Cisco HyperFlex with Cohesity Data Protection

Deployment and Configuration Guide for Cohesity Data Protection on Cisco C240 M5 LFF Servers for Protection of Cisco HyperFlex ESXi Clusters

Last Updated: January 11, 2019

About the Cisco Validated Design (CVD) Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information see:

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

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

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

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

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2019 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary

Across all industries, new applications, file formats and technologies are driving data growth, increasing infrastructure costs and complexity for both primary and secondary data. However, legacy point products cannot scale, are siloed, slow to deploy, are costly and hard to use. Cisco and Cohesity quickly and easily consolidate all data; both primary and secondary via integrated, hyperconverged platforms to provide greater visibility, management, protection and access in any cloud at any scale. No other combined solution offers the simplicity, visibility, agility and data access across clouds, while reducing costs, increasing efficiencies, and accelerating innovation for businesses worldwide.

Hyperconverged infrastructures (HCI) continue to gain a significant foothold across modern datacenters. The integrated combination of networking, storage, and compute resources leads to significant improvements in ease of management, rapid deployment, and lower total costs of ownership (TCO). Hyperconverged infrastructures provide a more streamlined and simplified architecture, one which offers new levels of flexibility to grow and adapt on demand to changing workloads and requirements, and consolidation of workloads into fewer and less distinct platforms.

Initially, hyperconverged infrastructures have been utilized as a primary workload and data platform. Cisco HyperFlex systems are based on the Cisco UCS platform, combining Cisco HX-Series x86 servers and integrated networking technologies through the Cisco UCS Fabric Interconnects, into a single management domain, along with industry leading virtualization hypervisor software from VMware, and next-generation software defined storage technology. The combination creates a complete modernized virtualization platform, which provides the network connectivity for the guest virtual machine (VM) connections, and the distributed storage to house the VMs, spread across all of the Cisco UCS x86 servers without using specialized storage or networking components. However, more than 80 percent of an organization's enterprise data is considered secondary, or non-mission critical. Backups, archives, test/dev, files, objects, as well as analytics are included in this category. The design principles that make Cisco HyperFlex a leading product in the hyperconverged infrastructure market as a primary storage system can also be applied to systems that target secondary storage workloads.

Enterprises today struggle with the growing problem of mass data fragmentation created by a legacy approach to secondary data and applications. This approach relies on a complex patchwork of point products to try and manage these secondary workflows. According to a recent global market study by Vanson Bourne, more than a third of organizations use six or more solutions for all of their secondary data operations and of that, 10 percent use 11 or more solutions to manage this complex web of infrastructure silos. For example, typical enterprise class secondary storage solutions rely on multiple components, such as backup agents, media servers, controller servers, on-disk storage servers, and other storage media systems such as tape libraries, or off-site cloud repositories. This leads to significant proliferation within the secondary system, beyond the sprawl which may already exist across the primary systems they are protecting. Cohesity offers a unique platform based on the principles of hyperconvergence, utilizing the power and flexibility of the Cisco Unified Computing System (UCS) architecture, and Software Defined Storage, to create a unified, web-scale secondary data and applications solution. This approach creates a consolidated secondary storage platform that avoids the common problems that plague traditional deployments. By leveraging the benefits of hyperconvergence offered by Cisco HyperFlex and Cohesity, the primary and secondary systems can be collapsed into a singular architecture within Cisco UCS, which provides both the primary workload storage, the data protection of those workloads, file/object services to those workloads, and more productive secondary data and applications.

This Cisco Validated Design and Deployment Guide provides descriptions and instruction for the design, setup, configuration and ongoing use of the Cohesity DataPlatform, protecting Cisco HyperFlex clusters, both operating within a Cisco UCS domain. This unique integrated solution is designed to solve the infrastructure, operational, and data management challenges, while reducing the fragmentation commonly seen across secondary storage silos, within many enterprise data centers. The best-of-breed solution combines the web-scale simplicity and

efficiency of Cohesity software with the power and flexibility of Cisco UCS servers. As a result, customers can more efficiently and effectively manage unstructured data growth, acquire new insights, and reduce costs and complexity with a single, integrated solution. For more information about the Cisco-Cohesity integrated solution, please see https://www.cohesity.com/products/cisco.

Solution Overview

Introduction

The Cisco HyperFlex system combines the industry-leading convergence of computing and networking provided by Cisco UCS, along with next-generation hyperconverged storage software, to uniquely provide the compute resources, network connectivity, storage, and hypervisor platform to run an entire virtual environment; all contained in a single uniform system. Some key advantages of hyperconverged infrastructures are the simplification of deployment, day to day management operations, as well as increased agility, thereby reducing operational costs. Since hyperconverged storage can be easily managed by an IT generalist, this can also reduce technical debt going forward that is often accrued by implementing complex systems that need dedicated management teams and skillsets. The Cisco HyperFlex HX Data Platform is a purpose-built, distributed log-based file system, delivering high-performance, along with many data management and optimization features required in enterprise-class storage systems. This platform offers independent scaling of storage and computing resources, continuous data optimization through in-line compression and deduplication, dynamic data distribution for increased data availability, plus integrated native snapshots, rapid cloning, encryption, and VM level replication. This agile system is quick to deploy, easy to manage, is scalable and flexible to adapt to changing workloads, and provides high levels of data security and availability.

Cohesity offers a revolutionary hyperconverged secondary data and applications platform, which empowers enterprises to consolidate and simplify their secondary data architectures, into a more concise and easier to use system. Typical secondary storage use cases, such as backup, recovery, file services, test/dev, and analytics environments, can be collapsed into a single cohesive system, managed via a web-based interface, and using policy driven controls. The unique hyperconverged secondary data platform offers a highly resilient, web-scale secondary data solution that reduces complexity, operational challenges, and storage inefficiencies.

