

Cisco HyperFlex with Veeam Availability Suite

Design Guide for Data Protection on Cisco HyperFlex Systems Through Veeam Availability Suite 9.02 and Cisco UCS S3260 Storage Servers in Single Data Center, ROBO and Multi Data Center Deployments

Last Updated: November 19, 2016



About Cisco Validated Designs

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 Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, 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)

© 2016 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	5
Solution Overview	6
Introduction	6
Audience	6
Purpose of the Document	6
Solution Summary.....	6
Technology Overview	9
Cisco Unified Computing System.....	10
Cisco Unified Computing System Components.....	11
Cisco Fabric Interconnects	11
Cisco UCS S-Series Storage Server	12
Cisco UCS C-Series Rack Mount Servers	13
Cisco HyperFlex HX-Series Nodes	14
Cisco HyperFlex HX220c-M4S Node.....	14
Cisco HyperFlex HX240c-M4SX Node.....	14
Cisco VIC 1227 MLOM Interface Card	15
Cisco HyperFlex Converged Data Platform Software	15
Cisco HyperFlex HX Data Platform Administration Plug-in	16
Cisco HyperFlex HX Data Platform Controller.....	16
Data Operations and Distribution	17
Veeam Availability Suite.....	20
Backup	20
Restore.....	21
Replication.....	22
Failover and Failback	25
Backup Server	26
Backup Proxy	31
Veeam Repository Sizing	31
Solution Design.....	33
Cisco HyperFlex Single Site Backup and Replication.....	34
Deployment Architecture	35
Remote office - Branch Office Replication for Cisco HyperFlex.....	37
Deployment Architecture	38

Multi-Site Backup and Replication for Cisco HyperFlex.....	40
Deployment Architecture	41
Design Consideration.....	43
Validation.....	48
Validated Hardware and Software.....	48
Bill of Materials	50
References	55

Executive Summary

In addition to consolidation of network, compute and storage resources, Hyperconverged solutions allow ease of deployment, integrated management, ability to start small and most importantly low operating and inception cost. Thus, it has become an important building block for organizations looking to either consolidate their distinct infrastructure components or simplify their data centers. Hyperconvergence allows customers to install scalable compute and storage infrastructure across small to large distributed data center or remote offices collocated across geographies.

One of the major challenge customers are facing today, is data protection provisioning for their virtualized applications deployed in Hyperconverged Infrastructure, across several remote offices and data centers in different geographies. Moreover, the operating cost to manage and deploy several data protection end points across the globe, rises exponentially. Hence, customers today are looking for flexible, agile, efficient and scalable data protection platform that is fast and easy to deploy. The data protection capabilities should also include centralized backup and replication, availability of data at all times including swift data recovery and disaster recovery.

The Cisco HyperFlex™ System's 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 Platforms built-in protection tools, Veeam Availability Suite expands the protection of your data with local and remote **backups. 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.

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, that addresses the challenge of delivering agile protection for Cisco HyperFlex platform.

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 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, manage, and scale 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 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 the Document

This document elaborates on various design architecture and guidelines to deploy Veeam Backup and Recovery Solution on Cisco HyperFlex Cluster. The document illustrates, various deployment scenarios such as Single Site, Multi Site across different data centers and centralized replication for Remote Office and Branch Offices.

Solution Summary

The solution for Cisco HyperFlex, Cisco S3260 Storage Server and Cisco UCS C240 M4 Rack 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.

As Backup and Replication of virtualized environments can be complex and may involve a few virtualized applications located in a singled data center domain or campus to more complex distributed virtualized environments located either in branch offices or in remote data centers, the present solution is divided into three important deployment scenarios. These deployment scenarios are:

- **Cisco HyperFlex Single Site Backup and Replication:** This design provides Backup and Replication of application 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 either on a separate Cisco HyperFlex Cluster or on a standalone ESXI cluster.
- **Remote Office Branch Office (ROBO) Replication:** Small to Large Scale organizations have Remote or Branch site infrastructure located across geographical locations. Cisco HyperFlex cluster are the ideal choice for Remote Office deployments as customers can start with few nodes in a cluster and thereafter can increase the compute or storage resources, independently. Moreover, with centralized management of Cisco UCS infrastructure, customers can manage Cisco HyperFlex infrastructure centrally and thus dramatically lower their IT operational cost. With Veeam solution for Cisco HyperFlex ROBO deployments, customers can now benefit from centralized backup on Cisco UCS S3260 server located in their main Data Center. Above-all, customers can opt-in for Veeam WAN accelerators on the Remote sites and can benefit from faster backup and replication of their Remote enterprise VMs.

The design of Cisco HyperFlex with Veeam for Remote Offices, details the key infrastructure components, design guidelines and validation, which help the customers to provide availability of their virtualized applications in Remote Office or Branch Offices.

- **Multi-Site Backup and Replication for Cisco HyperFlex:** This design helps customers to provision Backup and Replication solution for their virtualized infrastructure located across multiple Data Centers deployments across different geographies. In Cisco HyperFlex with Veeam, each of the Data Center or campus creates Backup and Replication task for their virtualized infrastructure. The remote Backup Repository is synced with the Primary Repository and in the event of Remote site failure; Application VMs can be recovered from the Primary Data Center Repository.

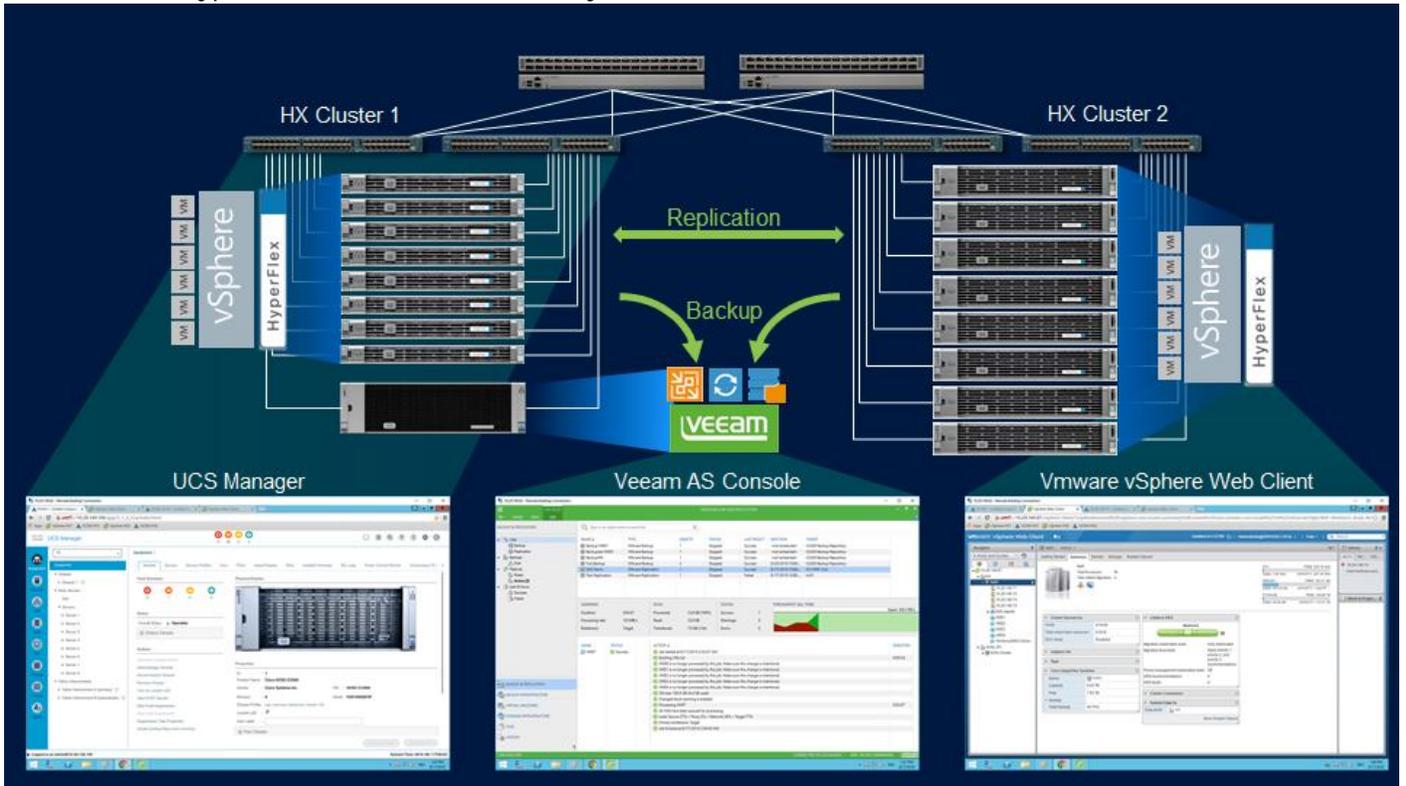
Besides the functional benefit of this design, Cisco UCS Central and Veeam Enterprise Manager help customers to ease the management of their distributed Cisco HyperFlex infrastructure and Veeam Backup environments. With Cisco UCS Central, customers can manage their UCS infrastructure distributed across multiple UCS domains or Data Centers whereas with Veeam Enterprise Manager, distributed Veeam Deployments across Data Centers can be managed through a single management windows, allowing customers to perform backup and replication jobs across the entire backup infrastructure, and providing enhanced reporting options.

Figure below provides a high-level view of Cisco HyperFlex with Cisco S3260 Storage server and Veeam Availability Suite and elaborates on:

- 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 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
- Veeam Availability Suite
- Windows 2012 R2 Data Centre Edition for Veeam Availability Suite

This design document uses the following models of above mentioned infrastructure components.

- Cisco UCS S3260 Storage Server
- Cisco UCS C240M4 Rack 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. 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 WAN Accelerator are dedicated components responsible for global data caching and data deduplication. On each WAN accelerator, Veeam Backup and Replication installs the Veeam WAN Accelerator Service responsible for WAN acceleration tasks.
- Veeam Backup Enterprise Manager collects data from backup server and enables you to run backup **and replication jobs across the entire backup infrastructure through a single “pane of glass”,** edit them 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.

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

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. It is an integrated, scalable, multi chassis platform in which all resources participate in a unified management domain.

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, which 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, which 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 one expansion slot.

Figure 2 Cisco UCS 6248UP Fabric Interconnect



Cisco UCS S-Series Storage Server

Cisco UCS S3260 Storage Server



Cisco UCS S3260 Storage Server was used for this Design Guide

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 S3260 Storage Server



Cisco UCS C-Series Rack Mount Servers

Cisco UCS C240 M4 Rack Server



Cisco UCS C240 M4 LFF rack server was used for this Design Guide

The enterprise-class Cisco UCS C240 M4 server extends the capabilities of the Cisco UCS portfolio in a 2RU form factor. Based on the Intel Xeon processor E5-2600 v4 and v3 series, it delivers an outstanding combination of performance, flexibility, and efficiency. In addition, the Cisco UCS C240 M4 offers outstanding levels of internal memory and storage expandability with exceptional performance. It delivers:

- Up to 24 DDR4 DIMMs at speeds up to 2400 MHz for improved performance and lower power consumption
- Up to 6 PCI Express (PCIe) 3.0 slots (4 full-height, full-length)
- Up to 24 small form factor (SFF) drives or 12 large form factor (LFF) drives, plus two (optional) internal SATA boot drives
- Support for 12-Gbps SAS drives
- A modular LAN-on-motherboard (mLOM) slot for installing a next-generation Cisco virtual interface card (VIC) or third-party network interface card (NIC) without consuming a PCIe slot
- 2 x 1 Gigabit Ethernet embedded LOM ports
- Supports up to two double-wide NVIDIA graphics processing units (GPUs), providing a graphics-rich experience to more virtual users
- Excellent reliability, availability, and serviceability (RAS) features with tool-free CPU insertion, easy-to-use latching lid, hot-swappable and hot-pluggable components, and redundant Cisco® Flexible Flash (FlexFlash) SD cards

For more information, see: <http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c240-m4-rack-server/datasheet-c78-732455.html>

Figure 4 Cisco UCS C240 M4 LFF Rack Server



Cisco UCS C220 M4 Rack Server



Cisco UCS C220 M4 rack server was used for this Design Guide

The enterprise-class Cisco UCS C220 M4 server extends the capabilities of the Cisco Unified Computing System (UCS) portfolio in a one rack-unit (1RU) form-factor. The Cisco UCS C220 M4 uses the power of the latest Intel® Xeon® E5-2600 v3 and v4 Series processor family CPUs with up to 1536 GB of RAM (using 64 GB DIMMs), 8 Small Form-Factor (SFF) drives or 4 Large Form-Factor (LFF) drives, and up to 80 Gbps throughput connectivity. The Cisco UCS C220 M4 Rack Server can be used standalone, or as integrated part of the Unified Computing System. It has 24 DIMM for up to 1536 GB total memory capacity. It supports one connector for the Cisco VIC 1225, 1227 or 1380 adapters, which provide Ethernet and FCoE.

