REPLICATION

HOMEVIEW

SESSION

Stop

Statistics Report

Actions

Details

BACKUP & REPLICATION

Jobs

Backup

Replication

Backups

Disk

Replicas

Ready

Last 24 Hours

Running (1)

Success

Failed

BACKUP & REPLICATION

BACKUP INFRASTRUCTURE

VIRTUAL MACHINES

STORAGE INFRASTRUCTURE

TAPE INFRASTRUCTURE

FILES

Type in an object name to search for

JOB NAME	SESSION TYPE	STATUS	START TIME ↑
Validate Backup (Active Full)	Backup	10% completed	1/26/2017 2:51 PM

Job progress:

10%

0 of 4 VMs

SUMMARY

DATA

STATUS

THROUGHPUT

Duration: 0:02:44

Processed: 33.1 GB (10%)

Success: 0

ed: 492.9 MB/s

Processing rate: 414 MB/s

Read: 33.1 GB

Warnings: 0

Bottleneck: Source

Transferred: 9.6 GB (3.5x)

Errors: 0

NAME

STATUS

ACTION

DURATI...

Backup-Validate-2

12%

Processing Backup-Validate-10

0:02:21

Backup-Validate-3

11%

All VMs have been queued for processing

4. Validate Successful completion of Backup job.

The screenshot displays the Veeam Backup and Replication interface. The top navigation bar includes 'HOME', 'VIEW', and 'JOB' tabs. The 'JOB' tab is active, showing a 'Validate Backup' job. The job is in a 'Stopped' state with a 'Success' status. The summary section provides details on the job's duration (0:11:19), processing rate (352 MB/s), and bottleneck (Source). The data section shows that 201.7 GB (100%) was processed, 190.3 GB was read, and 19.3 GB (9.8x) was transferred. The status section indicates 4 successes, 0 warnings, and 0 errors. The throughput section shows a peak of 492.9 MB/s. The bottom section lists several backup jobs, all of which are successful.

NAME	TYPE	OBJECTS	STATUS	LAST RES...	NEXT RUN	TARGET
Validate Backup	VMware Back...	4	Stopped	Success	<not scheduled>	Default Backup Rep

SUMMARY		DATA		STATUS		THROUGHPUT
Duration:	0:11:19	Processed:	201.7 GB (100%)	Success:	4	492.9 MB/s
Processing rate:	352 MB/s	Read:	190.3 GB	Warnings:	0	
Bottleneck:	Source	Transferred:	19.3 GB (9.8x)	Errors:	0	

NAME	STATUS	ACTION	DURATI...
Backup-Validate-2	Success	Job started at 1/26/2017 2:51:41 PM	
Backup-Validate-3	Success	Building VMs list	0:00:06
Backup-Validate-4	Success	VM size: 320.0 GB (83.5 GB used)	
Backup-Validate-10	Success	Changed block tracking is enabled	

Validated Hardware and Software

Table 4 lists all the software and hardware components deployed to validate the design for Cisco HyperFlex with Veeam Backup and Replication

Table 4 Software and Hardware Components Deployed in this CVD

	Components	Software Version	Comments
Compute & Storage	Cisco UCS 3260 M4 Rack Server	3.1(2b)	Directly managed through Fabric Interconnect. Veeam AS is installed on the same. Provides Storage Veeam Repository
	Cisco HX220c M4		Hyper Converged node for HX Cluster
	Cisco HX240c M4		Hyper Converged Node for HX Cluster
Management	Cisco UCS Manager	3.1(2b)	UCS Management for all servers directly attached to Fabric Interconnects

	Components	Software Version	Comments
Backup and Replication	Veeam Availability Suite	9.5	Pre-configured with Veeam Backup Server, Veeam Proxy , Veeam Repository
	Operating System	Windows 2016 DataCenter Edition	
Hyper Converged Software	Cisco HX Data Platform	HX Data Platform Release 1.8a	
Virtualization	VMWare VSphere	6.0 U2	
	VMWare vCenter	6.0 U2	
Network	Cisco Nexus 9372PX (N9k-9372PX)	6.1(2)I3(4b)	Cisco Platform Switch for ToR, MoR, EoR deployments; Provides connectivity to users and other networks and deployed in NX-OS Standalone mode
	Cisco UCS 6248UP FI	3.1(2b)	Fabric Interconnect with embedded Cisco UCS Manager

Bill of Materials

The BOM below lists the major components validated, but it is not intended to be a comprehensive list.