Cohesity running alongside Cisco HyperFlex within a Cisco UCS domain, offers a consolidated system that provides the primary storage, workload hosting, data protection, and file services required for most virtualized datacenters, all within a single unified architecture. Cohesity and Cisco HyperFlex share complementary datacenter technologies, with both of them utilizing a distributed filesystem architecture that is designed for high availability. Through a shared-nothing topology, there is no single point of failure or inherent bottlenecks, therefore both performance and capacity can scale linearly as more physical nodes are added to the clusters. The distributed file system spans across all nodes in the cluster and natively provides global deduplication, compression and encryption. Both systems are deployed on Cisco UCS x86 rack mount server hardware, connected to and managed by Cisco UCS Fabric Interconnects, which offer stateless, policy-based, programmatic control of the server configurations.

An additional benefit of the Cohesity system is its independence from the hardware platform which underlies the VMware virtualization environment. While this document targets the use case of Cohesity to protect Cisco HyperFlex with VMware ESXi hypervisors, in reality the Cohesity system deployed according to the guidelines in this document can also be used to protect multiple disparate VMware virtualized environments, plus bare-metal servers, database clusters, network attached storage (NAS) volumes, and offer file services via NFS, SMB and S3 object storage to the entire enterprise.

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, IT engineers, partners, and customers who are interested in learning about and deploying the Cohesity DataPlatform for secondary storage use cases, such as backup & restore, replication, archiving, plus file services with NFS, SMB/CIFS and S3 backed object storage.

Purpose of this Document

This document describes the installation, configuration and use of the Cohesity DataPlatform and Cohesity DataProtect, protecting VMware ESXi based Cisco HyperFlex systems, and providing file services to the environment. A reference architecture is provided for the combination of the two systems operating within a single Cisco UCS domain. The document does not specifically cover the design, installation and configuration of the Cisco HyperFlex system, VMware ESXi, or VMware vCenter, as these are covered in other Cisco Validated Design documents which are previously published. As such, it is recommended that a Cisco UCS environment that would contain both Cisco HyperFlex and Cohesity be initially set up to host the Cisco HyperFlex system. The installation of Cisco HyperFlex can require reboots of some components, therefore the subsequent configuration of the Cohesity system is more easily done as a secondary task.

What's New in this Release?

Unlike VMware snapshots (which uses redo log technology), the Cisco HyperFlex API allows Cohesity to integrate their backup solution with HyperFlex native snapshots. Backup products typically use virtual machine (VM) snapshots as a basis for implementing backups. Snapshots are an inherent virtualization feature that saves versions (states) of active VMs and can be used in a few scenarios:

- As a local point in time capture that is revertible, as needed. Typically, this is used when a guest OS or application fails.

- To establish a consistent point in time for seeding a backup. By saving the changes made to the VM since the time of the last backup, a new snapshot is analogous to an incremental backup.

- For seeding virtual machine clones and Virtual Desktop solutions.

With Cisco HyperFlex integration, Cohesity DataProtect software takes VM snapshots directly on HyperFlex, which creates a storage-native snapshot for the VM. Since this snapshot is native to HyperFlex, it has very similar performance characteristics as that of the original base disk, when compared to the performance when using standard VMware redo-log based snapshots. After the snapshot is taken, Cohesity DataProtect proceeds to back up the VM data, and once complete the snapshot is deleted via HyperFlex API. Using native snapshots eliminates common delays and I/O penalties, and improves application performance by using the underlying HyperFlex distributed storage technology to create and consolidate the snapshots.

Solution Summary

The following are the components of a Cohesity cluster running in a Cisco UCS domain:

- One pair of Cisco UCS model 6454 Fabric Interconnects

- A minimum of three Cisco C240-M5L Rack-Mount Servers with the Cisco UCS VIC 1457 card

- Cohesity DataPlatform Software

Technology Overview

Cisco Unified Computing System

The Cisco Unified Computing System (Cisco UCS) 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, 25 Gigabit Ethernet, or 40 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, 25-Gbps, or 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are often 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. Uplink connections from the UCS fabric can be 10-Gbps, 40-Gbps or 100-Gbps connections.

- 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 SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying 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 their choice of storage protocol and physical architecture, and enhanced 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 UCS Fabric Interconnect

The Cisco UCS Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. Depending on the model chosen, the Cisco UCS Fabric Interconnect offers line-rate, low-latency, lossless 10 Gigabit, 25 Gigabit, 40 Gigabit, or 100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel connectivity. Cisco UCS Fabric Interconnects provide the management and communication backbone for the Cisco UCS C-Series, S-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 Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabrics, the Cisco UCS Fabric Interconnects provide both the LAN and SAN connectivity for all servers within its domain.

For networking performance, the Cisco UCS 6454 Series uses a cut-through architecture, supporting deterministic, low latency, line rate 10/25/40/100 Gigabit Ethernet ports, 3.82 Tbps of switching capacity, and 320 Gbps bandwidth per Cisco 5108 blade chassis when connected via the IOM 2208 model. The product family supports Cisco low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over the Ethernet fabric from the servers to the uplinks. 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 6454 Fabric Interconnect

The Cisco UCS 6454 Fabric Interconnect is a one-rack-unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has eight (8) 10/25-Gbps fixed Ethernet ports, which can optionally be configured as 8/16/32-Gbps FC ports (ports 1 to 8), thirty-six (36) 10/25-Gbps fixed Ethernet ports (ports 9 to 44), four (4) 1/10/25-Gbps Ethernet ports (ports 45 to 48), and finally six (6) 40/100-Gbps Ethernet uplink ports (ports 49 to 54).

Figure 1   Cisco UCS 6454 Fabric Interconnect

Cisco UCS C-Series Cohesity Nodes

A Cohesity cluster requires a minimum of three Cisco UCS C-Series "converged" nodes (with disk storage). 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. Each node is equipped with two high-performance SSD drives for data caching and rapid acknowledgment of write requests. Each node also is equipped with additional hard disks for long term storage and overall capacity.