For more information, see: <http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c220-m4-rack-server/datasheet-c78-732386.html>

Figure 5 Cisco UCS C220 M4 Rack Server



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 also 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.

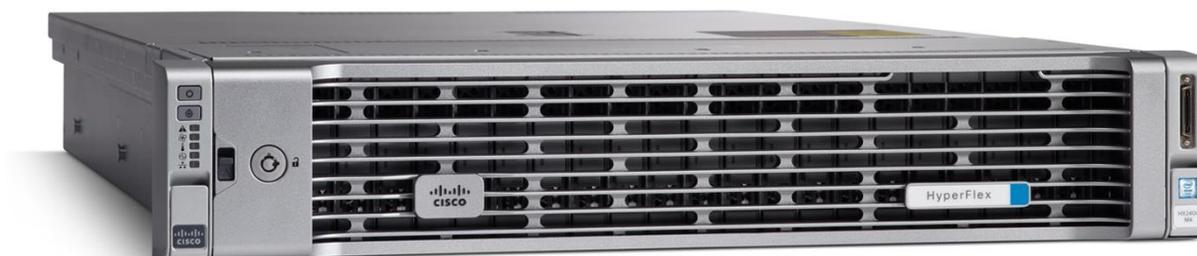
Figure 6 Cisco HyperFlex HX220c-M4S Node



Cisco HyperFlex HX240c-M4SX Node

The Cisco HyperFlex HX240c-M4SX 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 Cisco HyperFlex 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:

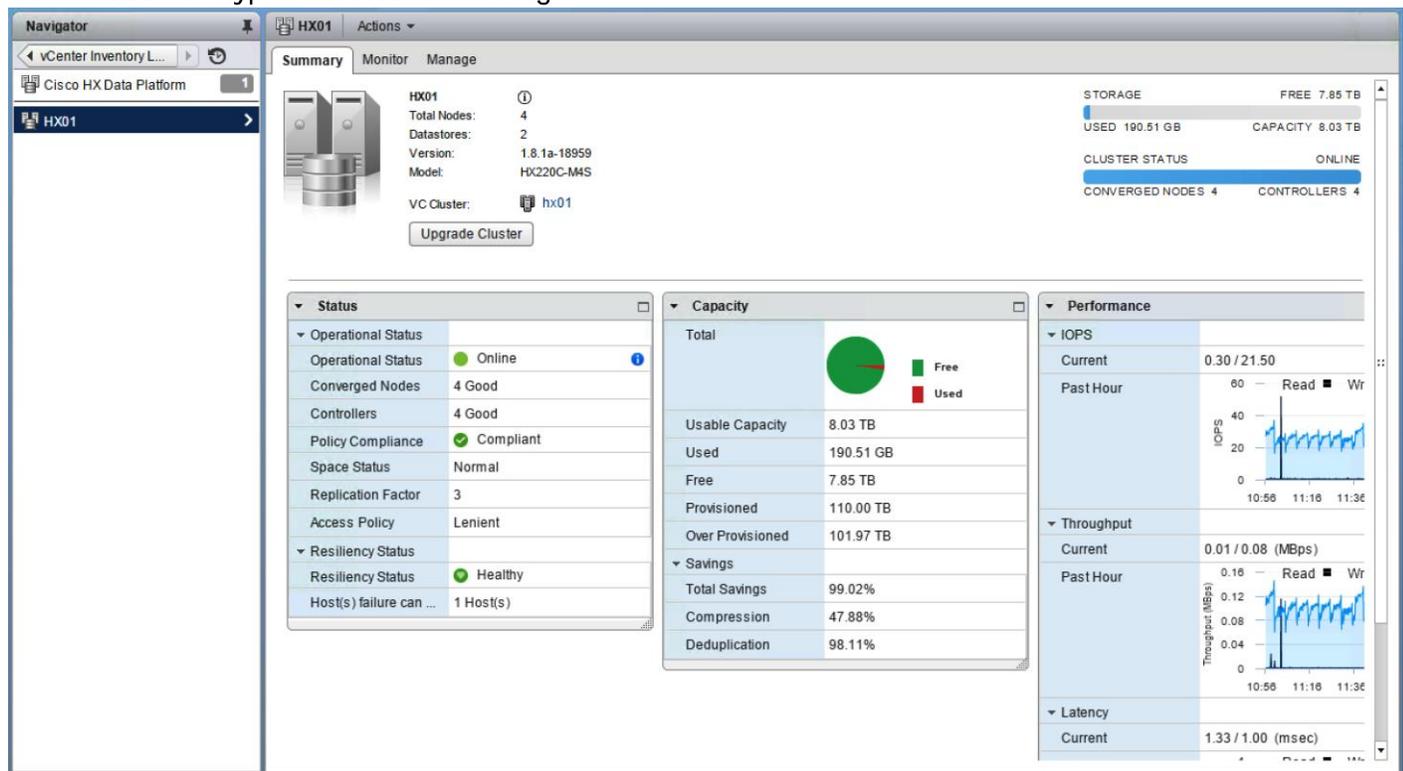
- Replication of all written data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- Deduplication is always on, helping reduce storage requirements in which multiple operating system instances in client 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 data stores, 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 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 data stores 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 also 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 three (RF=3), is a default setting and also a recommended best practice for HyperFlex systems.

- Replication Factor 3: For every I/O write committed to the storage layer, two 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 neither stored on the same disks, nor on the same nodes of the cluster. This setting can tolerate simultaneous failures of two disks, or two 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, one 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 neither 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.

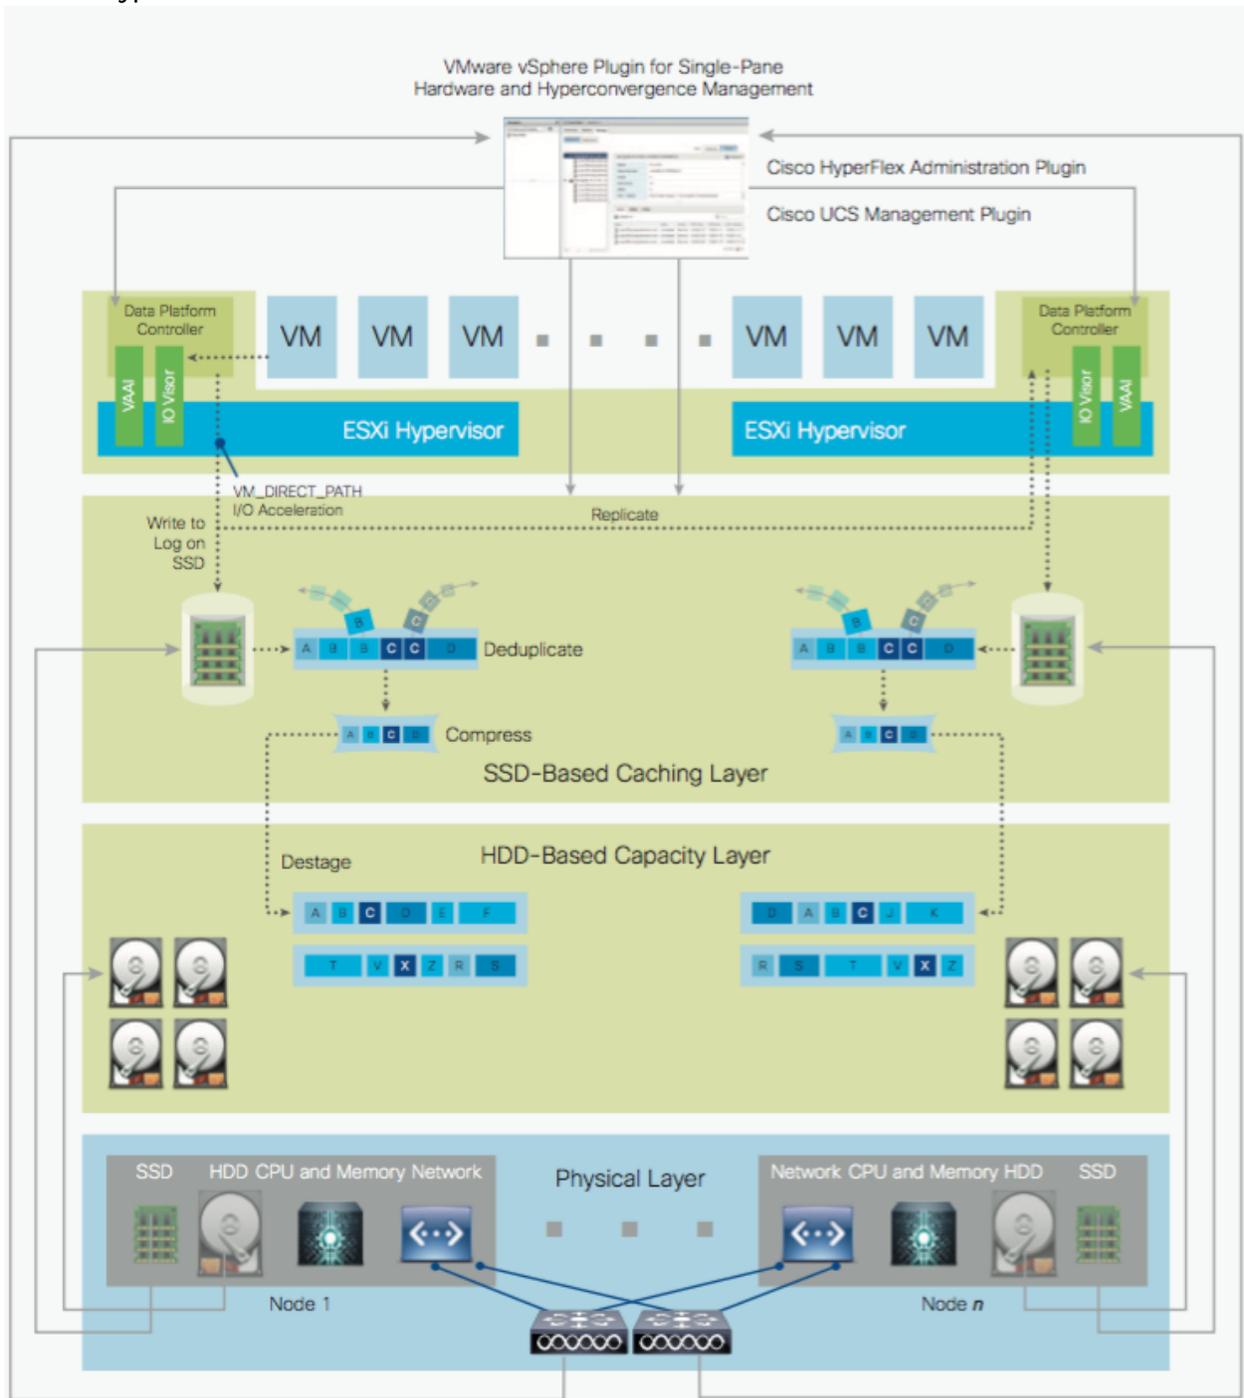
Data Write Operations

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

Veeam Availability Suite

Backup

Veeam Backup and 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 and 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 and 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 other.

Use of the image-based approach allows Veeam Backup and 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 and 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 and 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 and Replication copies only those data blocks that have changed since the last backup job run. To keep track of changed data blocks, Veeam Backup and Replication uses different approaches, including VMware's Changed Block Tracking (CBT) technology.

Changed Block Tracking

To perform incremental backup, Veeam Backup and 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 and Replication employs VMware vSphere Changed Block Tracking (CBT) – a native VMware feature. Instead of scanning VMFS, Veeam Backup and

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 and Replication offers a number of recovery options for various disaster recovery scenarios:

- Veeam Explorer enables you to restore Single Application 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 and 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 Backup and Replication can perform file-level restore as a preparatory step for application items restore. However, the database files may be huge and require many network resources. For this reason, if you restore application items from Microsoft SQL and Oracle VMs, Veeam Backup and Replication can mount the content of the backup file directly to the original VM.

Veeam Explorers are tools included in all editions of Veeam Backup and Replication. As of v9, following Explorers are available:

- Veeam Explorer for Active Directory
- Veeam Explorer for SQL Server
- Veeam Explorer for Exchange
- Veeam Explorer for SharePoint
- Veeam Explorer for Oracle

Each Explorer has a corresponding user guide available.

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 and 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 and Replication can help you to restore specific VM files (.vmdk, .vmx and others) if any of these files are deleted or the data store 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 and 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 and 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 and 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 and 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 and 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 and 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 and Replication supports a number of replication scenarios that depend on the location of the target host and the data transport path.