Line Number	Part Number	Description	Qty
1.0	HX-SP-220M4SBP1-1A	UCS SP HX220c HyperFlex System w/2xE52690v4,16x32Gmem,1yrSW	1
1.0.1	CON-PSJ1-220SBP1A	UCS SUPP PSS 8X5XNBD, UCS SP HX220c HyperFlex System w2xE526	1
1.1	UCS-CPU-E52690E	2.60 GHz E5-2690 v4/135W 14C/35MB Cache/DDR4 2400MHz	2
1.2	UCS-MR-1X322RV-A	32GB DDR4-2400-MHzRDIMM/PC4-19200/dual rank/x4/1.2v	16
1.3	UCS-HD12TB10K12G	1.2 TB 12G SAS 10K RPM SFF HDD	6
1.4	UCS-SD480G12S3-EP	480GB 2.5 inch Ent. Performance 6GSATA SSD(3X endurance)	1
1.5	UCS-SD120GBKS4-EV	120 GB 2.5 inch Enterprise Value 6G SATA SSD	1
1.6	UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	1
1.7	UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	1
1.8	UCS-SD-64G-S	64GB SD Card for UCS Servers	2
1.9	UCSC-PSU1-770W	770W AC Hot-Plug Power Supply for 1U C-Series Rack Server	2
1.1	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	2
1.11	HXDP-001-1YR	Cisco HyperFlex HX Data Platform SW 1 year Subscription	1
1.11.0.1	HXDP001-1YR	Cisco HyperFlex HX Data Platform SW Subscription 1 Year	1
1.12	UCS-M4-V4-LBL	Cisco M4 - v4 CPU asset tab ID label (Auto-Expand)	1
1.13	UCSC-HS-C220M4	Heat sink for UCSC220 M4 rack servers	2
1.14	HX220C-BZL-M4	HX220C M4 Security Bezel	1
1.15	SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	2
1.16	UCSC-SAS12GHBA	Cisco 12Gbps Modular (non-RAID) SAS HBA	1
1.17	HX-VSP-FND-D	Factory Installed - vSphere SW (End user to provide License)	1
1.18	HX-VSP-FND-DL	Factory Installed - VMware vSphere6 Fnd SW Download	1
2.0	UCS-FI-6248E16-ALL	UCS 6248UP and 16P Expansion Module with 48 Port Licenses	1
2.0.1	CON-PSJ7-F6248ALL	UCS PSS 24X7X4 OS UCS 6248UP and 16P E	1
2.1	UCS-ACC-6248UP	UCS 6248UP Chassis Accessory Kit	1
2.2	UCS-FAN-6248UP	UCS 6248UP Fan Module	2
2.3	UCS-FI-DL2	UCS 6248 Layer 2 Daughter Card	1
2.4	UCS-LIC-10GE	UCS 6200 Series ONLY Fabric Int 1PORT 1/10GE/FC-port license	28
2.5	UCS-FI-E16UP	UCS 6200 16-port Expansion module/16 UP/8p LIC	1
2.5.0.1	CON-PSJ7-FIE16UP	UCS PSS 24X7X4 OS 16prt 10Gb UnifiedPrt/Expnsn mod UCS6200	1
2.6	UCS-PSU-6248UP-AC	UCS 6248UP Power Supply/100-240VAC	2
2.7	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	2
3.0	HX-SP-240M4SBP1-5A	UCS SP HX240c HyperFlex System w/2xE52690v4,16x32Gmem,5yrSW	1
3.0.1	CON-PSJ1-240SBP5A	UCS SUPP PSS 8X5XNBD, UCS SP HX240c HyperFlex System w2xE526	1