Cisco UCS C240 M5 LFF Server

This two-rack-unit (2RU) Cisco C240 M5 Large Form Factor (LFF) model server contains a pair of 240 GB M.2 form factor solid-state disk (SSD) that acts as the boot drives, a pair of 1.6 TB or 3.2 TB NVMe SSD drives installed in the rear drive slots, and twelve 4 TB or 10 TB SATA HDD drives for storage capacity.

Figure 2   Cisco C240 M5 LFF Server

Cisco UCS VIC 1457 MLOM Interface Card

The Cisco UCS VIC 1457 Card is a quad-port Enhanced Small Form-Factor Pluggable (SFP+) 10/25-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS C-Series Rack Servers. The VIC 1457 is used in conjunction with the Cisco UCS 6454 model Fabric Interconnects. 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, providing investment protection for future feature releases. The card enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host, each dynamically configured as either a network interface card (NICs) or host bus adapter (HBA). The personality of the interfaces is set programmatically using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, adapter settings, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all specified using the service profile.

Figure 3   Cisco UCS VIC 1457 mLOM Card
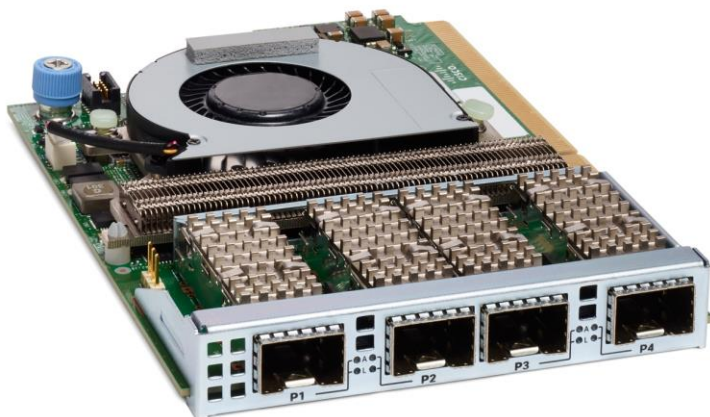
Cohesity Software

Cohesity addresses the growing challenges of secondary data sprawl and infrastructure silos by consolidating all secondary data and apps—including backups, disaster recovery, file services, object storage, test/dev, archives, and analytics—on a simple, software-defined, web-scale platform that extends from on-premises infrastructure to

the cloud and remote or branch offices. Complementing Cisco HyperFlex for primary data, Cohesity software integrates with Cisco UCS for better protected and more productive secondary data and applications.

- The Only Hyperconverged, Web-scale Platform for All Secondary Data and Apps

  Cohesity offers the industry's only secondary data and applications solution built on a hyperconverged, web-scale platform. Designed with Google-like principles, it delivers true global deduplication and impressive storage efficiency that spans edge to core to cloud, solving for backups, archives, files, objects, test/dev, and analytics. From the edge to the cloud, Cohesity delivers web-scale simplicity for secondary data and applications that eliminates silos and puts data to work. Global 2000 companies and federal agencies are modernizing and scaling with award-winning Cohesity solutions.

- End-to-end Data Protection, Archiving and Instant Recovery

  Cohesity provides the only end-to-end backup and recovery solution, providing simple data protection, archiving, recovery point objectives (RPOs) within minutes, instantaneous recovery time objectives (RTOs), and instant mass restore while cutting the cost of data protection in by more than half.

- Native Public Cloud Integration

  The Cohesity platform is built with the public cloud in mind. Cohesity's cloud-first architecture furthers enterprise objectives to leverage the cost and economic efficiencies of public cloud while staying in control. With Cohesity, customers can enable long-term retention, archival, disaster recovery, test/dev, and cloud backup all on the same platform.

- Unified Management

  Only Cohesity provides a SaaS-based management solution that unifies all secondary data and application infrastructures globally across edge, core and clouds and manages centrally with a single pane of glass. Powered by machine learning, it helps accelerate global IT productivity, improve business planning and continuity, and derive valuable insights into previously untapped secondary data.

Solution Design

Requirements

The following sections detail the physical hardware, software revisions, and firmware versions required to install a single cluster of the Cohesity DataPlatform running in Cisco Unified Computing System.

Physical Components

Table 1   Cohesity DataPlatform System Components

Component	Hardware Required
Fabric Interconnects	Two (2) Cisco UCS 6454 Fabric Interconnects
Servers	Minimum of three (3) Three Cisco UCS C240 M5 LFF rack servers

For complete server specifications and more information, please refer to the links below:

Cisco Fabric Interconnect 6454:

https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/ucs-6454-fab-int-specsheet.pdf

Cisco UCS C240 M5 LFF Server:

https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m5-lff-specsheet.pdf

Table 2lists the required hardware components and disk options for the Cisco UCS C240-M5L server model, which are required for installing the Cohesity DataPlatform:

Table 2   Cisco UCS C240 M5 LFF Server Options

C240 M5 LFF options		Hardware Required
Processors		Two Intel Xeon Processor 6142 Scalable Family CPUs
Memory		128 GB of total memory using four (4) 32 GB DDR4 2666 MHz 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA
Storage	SSDs	Two (2) 1.6 TB 2.5 Inch, High Performance, High Endurance NVMe SSDs (installed in the two read drive slots), or Two (2) 3.2 TB 2.5 Inch, High Performance, High Endurance NVMe SSDs (installed in the two read drive slots)
	HDDs	Twelve 4 TB 3.5 Inch, 12G SATA, 7200 RPM, 4K sector HDDs, or Twelve 10 TB 3.5 Inch, 12G SATA, 7200 RPM, 4K sector HDDs
Network		Cisco UCS VIC1457 MLOM
Boot Device		Two 240 GB M.2 form factor SATA SSDs

Note: Choose the 1.6 TB NVMe SSDs when also choosing the 4 TB HDD capacity drives. Choose the 3.2 TB NVMe SSDs along with the 10 TB HDD capacity drives.