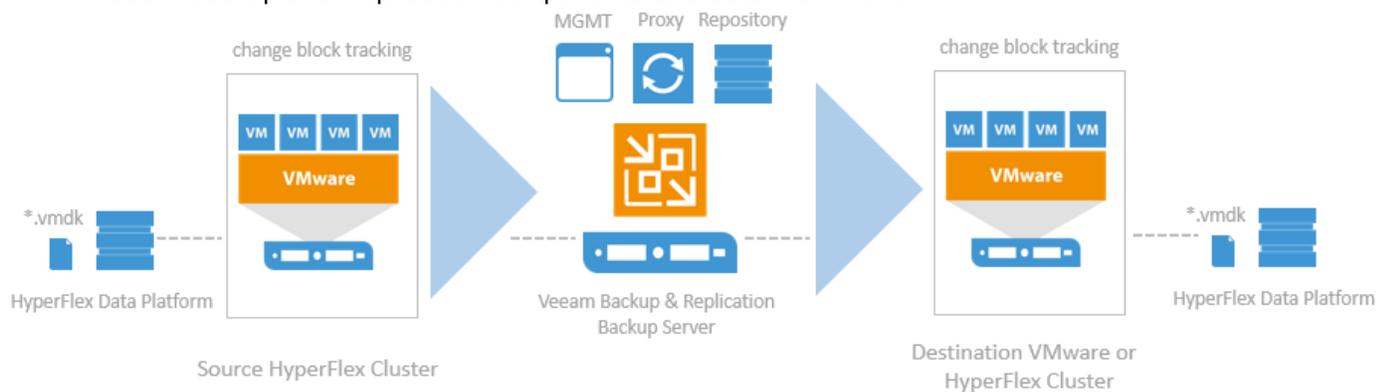
Onsite Replication

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 and Replication components and data movement



In the onsite replication scenario, Veeam Backup and 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

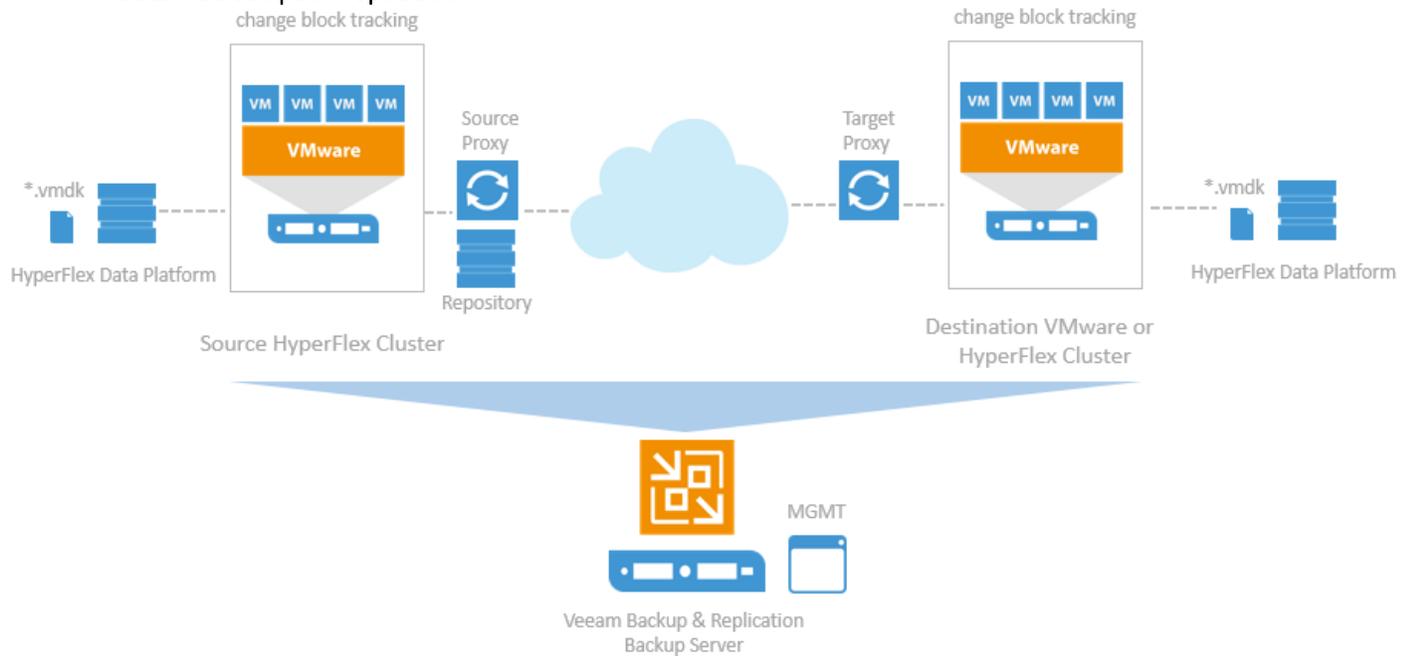
If the source host is located in the primary site and the target host is located in the DR site, you can perform offsite replication.

Offsite replication can run over two data paths:

- Direct data path
- Via a pair of WAN accelerators

Direct data path

Figure 13 Veeam Direct path replication



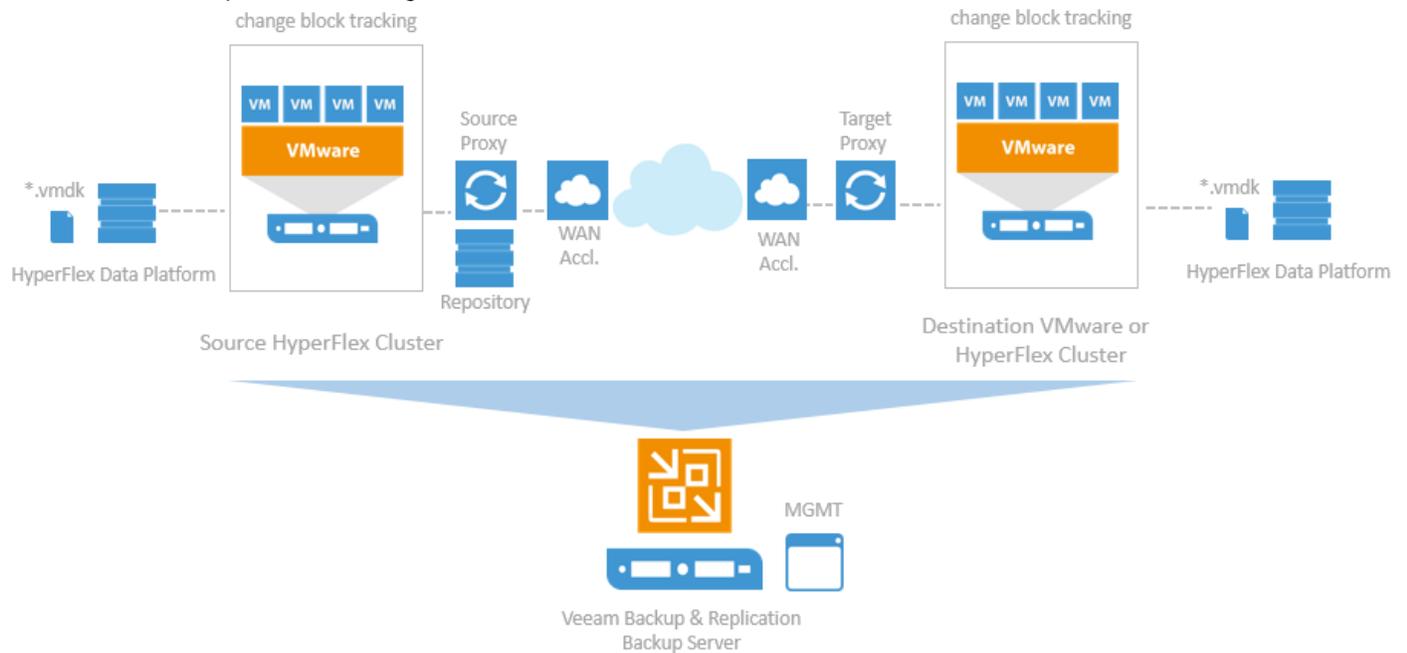
The common requirement for offsite replication is that one Data Mover Service runs in the production site, closer to the source host, and another Data Mover Service runs in the remote DR site, closer to the target host. During backup, the Data Mover Services maintain a stable connection, which allows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

Via WAN accelerators

If you have a weak 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

Figure 14 Veeam Replication through WAN accelerators



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, destines VM data to the target host.

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 and Replication, failover is a temporary intermediate step that should be further finalized. Veeam Backup and 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 and Replication supports failover and failback operations for one VM and for several VMs. In case one or several hosts fail, you can use batch processing 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 and 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.

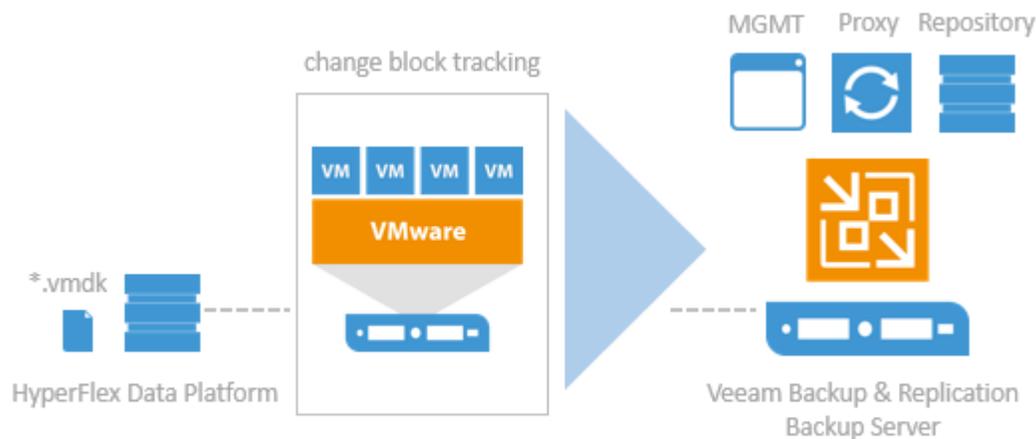
Backup Server

Components

Veeam Availability Suite combines the backup, restore and replication capabilities of Veeam Backup and **Replication™ with the advanced monitoring, reporting and capacity planning functionality of Veeam ONE™**. Veeam Availability Suite delivers everything you need to reliably ensure and manage your Cisco HyperFlex VMware environment. Veeam Backup and 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 and 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 and Replication components on the same machine, physical or virtual, or you can set them up separately for a more scalable approach.

The following drawing shows an overview on the main Veeam components:

Figure 15 Veeam Backup and Replication components



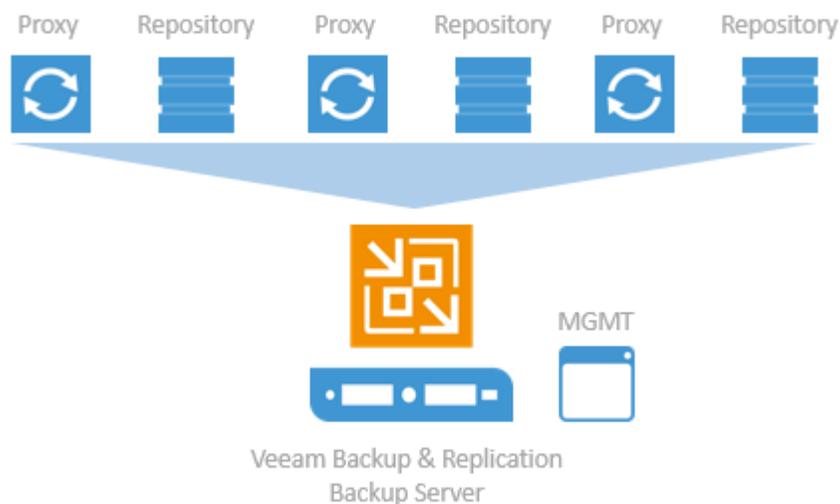
Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup and 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

It 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:

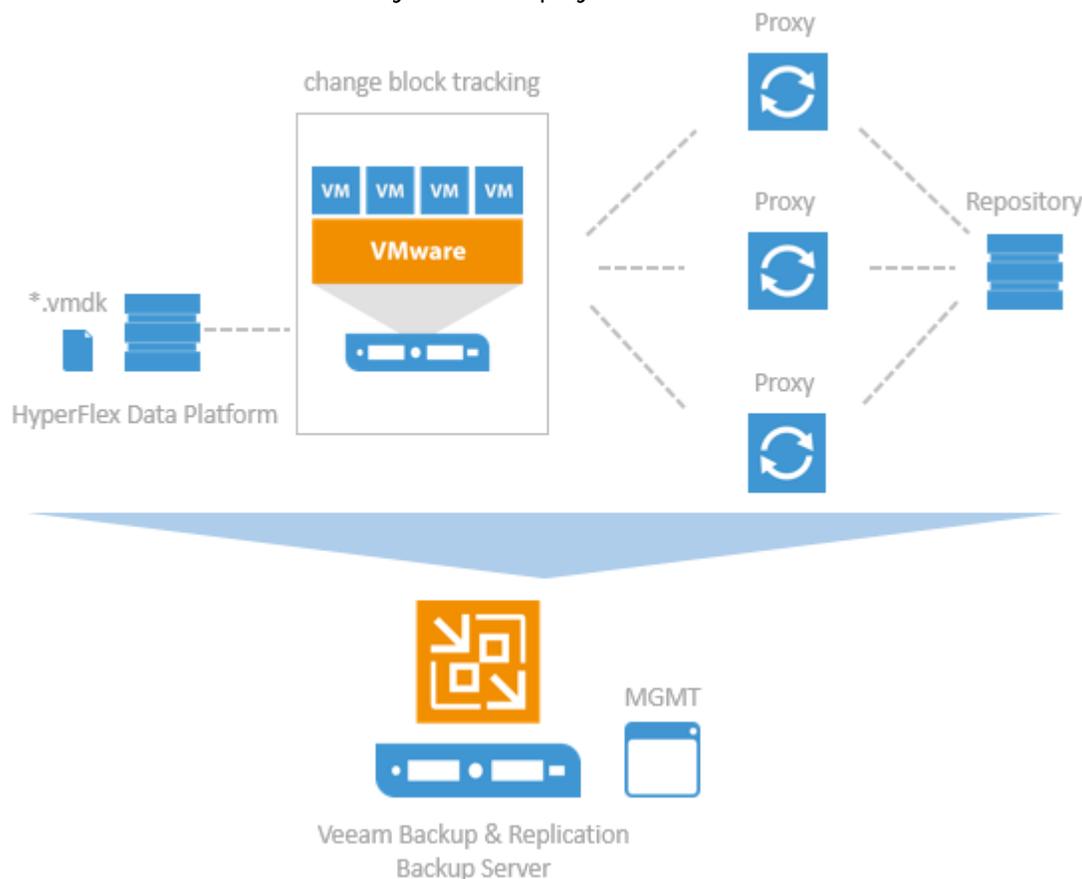
- Veeam Backup Service is a Windows service that coordinates all operations performed by Veeam Backup and Replication such as backup, replication, recovery verification and restore tasks. The Veeam Backup Service runs under the Local System account or account that has the Local Administrator permissions 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 and 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 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 and 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 light-weight 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 and 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 and 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 and 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.

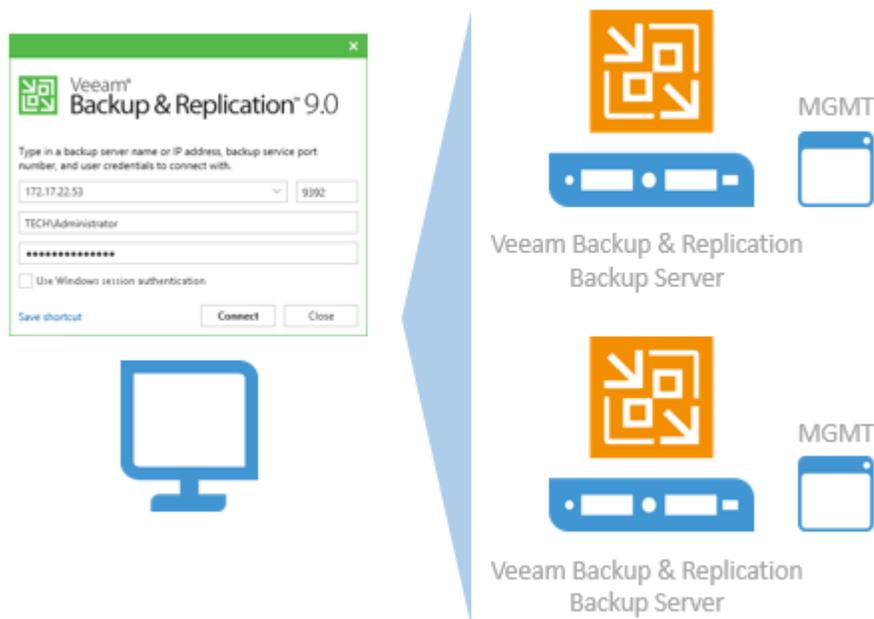
You can configure one of the following types of backup repositories:

- 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 performs a backup proxy that is used by the job for data transport.
- Deduplicating storage appliance. Veeam Backup and Replication supports different deduplicating storage appliances.

Backup and Replication Console

The Veeam Backup and 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 and Replication remotely over the network. The console lets you log in to Veeam Backup and Replication and perform all kind of data protection and disaster recovery operations as if you work on the backup server.

Figure 18 Veeam Backup and Replication Console



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

Backup Proxy

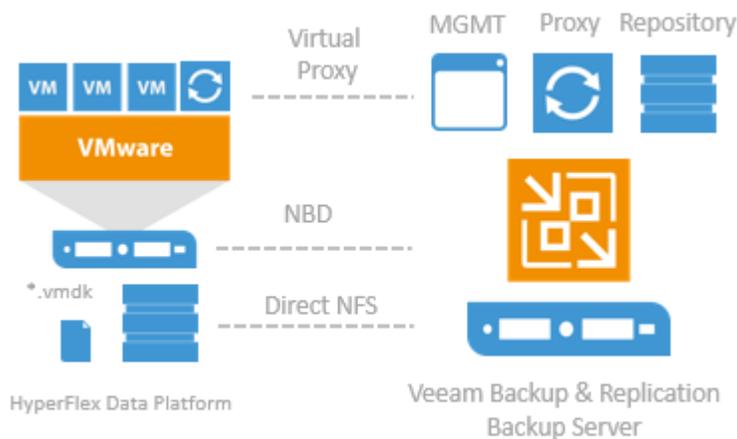
Transport Modes

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 and Replication offers the following modes:

- Direct storage access
- Virtual appliance
- Network (NBD)

Figure 19 Veeam Backup and Replication transport modes



In the Direct storage access mode, Veeam Backup and Replication reads/writes data directly from/to the storage system where VM data or backups are located. With the Direct NFS access mode, Veeam Backup and Replication bypasses the ESX(i) host and reads/writes data directly from/to NFS data stores. To do this, Veeam Backup and 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 and 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 data store, 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: - Total size of VMs being backed up - Frequency of backups - Retention period for backups - Will jobs use forward or reverse incremental

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
- Instance VM recovery of failed VM in less than two minutes
- Multiple backup end points such as tape drives, on cloud or on local repository

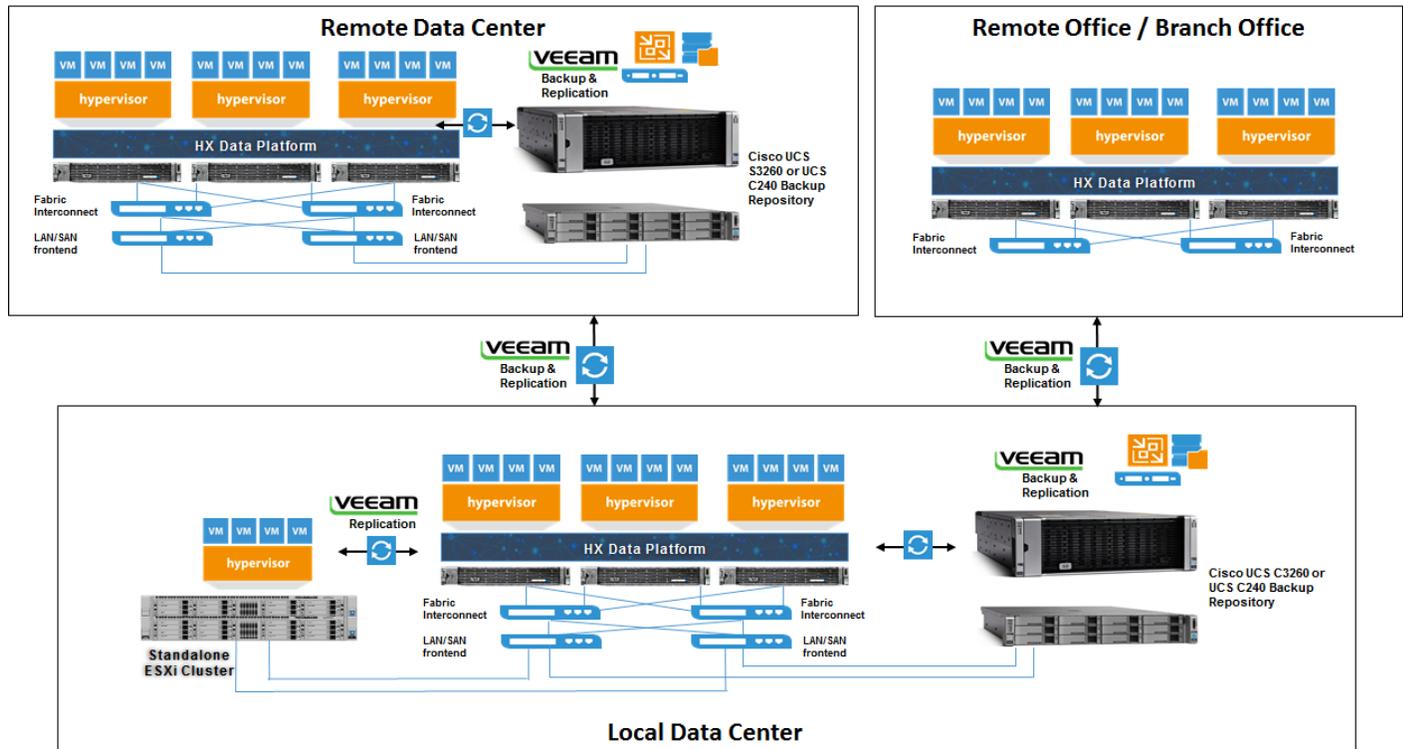
This section elaborates on the deployment architecture and design considerations to protect application data through Veeam Availability Suite. The application VMs can reside, across multiple HyperFlex Clusters within same Data Center or HyperFlex Clusters deployed across Data Centers such as Remote Office Branch Office (ROBO) deployments.

The key deployment scenarios to protect Cisco HyperFlex cluster with Veeam Availability Suite are listed as below

- Cisco HyperFlex Single Site Backup and Replication
- Cisco HyperFlex Remote office - Branch Office Replication
- Cisco HyperFlex multi-site Backup and Replication

The end to end deployment scenarios for Cisco HyperFlex with Veeam Availability Suite is detailed in the figure below.

Figure 20 Deployment Overview: Cisco HyperFlex with Veeam Availability Suite



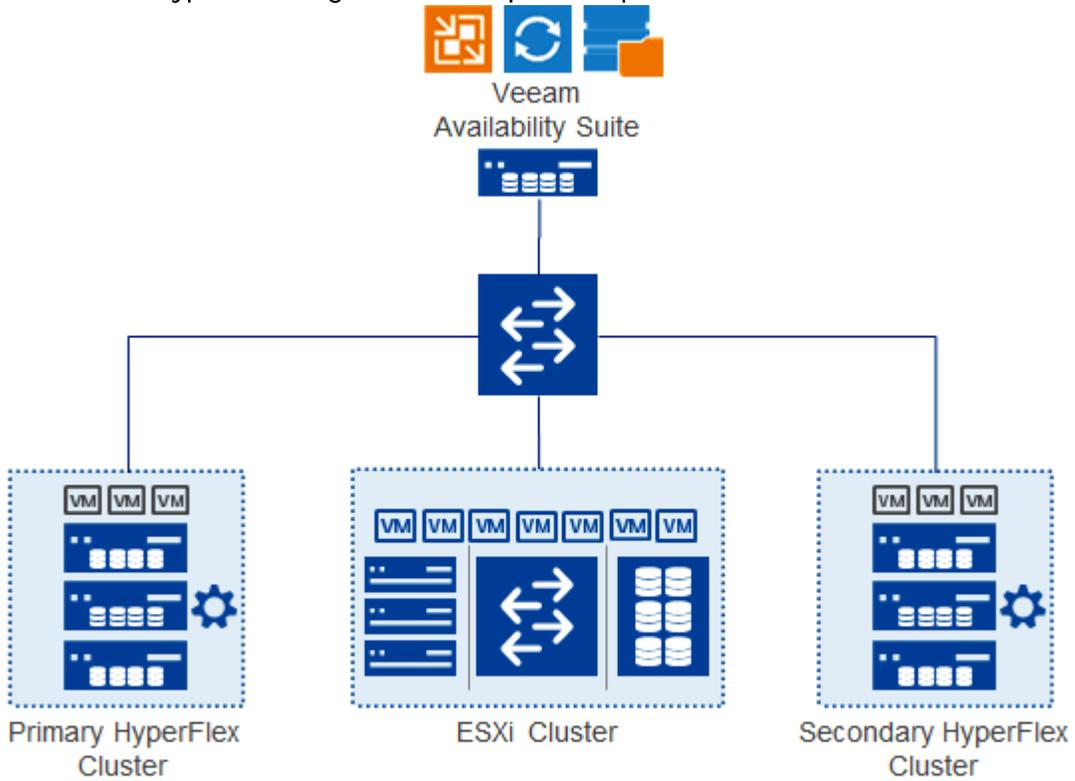
Cisco HyperFlex Single Site Backup and Replication

Single Site Backup and Replication for HyperFlex provides protection of application, data and VMs within the same deployment site or Data Center. The key features of Single Site Backup and Replication are:

- Veeam Availability Suite which includes Veeam Backup Server, Veeam Proxy and Veeam Repository resides on either Cisco UCS S3260 or Cisco UCS C240 M4 server. The choice for either of the Rack Server is dependent on network bandwidth, compute and storage resources required for Backup repository.
- 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 Veeam Proxy Server
- Backup of secondary HyperFlex Cluster VMs located in the same Data Center or campus to common Veeam Repository
- Replication of Primary HyperFlex Cluster VMs to Secondary HyperFlex cluster or vice-versa, with common Veeam Proxy

The figure below, elaborates on the use case for Cisco Hyper Flex Single Site Backup and Replication.