Line Number	Part Number	Description	Qty
3.1	UCS-CPU-E52690E	2.60 GHz E5-2690 v4/135W 14C/35MB Cache/DDR4 2400MHz	2
3.2	UCS-MR-1X322RV-A	32GB DDR4-2400-MHzRDIMM/PC4-19200/dual rank/x4/1.2v	16
3.3	UCS-HD12TB10K12G	1.2 TB 12G SAS 10K RPM SFF HDD	15
3.4	UCS-SD16TB12S3-EP	1.6TB 2.5 inch Ent. Performance 6GSATA SSD(3X endurance)	1
3.5	UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	1
3.6	UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	1
3.7	UCSC-PSU2V2-1400W	1400W V2 AC Power Supply (200 - 240V) 2U & 4U C Series	2
3.8	UCS-SD-64G-S	64GB SD Card for UCS Servers	2
3.9	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	2
3.1	UCSC-PCI-1C-240M4	Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCIs	1
3.11	UCS-SD120GBKS4-EB	120 GB 2.5 inch Enterprise Value 6G SATA SSD (boot)	1
3.12	HXDP-001-5YR	Cisco HyperFlex HX Data Platform SW 4 Yr Subscription Add On	1
3.12.0.1	HXDP001-5YR	Cisco HyperFlex HX Data Platform SW Subscription 5 Year	1
3.13	HX240C-BZL-M4SX	HX240C M4 Security Bezel	1
3.14	UCS-M4-V4-LBL	Cisco M4 - v4 CPU asset tab ID label (Auto-Expand)	1
3.15	UCSC-HS-C240M4	Heat sink for UCS C240 M4 rack servers	2
3.16	SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	2
3.17	N20-BBLKD	UCS 2.5 inch HDD blanking panel	8
3.18	UCSC-SAS12GHBA	Cisco 12Gbps Modular (non-RAID) SAS HBA	1
3.19	HX-VSP-FND-D	Factory Installed - vSphere SW (End user to provide License)	1
3.2	HX-VSP-FND-DL	Factory Installed - VMware vSphere6 Fnd SW Download	1
4.0	N9K-C9372PX	Nexus 9300 with 48p 10G SFP+ and 6p 40G QSFP+	1
4.0.1	CON-PSRT-9372PX	PRTNR SS 8X5XNBD Nexus 9300 with 48p 10G SFP+ and 6p 40G	1
4.1	NXOS-703I2.3	Nexus 9500, 9300, 3000 Base NX-OS Software Rel 7.0(3)I2(3)	1
4.2	N3K-C3064-ACC-KIT	Nexus 3K/9K Fixed Accessory Kit	1
4.3	NXA-FAN-30CFM-F	Nexus 2K/3K/9K Single Fan, port side exhaust airflow	4
4.4	N9K-PAC-650W-B	Nexus 9300 650W AC PS, Port-side Exhaust	2
4.5	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	2
5.0	UCSS-S3260	Cisco UCS S3260 Storage Server Base Chassis	1
5.1	CON-OSP-UCSS3260	SNTC 24X7X4OS, Cisco UCS S3260 Storage Server Base Chassis	1
5.2	N20-BBLKD-7MM	UCS 7MM SSD Blank Filler	2
5.3	N20-BKVM	KVM local IO cable for UCS servers console port	1
5.4	UCS-C3K-28HD10	UCS C3X60 2 row of 10TB NL-SAS drives (28 Total) 280TB	1
5.5	UCS-C3K-M4RAID	Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	1
5.6	UCS-C3X60-G2SD48	UCSC C3X60 480GB Boot SSD (Gen 2)	2
5.7	UCS-CPU-E52695E	2.10 GHz E5-2695 v4/120W 18C/45MB Cache/DDR4 2400MHz	2
5.8	UCS-MR-1X161RV-A	16GB DDR4-2400-MHzRDIMM/PC4-19200/single rank/x4/1.2v	8
5.9	UCSC-C3260-SIOC	Cisco UCS C3260 System IO Controller with VIC 1300 incl.	1

Line Number	Part Number	Description	Qty
5.10	UCSC-C3K-M4SRB	UCS C3000 M4 Server Node for Intel E5-2600 v4	1
5.11	UCSC-C3X60-10TB	UCSC C3X60 10TB 4Kn for Top-Load	28
5.12	UCSC-C3X60-BLKP	Cisco UCS C3X60 Server Node blanking plate	1
5.12	UCSC-C3X60-RAIL	UCS C3X60 Rack Rails Kit	1
5.13	UCSC-C3X60-SBLKP	UCS C3x60 SIOC blanking plate	1
5.14	UCSC-HS-C3X60	Cisco UCS C3X60 Server Node CPU Heatsink	2
5.15	UCSC-PSU1-1050W	UCS C3X60 1050W Power Supply Unit	4
5.16	UCSS-S3260-BBEZEL	Cisco UCS S3260 Bezel	1
5.17	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	4

References

Cisco HyperFlex HX220c M4 Node Installation Guide:

http://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HX_series/HX220c_M4/HX220c.html

Cisco HyperFlex HX240c M4 Node Installation Guide:

http://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HX_series/HX240c_M4/HX240c.html

Design and Deployment Guide for Cisco HyperFlex Systems:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/HX171_VSI_ESXi6U2.html

HyperFlex Hardware and Software Interoperability Matrix: <http://www.cisco.com/c/en/us/support/hyperconverged-systems/hyperflex-hx-data-platform-software/products-technical-reference-list.html>

Veeam Availability Suite v9 Installation and Deployment guide: <https://www.veeam.com/videos/veeam-availability-suite-v9-installment-deployment-7554.html>

About the Authors

Anil Dhiman, Technical Marketing Engineer, Cisco Unified Computing Systems, Cisco Systems, Inc.

Anil Dhiman has over 16 years of experience specializing in Data Center solutions on Cisco UCS servers, and Performance Engineering of large-scale enterprise applications. Over the past 6 years, Anil has authored several Cisco Validated Designs for Enterprise Solutions on Cisco Data Center Technologies.

Currently his focus is on Cisco's portfolio of Hyperconverged Infrastructure and Backup Solutions.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the significant contribution and expertise that resulted in developing this document:

- Ulrich Kleidon, Principal Engineer, Cisco Systems, Inc.
- Shawn Lieu, Solutions Architect with the Global Alliances Team, Veeam Software
- Stefan Renner, EMEA Alliance Systems Engineer, Veeam Software
- Jacques Thomas, Solutions Architect, Veeam Software