Software Components

Table 3lists the software components and the versions required for a single cluster of the Cohesity DataPlatform running in the Cisco UCS system, as tested and validated in this document:

Table 3   Software Components

Component	Software Required
Cohesity	Cohesity 6.1.1a software or later
Cisco HyperFlex Data Platform	Cisco HyperFlex HX Data Platform Software 3.5(2a) or later
Cisco UCS Firmware	Cisco UCS Infrastructure software, B-Series and C-Series bundles, revision 4.0(1c) or later. Note: Cisco UCS Firmware 4.0(1a) is the minimum version required for any cluster containing Cisco Fabric Interconnect model 6454, and any server containing the VIC 1457.

Licensing

Cisco UCS systems and the Cohesity software must be properly licensed for all software features in use, and for all ports in use on the Cisco UCS Fabric Interconnects. Please contact your resale partner or your direct Cisco and Cohesity sales teams to ensure you order all of the necessary and appropriate licenses for your solution.

Physical Topology

Topology Overview

The Cisco Unified Computing System is composed of a pair of Cisco UCS Fabric Interconnects along with up to 160 B-Series blade servers, Cisco UCS C-Series rack-mount servers, HX-Series hyperconverged servers, or S-Series storage servers per UCS domain. Inside of a Cisco UCS domain, multiple environments can be deployed for differing workloads. For example, a Cisco HyperFlex cluster can be built using Cisco HX-Series rack-mount servers, a Cohesity cluster can be built using Cisco C-Series rack-mount servers, Cisco UCS B-Series blade servers inside of Cisco 5108 blade chassis can be used for various bare-metal or virtualized environments, and Cisco UCS S-Series storage servers can be deployed for high density storage. The two Fabric Interconnects both connect to every Cisco UCS C-Series, HX-Series, or Cisco UCS S-Series rack-mount server, and both of them also connect to every Cisco UCS 5108 blade chassis. Upstream network connections, also referred to as "northbound" network connections are made from the Fabric Interconnects to the customer datacenter network at the time of installation.

Figure 4   Cisco UCS Example Physical Topology

3-32 HyperFlex Converged Nodes   Plus Cohesity Rack Mount Servers   Plus Blade Servers
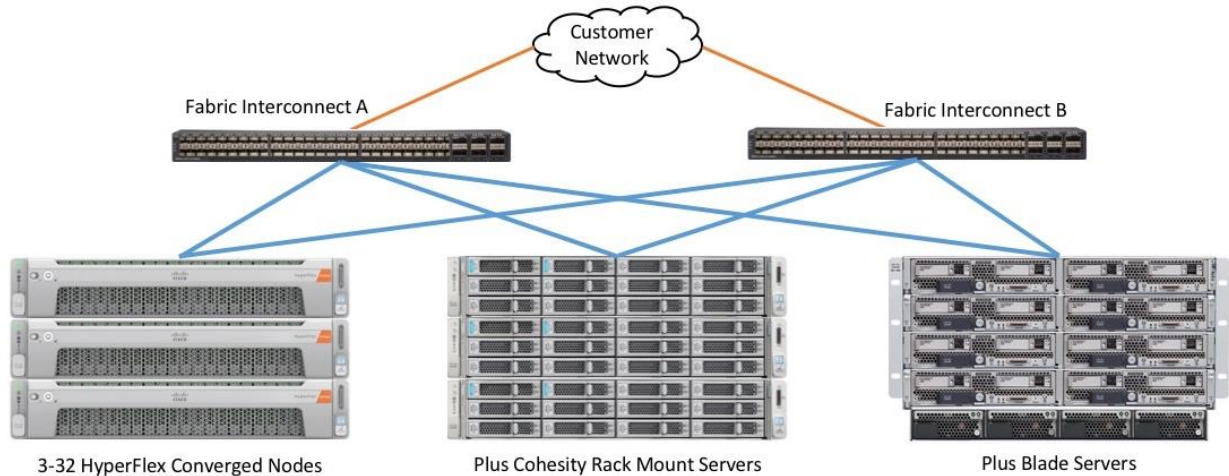
Fabric Interconnects

Cisco UCS Fabric Interconnects (FI) are deployed in pairs, wherein the two units operate as a management cluster, while forming two separate network fabrics, referred to as the A side and B side fabrics. Therefore, many design elements will refer to FI A or FI B, alternatively called fabric A or fabric B. Both Fabric Interconnects are active at all times, passing data on both network fabrics for a redundant and highly available configuration. Management services, including Cisco UCS Manager, are also provided by the two FIs but in a clustered manner, where one FI is the primary, and one is secondary, with a roaming clustered IP address. This primary/secondary relationship is only for the management cluster, and has no effect on data transmission.

Fabric Interconnects have the following ports, which must be connected for proper management of the Cisco UCS domain:

- Mgmt: A 10/100/1000 Mbps port for managing the Fabric Interconnect and the Cisco UCS domain via GUI and CLI tools. This port is also used by remote KVM, IPMI and SoL sessions to the managed servers within the domain. This is typically connected to the customer management network.

- L1: A cross connect port for forming the Cisco UCS management cluster. This port is connected directly to the L1 port of the paired Fabric Interconnect using a standard CAT5 or CAT6 Ethernet cable with RJ45 plugs. It is not necessary to connect this to a switch or hub.

- L2: A cross connect port for forming the Cisco UCS management cluster. This port is connected directly to the L2 port of the paired Fabric Interconnect using a standard CAT5 or CAT6 Ethernet cable with RJ45 plugs. It is not necessary to connect this to a switch or hub.

- Console: An RJ45 serial port for direct console access to the Fabric Interconnect. This port is typically used during the initial FI setup process with the included serial to RJ45 adapter cable. This can also be plugged into a terminal aggregator or remote console server device.