Figure 21 Cisco HyperFlex Single Site Backup and Replication

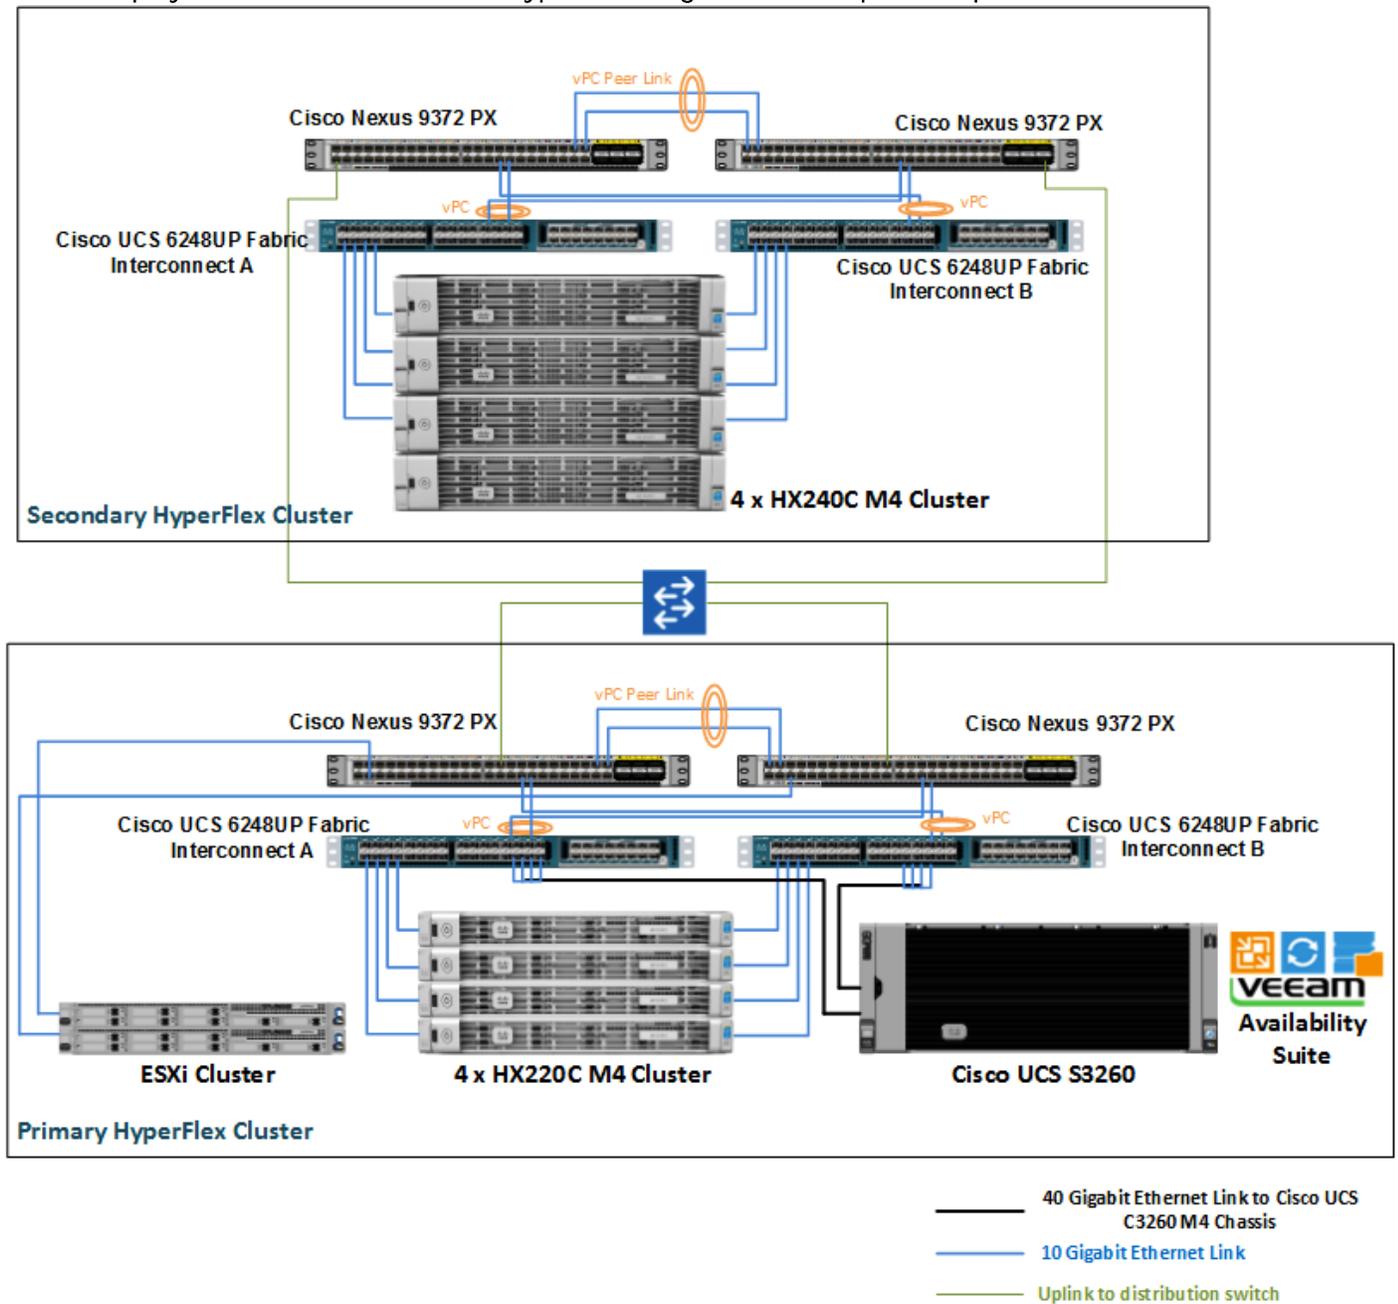


Deployment Architecture

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 elaborates on the reference architecture for Veeam AS with Cisco HyperFlex deployed in same campus or Data Center. The figure below 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 update 2 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)

- Cisco HX Cluster with Cisco HX Data Platform 1.8a and VMWare 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 UCS S3260 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 VMWare 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 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

Remote office - Branch Office Replication for Cisco HyperFlex

Several organizations today have Remote office and Branch offices (ROBO) spread across geographies, which provide localized data availability and allow businesses to execute critical workloads locally. ROBO deployments typically require fewer compute and storage resources with just few servers running workloads to support local needs. Organizations have several ROBO deployments spread across regions and a major challenge faced by these deployments is provisioning of application availability, deployed remotely.

The present design overcomes these challenges, by providing Replication of application VMs deployment in Remote Offices through Veeam Availability Suite. Application VMs in ROBO deployments are replicated to the primary Data Center and hence provide Failover and Fail Back at all times. This requires minimal infrastructure requirements and the replication is executed by Veeam Proxy installed just in a Virtual Machine. Moreover, it allows ensuring remote sites are in compliance, and reducing IT management time at remote offices.

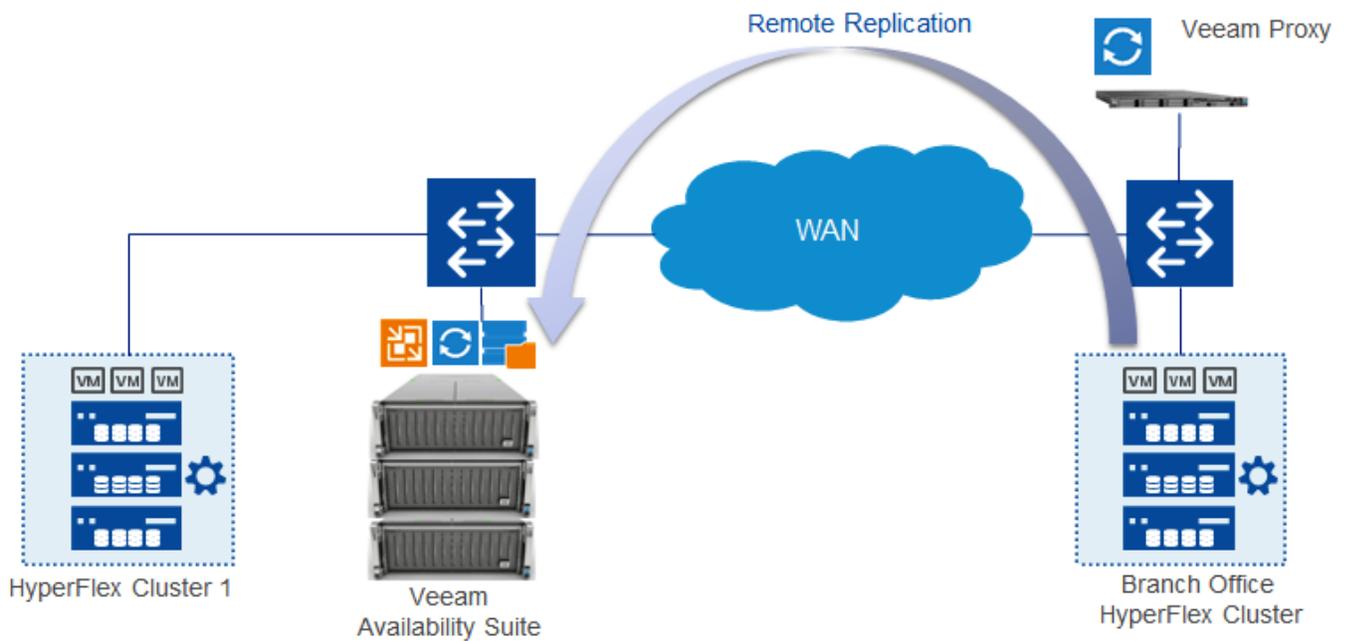
The key deployment features of ROBO Replication for Cisco HyperFlex are:

- Veeam Application Server, Veeam Proxy and Veeam Repository reside on either Cisco UCS S3260 or Cisco UCS C240 M4 server located in the Primary Data Center.
- Replication and Backup of HyperFlex Cluster VMs in Primary Data Center through Veeam Proxy and Veeam Backup Server.

- Replication of application VMs on Cisco HyperFlex Cluster located in Remote Office is executed on Veeam Availability Suite located in Primary Data Center.
- A Veeam Proxy Server is installed in the Remote Office. This could be installed either on a bare metal server or on a Virtual Machine. In present design, Veeam Proxy Server is installed in a Cisco UCS C220 M4 Server located in Remote Office. The choice of either a virtual machine or physical server for Veeam Proxy is dependent on several factors such as:
 - Number of Replication jobs executed on ROBO deployment.
 - Deployment of Veeam WAN Accelerator on the Remote Site. Customers can deploy Veeam WAN Accelerator, which allows faster replication and backup of application VMs. It is required to deploy Veeam WAN Accelerator on a Cisco UCS C220 M4 with SSDs for WAN Accelerator Cache.

The figure below, elaborates on the use case for Replication of application VM on ROBO site to the Primary Data Centre.