Cisco UCS C-Series Rack-Mount Servers

Cohesity UCS clusters require a minimum of three (3) Cisco UCS C240 M5 LFF Rack-Mount Servers. The Cisco UCS C-Series Rack-Mount Servers are connected directly to the Cisco UCS Fabric Interconnects in Direct Connect mode. Internally the Cisco UCS C-Series servers are configured with the Cisco VIC 1457 network interface card (NIC) installed in a modular LAN on motherboard (MLOM) slot, which has quad 10/25 Gigabit

Ethernet (GbE) ports. The standard and redundant connection practice is to connect port 1 of each server's VIC card to a numbered port on FI A, and port 3 of each server's VIC card to the same numbered port on FI B. The use of ports 1 and 3 are due to the fact that ports 1 and 2 form an internal port-channel, as does ports 3 and 4. This allows an optional 4 cable connection method, which is not used in this design.

◭	Note: Failure to follow this cabling practice can lead to errors, discovery failures, and loss of redundant connectivity.

◭	Warning: Do not connect port 1 of the VIC 1457 to Fabric Interconnect A, and then connect port 2 of the VIC 1457 to Fabric Interconnect B. Only use ports 1 and 3 on the VIC 1457 card. Using ports 1 and 2 only will lead to discovery and configuration failures.

Figure 5   Cisco UCS C-Series Server Connectivity

Cisco C240 M5 LFF Cohesity Node
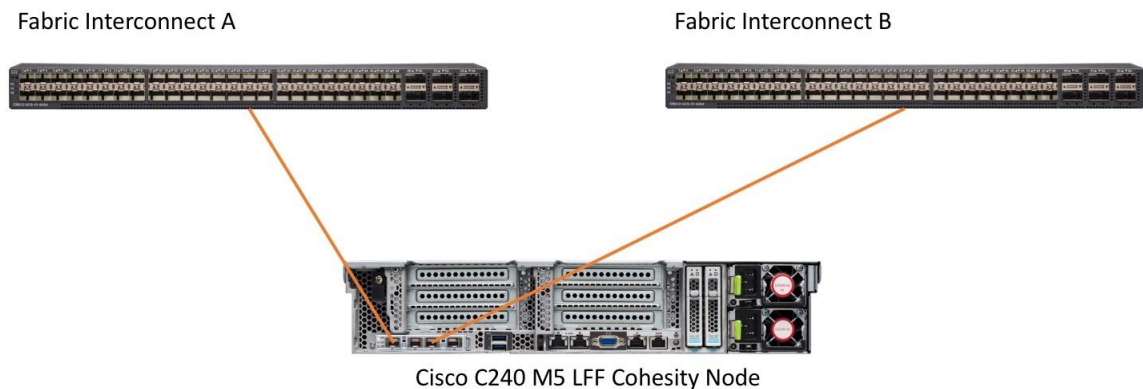
Logical Topology

Logical Network Design

The Cohesity DataPlatform running on Cisco UCS has communication pathways that fall into two defined zones (Figure 6):

- Management Zone: This zone comprises the connections needed to manage the physical hardware, and the configuration of the Cisco UCS domain. These interfaces and IP addresses need to be available to all staff who will administer the UCS system, throughout the LAN/WAN. All IP addresses in this zone must be allocated from the same layer 2 (L2) subnet. This zone must provide access to Domain Name System (DNS), Network Time Protocol (NTP) services, and allow communication via HTTP/S and Secure Shell (SSH). In this zone are multiple physical and virtual components:

  - Fabric Interconnect management ports.

  - Cisco Intelligent Management Controller (CIMC) management interfaces used by each the rack-mount servers and blades, which answer via the FI management ports.

- Application Zone: This zone comprises the connections used by the Cohesity Data Platform software and the underlying operating system on the nodes. These interfaces and IP addresses need to be able to communicate with each other at all times for proper operation, and they must be allocated from the same

L2 subnet. The VLAN used for Cohesity application traffic must be accessible to/from all environments which will be protected by Cohesity, such as the management interfaces of the Cisco HyperFlex nodes, and also its managing vCenter Server system. This zone must provide access to Domain Name System (DNS), Network Time Protocol (NTP) services, and allow communication via HTTP/S and Secure Shell (SSH). Additionally, clients which will connect to Cohesity for file services must also be able to access this VLAN. Finally, the VLAN must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B directly via the northbound switches, and vice-versa. In this zone are multiple components:

– A static IP address configured for the underlying Linux operating system of each Cohesity node. Two UCS vNICs are configured per node, one on the A side fabric, and one on the B side fabric. The two interfaces are configured as slave interfaces in a bond within the Linux operating system, using bond mode 1 (active/passive).

– A floating virtual IP address (VIP), one per node, that is used by Cohesity for all management, backup, and file services access. The assignment of the addresses is handled by the Cohesity software and will be re-assigned to an available node if any node should fall offline. These floating addresses are all assigned in DNS to a single A record, and the DNS server must respond to queries for that A record using DNS round-robin.