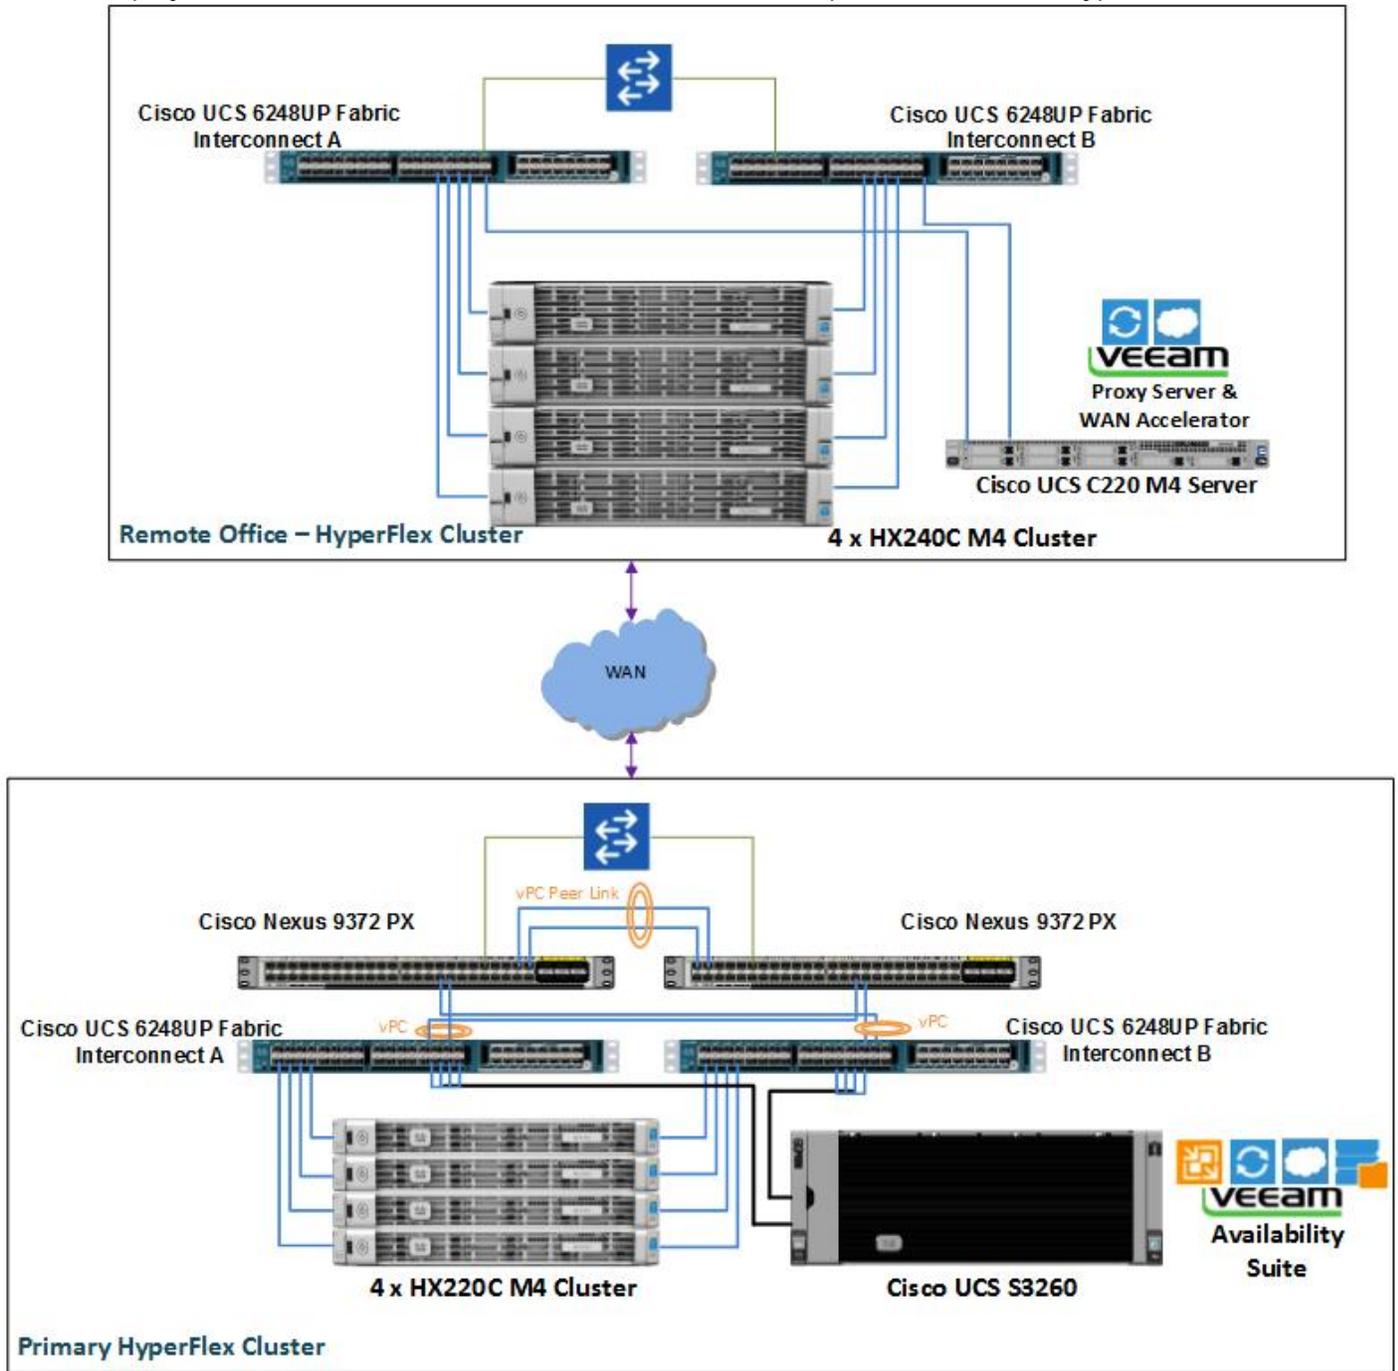
Figure 23 Remote office - Branch Office Replication for Cisco HyperFlex



Deployment Architecture

This section elaborates on the reference architecture for Veeam AS with Cisco HyperFlex existing in Remote Office locations. The figure below details the physical topology of the present deployment.

Figure 24 Deployment Architecture: Remote Office Branch Office Replication for Cisco HyperFlex



The solution detailed in the figure above includes Primary HyperFlex Cluster and Hyper Flex Cluster deployed in Remote Office. The application VMs are replicated through Veeam Availability Suite on Cisco UCS S3260 Storage server located in the Primary Data Center. The details on the deployment are as follows

- Primary HyperFlex Cluster

- Veeam Application Suite 9 update 2 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)
- Cisco HX Cluster with Cisco HX Data Platform 1.8a and VMWare Hypervisor 6.0 update 2. This is the primary HX Cluster wherein 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 UCS S3260 repository through Veeam AS
- Cisco Nexus 9300 switches
- Remote Office Branch Office deployment (ROBO)
- Cisco HX Cluster with Cisco HX Data Platform 1.8a and VMWare Hypervisor 6.0 update 2. This cluster provisions application VMs deployed in Remote Office location.
- Replication of the application VM on Remote Office are created on HX Cluster located in Primary Data Center
- Cisco UCS 6200 Series Fabric Interconnects
- Veeam Proxy Server and Veeam WAN Accelerator are deployed on Cisco UCS C220 M4 Rack Server. As mentioned, Veeam Proxy Server can be deployed on a VM, but it is recommended to have Veeam Proxy Server deployed on a bare metal server

Multi-Site Backup and Replication for Cisco HyperFlex

Multi-Site Backup and Replication comprises of distributed deployment scenarios for large geographically dispersed virtual environments with multiple backup servers installed across different sites. These backup servers can be federated under Veeam Backup Enterprise Manager; an optional component that provides centralized management and reporting for these servers through a web interface. Veeam Availability Suite for large multi-site Data Centers allows customers to more efficiently move large amount of data across the WAN network.

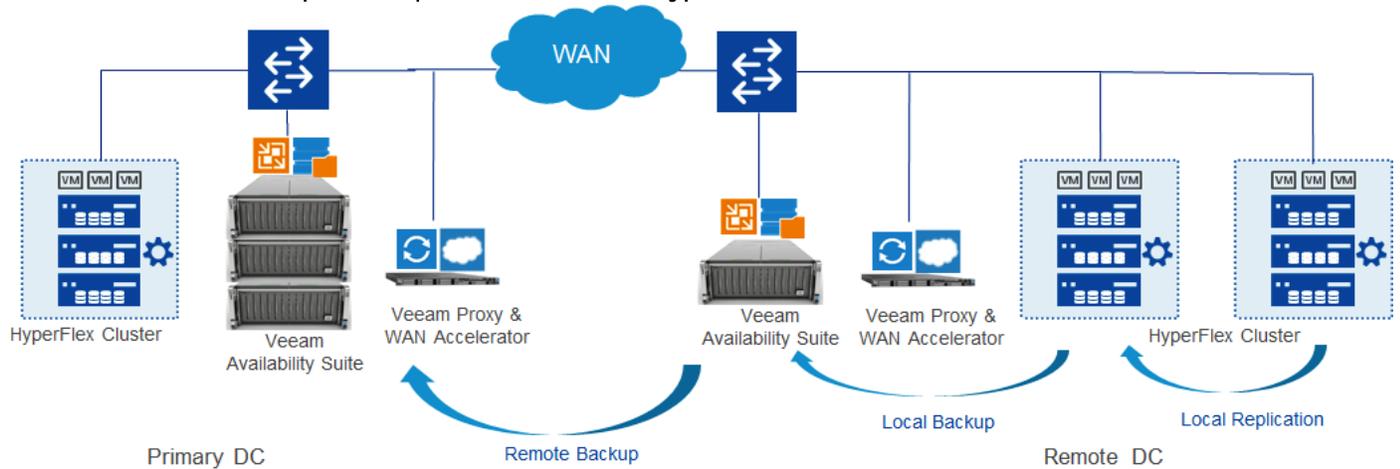
The key deployment features for Multi-Site Backup and Replication for Cisco HyperFlex are:

- Application VMs residing either on the Primary Data Center or on Remote Data Centers are backed up and replicated through Veeam Availability Suite deployed on Cisco S3260 Storage Server
- The distributed deployment environment is spread across geographies and can be managed through Veeam Backup Enterprise Manager – an optional component that provides centralized management and reporting for distributed Veeam Backup servers through a web interface
- The backup repository is replicated from remote Data Center repository to Primary Data Center repository through Veeam Availability Suite
- Application VMs can be replicated either to HyperFlex cluster within the same Data Center or to HyperFlex cluster in the Primary Data Center.

- Veeam WAN Accelerators are deployed in Primary and Remote sites, which reduce the amount of data that needs to flow back and forth across the WAN by using caching and data compression techniques. This allows reduction in the bandwidth required for transferring backups and replicas over the WAN. Built-in WAN Acceleration dramatically reduces the bandwidth required for transferring backups and replicas over the WAN. A WAN accelerator reduces the amount of data that needs to flow back and forth across by caching duplicate files (or parts of files) so they can be referenced in global cache instead of having to be sent across the WAN again

The figure below elaborates on the key uses cases for Backup and Replication for application VMs residing in multi-site HyperFlex deployment