Figure 6   Logical Network Design
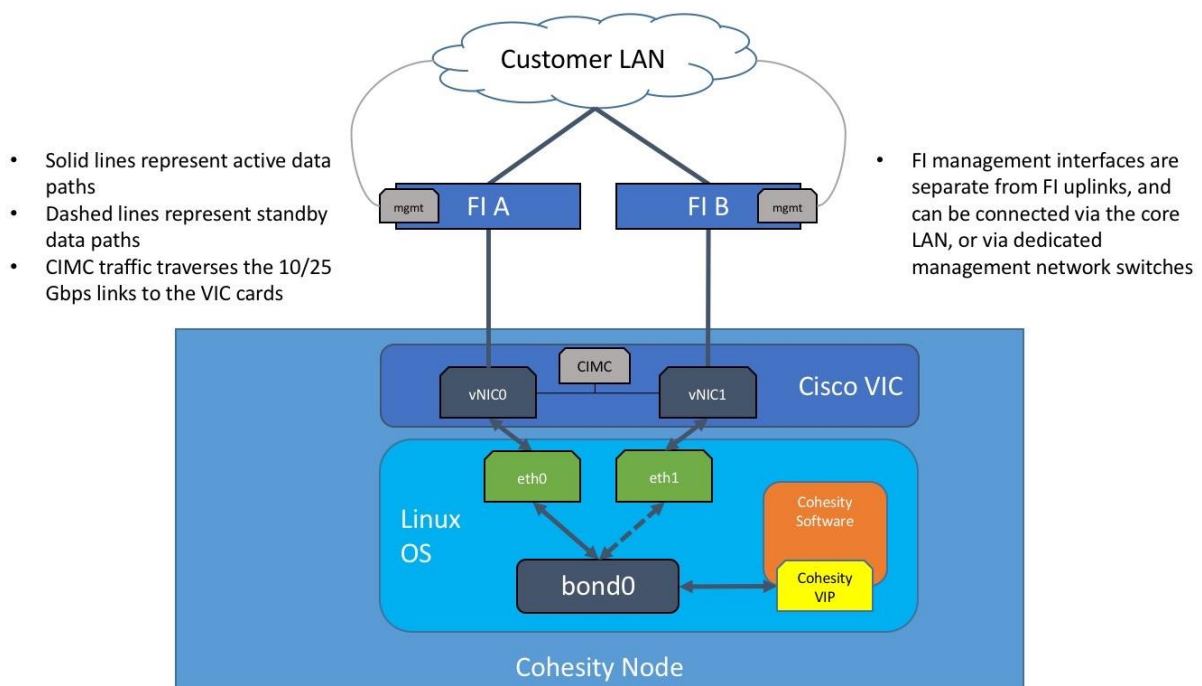
Network Design

Cisco UCS Uplink Connectivity

Cisco UCS network uplinks connect "northbound" from the pair of Cisco UCS Fabric Interconnects to the LAN in the customer datacenter. All Cisco UCS uplinks operate as trunks, carrying multiple 802.1Q VLAN IDs across the

uplinks. The default Cisco UCS behavior is to assume that all VLAN IDs defined in the Cisco UCS configuration are eligible to be trunked across all available uplinks.

Cisco UCS Fabric Interconnects appear on the network as a collection of endpoints versus another network switch. Internally, the Fabric Interconnects do not participate in spanning-tree protocol (STP) domains, and the Fabric Interconnects cannot form a network loop, as they are not connected to each other with a layer 2 Ethernet link. All link up/down decisions via STP will be made by the upstream root bridges.

Uplinks need to be connected and active from both Fabric Interconnects. For redundancy, multiple uplinks can be used on each FI, either as 802.3ad Link Aggregation Control Protocol (LACP) port-channels or using individual links. For the best level of performance and redundancy, uplinks can be made as LACP port-channels to multiple upstream Cisco switches using the virtual port channel (vPC) feature. Using vPC uplinks allows all uplinks to be active passing data, plus protects against any individual link failure, and the failure of an upstream switch. Other uplink configurations can be redundant, however spanning-tree protocol loop avoidance may disable links if vPC is not available.

All uplink connectivity methods must allow for traffic to pass from one Fabric Interconnect to the other, or from fabric A to fabric B. There are scenarios where cable, port or link failures would require traffic that normally does not leave the Cisco UCS domain, to now be forced over the Cisco UCS uplinks. Additionally, this traffic flow pattern can be seen briefly during maintenance procedures, such as updating firmware on the Fabric Interconnects, which requires them to be rebooted. The following sections and figures detail several uplink connectivity options.

Single Uplinks to Single Switch

This connection design is susceptible to failures at several points; single uplink failures on either Fabric Interconnect can lead to connectivity losses or functional failures, and the failure of the single uplink switch will cause a complete connectivity outage.

Figure 7   Connectivity with Single Uplink to Single Switch
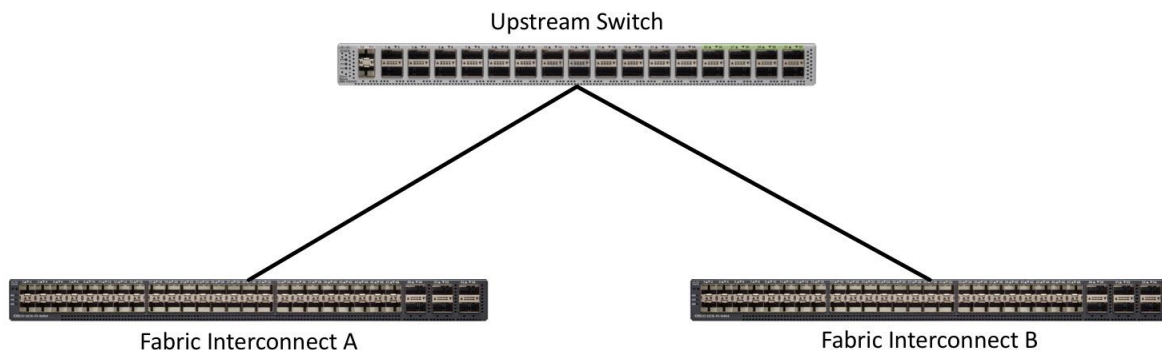
Port Channels to Single Switch

This connection design is now redundant against the loss of a single link but remains susceptible to the failure of the single switch.