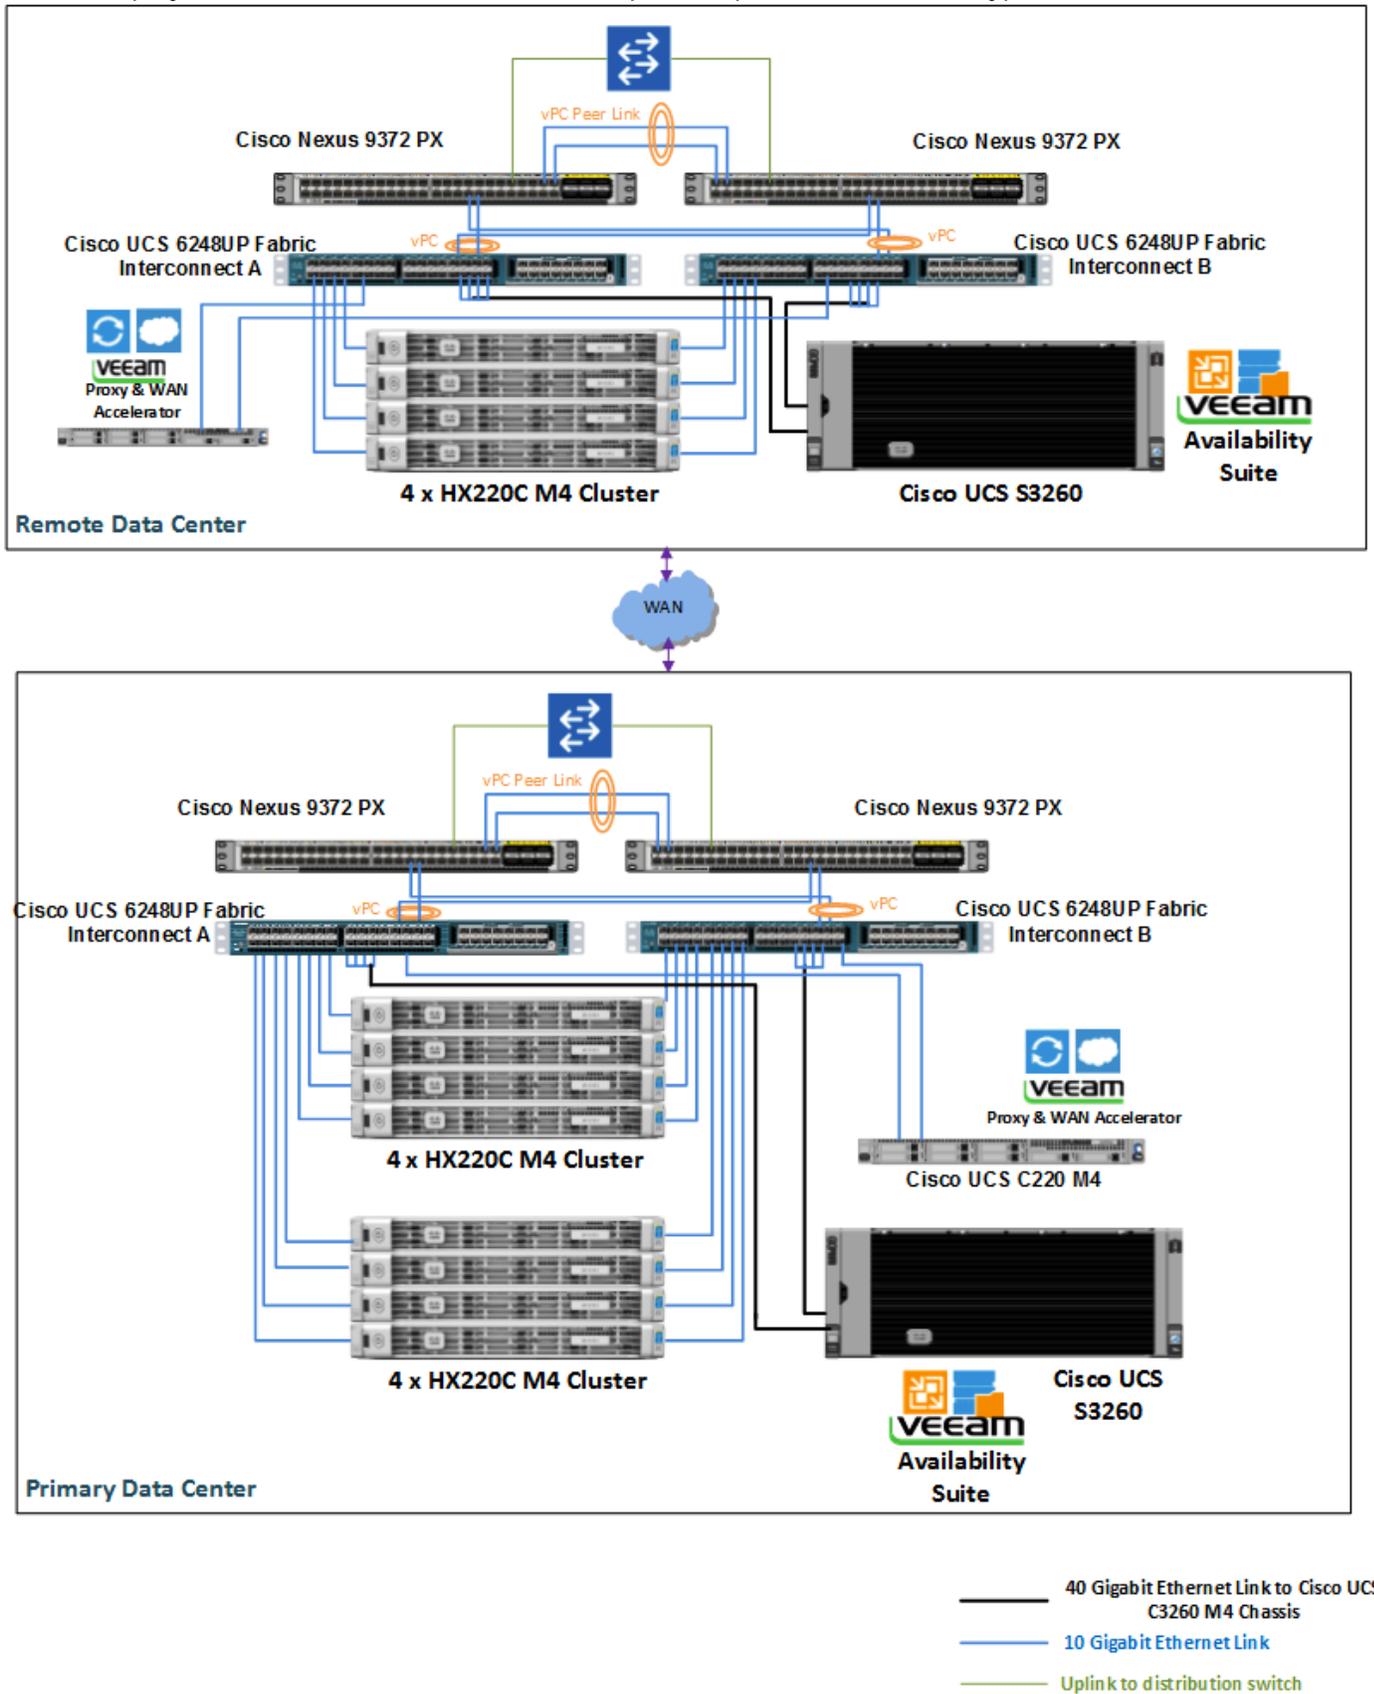
Figure 25 Multi-Site Backup and Replication for Cisco HyperFlex



Deployment Architecture

This section elaborates on the reference architecture for Veeam AS with in a multi-site Cisco Hyper Flex deployment. The figure below details the reference architecture for the same.

Figure 26 Deployment Architecture: Multi-site backup and replication for Cisco HyperFlex



The solution detailed in the figure above includes Primary and Remote Data Centers running virtualized applications on either Converged Infrastructure or Cisco HyperFlex Clusters. This deployment supports a combination of Converged and Cisco Hyper Flex Data Clusters with replication and backup on Cisco UCS S3260 Storage Server.

The infrastructure components remain same for both Primary and Remote Data Center. The details on the key infrastructure components deployed to validate this setup are as below.

- Veeam Availability Suite 9 update 2 is deployed on Cisco UCS S3260 Storage Server. This includes Veeam Repository as the primary repository for all backup jobs across distributed data centers. Veeam Repository and Veeam Backup Server are installed on Cisco UCS S3260 servers.
- Veeam Proxy Server and Veeam WAN Accelerator are deployed on a separate Cisco UCS C220 M4 server.
- Cisco HX Cluster with Cisco HX Data Platform 1.8a and VMWare 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 UCS S3260 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

Design Consideration

Unified Management

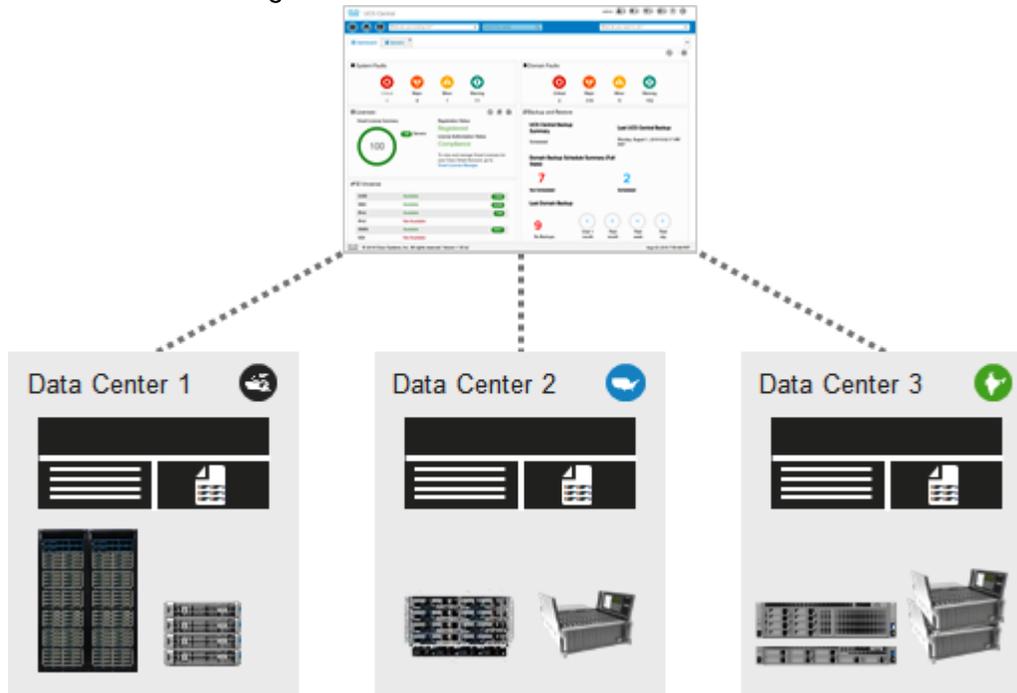
Cisco UCS Manager can manage B-series blade servers, C-series rack servers and S3260 Storage Servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware agnostic.

Moreover, Cisco UCS Central Software extends the policy-based functions and concepts of Cisco UCS Manager across multiple Cisco Unified Computing System (Cisco UCS) domains in one or more physical locations. This allows hardware configuration from a single UCS Central window, across multiple UCS domains, either in same Data Center or across Data Centers.

Cisco Unified Computing System™ (Cisco UCS) management helps significantly reduce management and administration expenses by automating routine tasks to increase operational agility. Cisco UCS management provides enhanced storage management functions for the Cisco UCS S3260 and all Cisco UCS servers. Cisco UCS Manager supports Storage profiles give you flexibility in defining the number of storage disks and the roles and uses of these disks and other storage parameters.

Figure below, elaborates on the single management window through UCS Central.

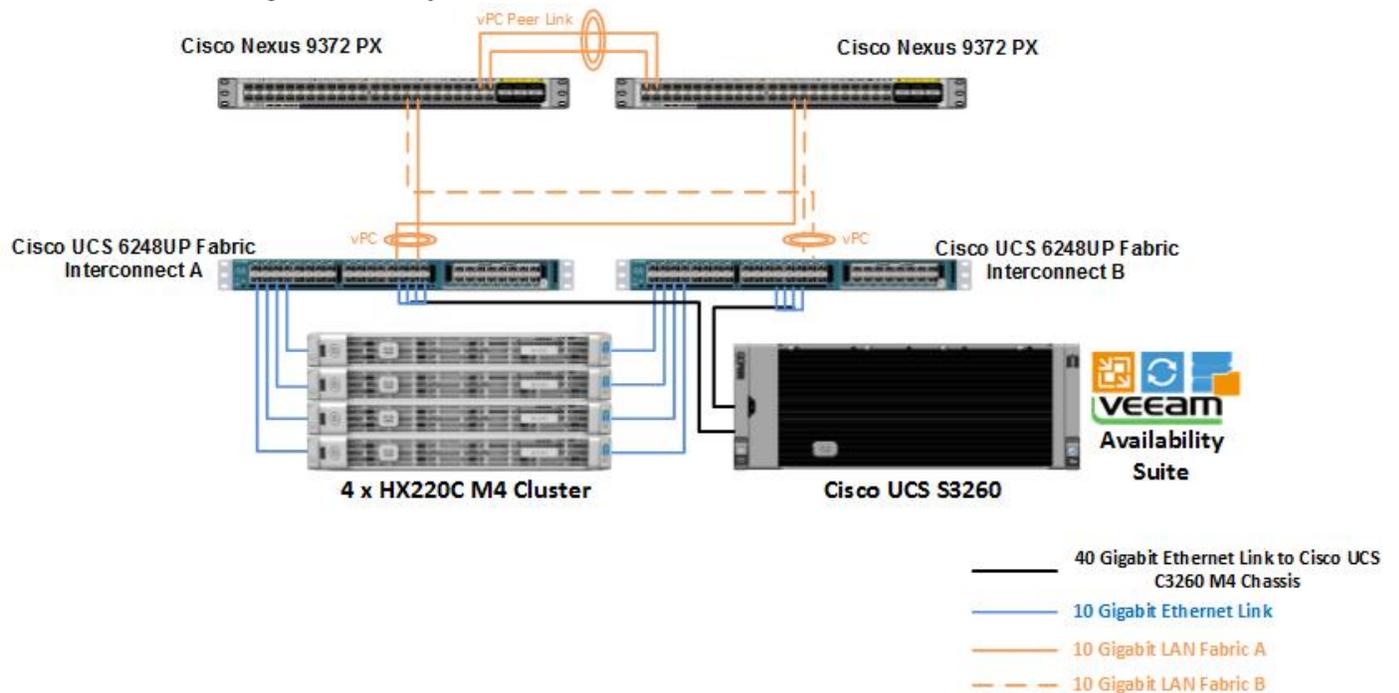
Figure 27 Unified Management across Data Centers



Network

The network design is shown in the figure below.