Figure 8   Connectivity with Port–Channels to Single Switch

Figure 8 diagram: Upstream Switch connected via po 10 to Fabric Interconnect A and via po 20 to Fabric Interconnect B.
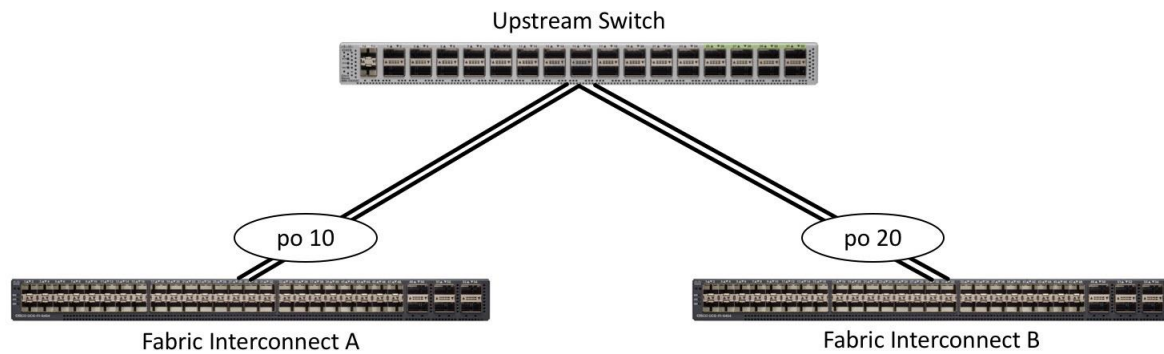
Single Uplinks or Port Channels to Multiple Switches

This connection design is redundant against the failure of an upstream switch, and redundant against a single link failure. In normal operation, STP is likely to block half of the links to avoid a loop across the two upstream switches. The side effect of this is to reduce bandwidth between the Cisco UCS domain and the LAN. If any of the active links were to fail, STP would bring the previously blocked link online to provide access to that Fabric Interconnect via the other switch. It is not recommended to connect both links from a single FI to a single switch, as that configuration is susceptible to a single switch failure breaking connectivity from fabric A to fabric B. For enhanced redundancy, the single links in the figure below could also be port–channels.

Figure 9   Connectivity with Multiple Uplink Switches

Figure 9 diagram: Two Upstream Switches connected to Fabric Interconnect A and Fabric Interconnect B, with two links Blocked by STP.
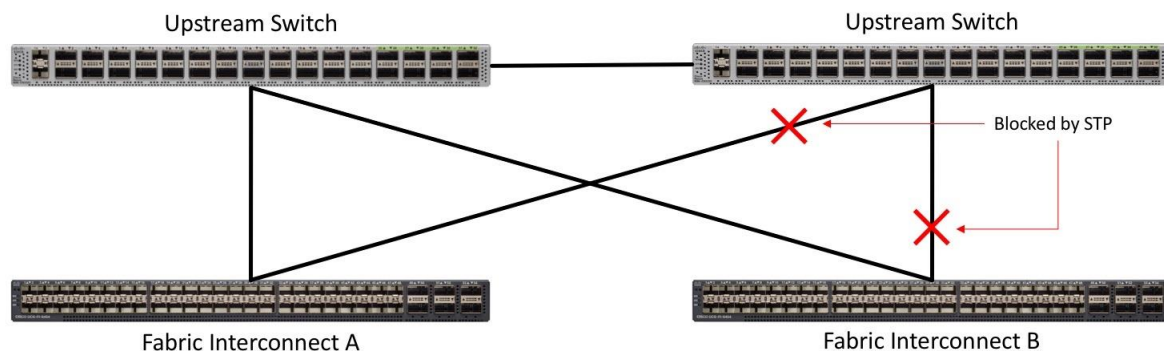
vPC to Multiple Switches

This recommended connection design relies on using Cisco switches that have the virtual port channel feature, such as Catalyst 6000 series switches running VSS, Cisco Nexus 5000 series, and Cisco Nexus 9000 series switches. Logically the two vPC enabled switches appear as one, and therefore spanning–tree protocol will not block any links. This configuration allows for all links to be active, achieving maximum bandwidth potential, and multiple redundancy at each level.

Figure 10   Connectivity with vPC
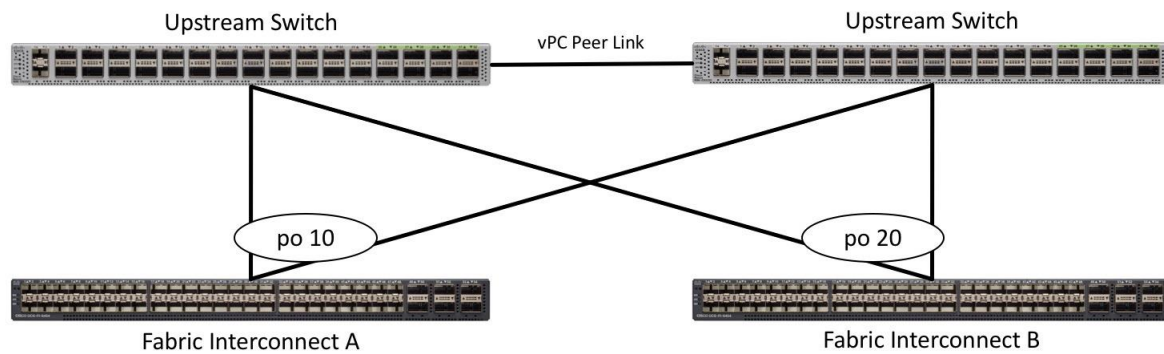
VLANs and Subnets

For the Cohesity system configuration, one only one VLAN is needed to be carried to the Cisco UCS domain from the upstream LAN, and this VLAN is also defined in the Cisco UCS configuration. Table 4lists the VLANs required by Cohesity in Cisco UCS, and their functions:

Table 4   VLANs

VLAN Name	VLAN ID	Purpose
<<cohesity_vlan>>	Customer supplied	Cohesity node Linux OS interfaces Cohesity node software virtual IP addresses

Jumbo Frames

All Cohesity traffic traversing the <<cohesity_vlan>> VLAN and subnet is configured by default to use standard ethernet frames.

Considerations

Prior to the installation of the cluster, proper consideration must be given to the number of nodes required for the Cohesity cluster, and the usable capacity that will result.

Scale

Cohesity clusters require a minimum of three (3) Cisco C240 M5 LFF rack-mount server nodes to create an initial cluster. From that point, the cluster can grow to any size of cluster that is required by the end user which meets their overall storage space requirements. This limitless scaling is a key feature present in Cohesity which allows future growth without the fears of reaching an overall capacity restriction.