Figure 28 Network Design for Primary Data Center



The LAN network provides network reachability to the applications hosted on Cisco UCS servers in the data center. The infrastructure consists of a pair of Cisco Nexus 9372 PX switches deployed in NX-OS standalone mode. Redundant 10Gbps links from each Cisco Nexus switch are connected to ports on each FI and provide 20Gbps of bandwidth through each Cisco Nexus. Virtual PortChannels (vPCs) are used on the Cisco

Nexus links going to each FI. Jumbo Frames are also enabled in the LAN network to support backup and replication of application VMs on Veeam Server.

The design also uses the following best practices:

- Jumbo frames on unified fabric links between Cisco UCS and fabric interconnects
- QoS policy for traffic prioritization on the unified fabric
- Port-channels with multiple links are used in the unified fabric for higher aggregate bandwidth and redundancy
- A HyperFlex snapshot should be taken on VMs before creating a Veeam Backup and Replication job

The Veeam AS is deployed on Cisco UCS 3260 Storage Server and provides an aggregated bandwidth of 80 Gbps for VM backup and Replication. The present deployment supports Veeam AS on Cisco UCS C240 M4 server connected to common pair of Fabric interconnects and provides network throughput of up to 40 Gbps.

Scale

In the present design for Single Site Backup/Replication and Backup/Replication for ROBO deployments, Veeam Proxy server exists on the same server as Veeam Backup Server and Repository. This leads to limited compute resources for Backup and Replica jobs executed through Veeam Proxy and restricts scaling of parallel Backup and Replica jobs. Thus, to scale parallel Backup and Replica jobs, it is recommended to distribute Veeam Proxy on multiple compute servers. Moreover, with UCS unified management, Veeam Proxy distributed on several UCS Rack server can be easily managed through single UCS Manager window. In the event Veeam Proxy is distributed across multiple UCS Fabric Interconnects, customers can utilize Cisco UCS Central, which enabled global and local management of Cisco UCS domains to promote consistency and standardization across domain.

In the present design for multi-site or multi Data Center Backup and Replication with Veeam, Veeam Proxy is installed on a separate Cisco UCS C220 M4, which is directly attached to same Fabric Interconnect as the HyperFlex cluster, or Veeam Repository deployed on Cisco UCS S3260 server.

WAN Acceleration for Backup and Replica

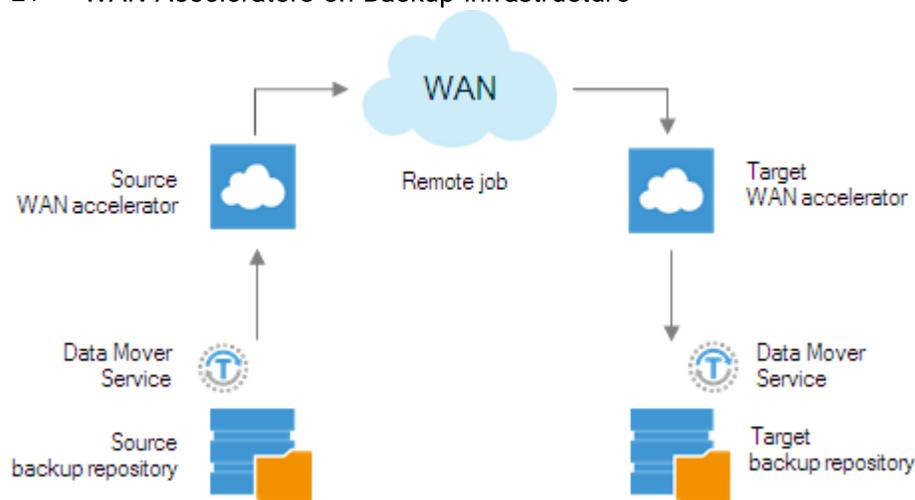
Remote Site Backup and Replication always involves moving large volumes of data between remote sites. The most common problems that backup administrators encounter during Remote Site Backup and Replication are:

- Insufficient network bandwidth to support VM data traffic
- Transmission of redundant data

Veeam Backup and Replication offers the WAN acceleration technology that helps optimize data transfer over WAN. Built-in WAN Acceleration utilizes global caching, variable block length data fingerprinting and traffic compression to significantly reduce bandwidth requirements, while multiple WAN optimization ensures that available bandwidth is leveraged to its fullest potential.

Figure below elaborates on the WAN Acceleration for faster Backup and Replication jobs across distributed Data Centers.

Figure 29 WAN Accelerators on Backup Infrastructure



In the present design, Veeam WAN Acceleration is utilized for Remote Site Backup and Replication. It is recommended to have a caching layer such as SSD to create global cache across the Primary and Remote Site; WAN Accelerator is installed on a Cisco UCS C220 M4.

Veeam Backup Enterprise Manager Console

Veeam Enterprise Manager is intended for centralized reporting and management of multiple backup servers. It provides delegated restore and self-service capabilities as well as the ability for users to request Virtual Labs from backup administrators. It provides a central management point for multiple backup servers from a single interface. Enterprise Manager is also a part of the data encryption and decryption processes implemented in the Veeam solution. For best practices, Veeam recommends deploying Enterprise Manager in the following scenarios:

- It is recommended to deploy Enterprise Manager if you are using encryption for backup or backup copy jobs. If you have enabled password loss protection (http://helpcenter.veeam.com/backup/em/index.html?em_manage_keys.html) for the connected backup servers backup files will be encrypted with an additional private key which is unique for each instance of Enterprise Manager. This will allow Enterprise Manager administrators to unlock backup files using a challenge/response mechanism effectively acting as a Public Key Infrastructure (PKI).
- If an organization has Remote Office/Branch Office (ROBO) deployments, then leverage Enterprise Manager to provide site administrators with granular restore access via web UI (rather than providing access to the Backup and Replication console).
- In enterprise deployments, delegation capabilities can be used to elevate the first line support to perform in-place restores without administrative access.
- For deployments spanning multiple locations with stand-alone instances of Veeam Backup and Replication, Veeam Enterprise Manager will be helpful in managing licenses across these instances to ensure compliance.
- Enterprise Manager is required when automation is essential to delivering IT services – to provide access to the Veeam RESTful API.

Veeam Transport Mode for HX Cluster

The best practice is to utilize the Network Transport Mode on the Veeam Proxy, and to ensure that each VM in a Veeam Backup Job has an HX Snapshot for HyperFlex. HX Snapshots allow rapid creation and deletion of snapshots with negligible performance impact.

The HX Native Snapshot creation and deletion is offloaded from ESX through the HyperFlex VAAI plugin to provide quiesced, crash-consistent, snapshots that leverage the log structured filesystem in HyperFlex. This allows space efficient and nearly instant snapshots without performance impact. Ensure the VM has no snapshots in the vSphere Snapshot Manager and then create an HX snapshot. Right click the VM, select **Cisco HX Data Platform**, and choose either **'Schedule Snapshot'** or **'Snapshot Now'**. **VMs can also be added** to a Folder in the vSphere Web Client and then the Folder can be right clicked and a HyperFlex snapshot can be created or scheduled.

Validation

This section elaborates on the test executed to validate the Veeam Backup and Recovery Solution on Cisco HyperFlex platform. This solution and its scenarios are validated with high level backup, Replication, Failover and Failback task across HyperFlex clusters.

Some of the important tests executed are:

- Single Site Backup and 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
- Remote Office Branch Office Replication
 - Create Backup of application VM on Remote Office HX Cluster to Veeam Backup Repository in Primary Site
 - Restore File, Instant Recovery and Restore Entire VM to HX Cluster in Remote Office
 - Create Replica, Failback and Failover of application VM on Remote Office HX Cluster to Primary Site HX Cluster

Validated Hardware and Software

Table below summarizes the software and hardware components deployed to validate the design for Cisco HyperFlex with Veeam Backup and Replication.

Table 1 **Hardware Software Component List**

	Components	Software Version	Comments
Compute and Storage	Cisco UCS S3260 Storage Server	3.1(2b)	Directly managed through Fabric Interconnect. Veeam AS is installed on the same. Provides Storage Veeam Repository
	Cisco UCS C240 M4 Rack Server	3.1(2b)	Directly attached to Fabric Interconnect. Veeam AS is deployed in the Remote Data Center
	Cisco UCS C220 M4 Rack Server		Directly attached to Fabric Interconnect. Veeam Proxy and WAN Accelerator are installed on this server
	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	9.0 update 2	Pre-configured with Veeam Backup Server, Veeam Proxy , Veeam Repository
	Operating System	Windows 2012 R2	
Hyperconverged 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	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

Bill of Materials

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

Table 2 **Bill of Materials**

Line Number	Part Number	Description	Quantity
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-MHz RDIMM/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 UCS C220 M4 rack servers	2
1.14	HX220C-BZL-M4	HX220C M4 Security Bezel	1

Bill of Materials

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/16UP/ 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
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-MHz RDIMM/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	1

Bill of Materials

		Port 10Gb SFP+	
3.7	UCSC-PSU2V2-1400W	1400W V2 AC Power Supply (200 - 240V) 2U and 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+ 2PCI slts	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	UCSC-C240-M4L	UCS C240 M4 LFF 12 HD w/o CPU,mem,HD,PCIe,PS,railkit w/expdr	1
5.0.1	CON-PSJ7-C240M4L	UCS PSS 24X7X4 OS UCS C240 M4 LFF 12 HD w/o CPU,mem	1
5.1	UCS-CPU-E52650E	2.20 GHz E5-2650 v4/105W 12C/30MB Cache/DDR4 2400MHz	2
5.2	UCS-MR-1X081RV-A	8GB DDR4-2400-MHz RDIMM/PC4-19200/single rank/x4/1.2v	4
5.3	UCS-SD480GBKS4-EB	480 GB 2.5 inch Enterprise Value 6G SATA SSD (Boot)	2
5.4	UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	1
5.5	UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	1
5.6	UCSC-PSU2V2-1400W	1400W V2 AC Power Supply (200 - 240V) 2U and 4U C Series	2
5.7	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	2
5.8	UCSC-PCI-1C-240M4	Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	1
5.9	UCS-HD4T7KL12G	4 TB 12G SAS 7.2K RPM LFF HDD	12
5.1	UCSC-SCCBL240	Supercap cable 250mm	1
5.11	UCS-M4-V4-LBL	Cisco M4 - v4 CPU asset tab ID label (Auto-Expand)	1
5.12	UCSC-HS-C240M4	Heat sink for UCS C240 M4 rack servers	2
5.13	UCSC-MRAID12G	Cisco 12G SAS Modular Raid Controller	1
5.14	UCSC-MRAID12G-1GB	Cisco 12Gbps SAS 1GB FBWC Cache module (Raid 0/1/5/6)	1
5.15	C1UCS-OPT-OUT	Cisco ONE Data Center Compute Opt Out Option	1
6.0	UCSC-S3260	Cisco UCS S3260 Base Chassis w/4x PSU, SSD, Railkit	1
6.0.1	CON-PSJ7-S3260BSE	UCS PSS 24X7X4 OS, Cisco UCS S3260 Base Chassis w/4x PSU	1
6.1	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	4

Bill of Materials

6.2	N20-BKVM	KVM local IO cable for UCS servers console port	1
6.3	UCSC-C3X60-BLKP	Cisco UCS C3X60 Server Node blanking plate	1
6.4	UCSC-C3160-BEZEL	Cisco UCS C3160 System Bezel	1
6.5	UCSC-C3X60-RAIL	UCS C3X60 Rack Rails Kit	1
6.6	UCSC-C3X60-SBLKP	UCS C3x60 SIOC blanking plate	1
6.7	UCSC-PSU1-1050W	UCS C3X60 1050W Power Supply Unit	4
6.8	N20-BBLKD-7MM	UCS 7MM SSD Blank Filler	2
6.9	UCSC-C3K-M4SRB	UCS C3000 M4 Server Node for Intel E5-2600 v4	1
6.1	UCS-CPU-E52650E	2.20 GHz E5-2650 v4/105W 12C/30MB Cache/DDR4 2400MHz	2
6.11	UCS-MR-1X161RV-A	16GB DDR4-2400-MHz RDIMM/PC4-19200/single rank/x4/1.2v	8
6.12	UCS-C3K-M4RAID	Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	1
6.13	UCSC-HS-C3X60	Cisco UCS C3X60 Server Node CPU Heatsink	2
6.14	UCSC-S3260-SIOC	Cisco UCS S3260 System IO Controller with VIC 1300 incl.	1
6.15	UCS-C3K-56HD10	UCS C3X60 4 row of 10TB NL-SAS drives (56 Total) 560TB	1
6.16	UCSC-C3X60-10TB	UCSC C3X60 10TB 4Kn for Top-Load	56
6.17	UCS-C3X60-G2SD48	UCSC C3X60 480GB Boot SSD (Gen 2)	2

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>