Capacity

Overall usable cluster capacity is based on a number of factors, primarily the number of nodes in the cluster, and the number and size of the capacity layer disks. Caching disk sizes are not calculated as part of the cluster capacity. Space consumption is a direct result of the configured replication factor in the Storage Domains created in the Cohesity cluster. In addition, deduplication and compression of the incoming data affects the efficiency of the data being stored. The authoritative view of the Cohesity cluster's available and used capacity is provided by the Cohesity HTML management dashboard.

Disk drive manufacturers have adopted a size reporting methodology using calculation by powers of 10, also known as decimal prefix. As an example, a 120 GB disk is listed with a minimum of 120 x 10^9 bytes of usable addressable capacity, or 120 billion bytes. However, many operating systems and filesystems report their space based on standard computer binary exponentiation, or calculation by powers of 2, also called binary prefix. In this example, 2^10 or 1024 bytes make up a kilobyte, 2^10 kilobytes make up a megabyte, 2^10 megabytes make up a gigabyte, and 2^10 gigabytes make up a terabyte. As the values increase, the disparity between the two systems of measurement and notation get worse, at the terabyte level, the deviation between a decimal prefix value and a binary prefix value is nearly 10 percent.

The International System of Units (SI) defines values and decimal prefix by powers of 10 as follows:

Table 5   SI Unit Values (Decimal Prefix)

Value	Symbol	Name
1000 bytes	kB	Kilobyte
1000 kB	MB	Megabyte
1000 MB	GB	Gigabyte
1000 GB	TB	Terabyte

The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) defines values and binary prefix by powers of 2 in ISO/IEC 80000-13:2008 Clause 4 as follows:

Table 6   IEC unit values (binary prefix)

Value	Symbol	Name
1024 bytes	KiB	Kibibyte
1024 KiB	MiB	Mebibyte
1024 MiB	GiB	Gibibyte
1024 GiB	TiB	Tebibyte

For the purpose of this document, the decimal prefix numbers are used only for raw disk capacity as listed by the respective manufacturers. For all calculations where raw or usable capacities are shown from the perspective of the Cohesity software, filesystems or operating systems, the binary prefix numbers are used. This is done primarily to show a consistent set of values as seen by the end user from within the Cohesity HTML management dashboard when viewing cluster capacity, allocation and consumption, and also within most operating systems.

Table 7lists a set of Cohesity DataPlatform cluster usable capacity values, using binary prefix, for an array of cluster configurations. These values provide an example of the capacity calculations, for determining the appropriate size of Cohesity cluster to initially purchase. Additional savings from deduplication and compression will raise the effective logical capacity far beyond the physical capacity of the nodes. Additionally, the choice of replication factor 2, or erasure coding, will determine the overall efficiency of the real data being stored on the nodes.

Table 7   Cohesity Cluster Usable Physical Capacities

C-Series Server Model	Capacity Disk Size (each)	Capacity Disk Quantity (per node)	Usable Capacity (per node)	Capacity per node @ RF2	Capacity per node with EC 2:1
C240-M5L	4 TB	12	45.18 TiB	22.59 TiB	30.12 TiB
	10 TB	12	112.95 TiB	56.48 TiB	75.3 TiB

Configuration and Installation

Installing the Cohesity DataPlatform system is done through mounting a virtual DVD image to each Cisco C240 M5 node, which is available for download from Cohesity as an ISO file. The installation DVD validates the hardware configuration, installs the Linux operating system, copies the Cohesity software packages, and completes with the nodes ready for their final configuration to form a cluster. Prior to using the installation DVD, the configuration of the Cisco UCS domain, its policies, templates, and service profiles to be associated to the servers must be completed. The following sections will guide you through the prerequisites and manual steps needed to configure Cisco UCS Manager prior to booting the Cohesity installation DVD, the steps to install Cohesity to each node, and how to perform the remaining post-installation tasks to configure the Cohesity cluster. Finally, a basic configuration example is given for configuring Cohesity Storage Domains, Sources, Policies, Protection Jobs, file services Views, and Test/Dev VM services.

Prerequisites

Prior to beginning the installation activities, perform the following necessary tasks and gather the required information:

IP Addressing

IP addresses for the Cohesity system on Cisco UCS need to be allocated from the appropriate subnets and VLANs to be used. IP addresses that are used by the system fall into the following groups:

- UCS Management: These addresses are used and assigned by Cisco UCS Manager. Three IP addresses are used by Cisco UCS Manager; one address is assigned to each Cisco UCS Fabric Interconnect, and the third IP address is a roaming address for management of the Cisco UCS cluster. In addition, at least one IP address per Cisco UCS C-series rack-mount server is required for the Cohesity external management IP address pool, which are assigned to the CIMC interface of the physical servers. Since these management addresses are assigned from a pool, they need to be provided in a contiguous block of addresses. These addresses must all be in the same subnet.

- Cohesity Application: These addresses are used by the Linux OS on each Cohesity node, and the Cohesity software. Two IP addresses per node in the Cohesity cluster are required from the same subnet. These addresses can be assigned from the same subnet at the Cisco UCS Management addresses, or they may be separate.

The following tables will assist with gathering the required IP addresses for the installation of a 4-node standard Cohesity cluster by listing the addresses required, plus an example IP configuration:

⚠️   Note: Table cells shaded in black do not require an IP address.

Table 8   Cohesity Cluster IP Addressing

Address Group:	UCS Management	Cohesity Application
VLAN ID:		
Subnet:		
Subnet Mask:		
Gateway:		

Device	UCS Management Addresses	OS Interface	Cohesity VIP
Fabric Interconnect A			
Fabric Interconnect B			
UCS Manager			
Cohesity Node #1			
Cohesity Node #2			
Cohesity Node #3			
Cohesity Node #4			

Table 9   Example Cohesity Cluster IP Addressing

Address Group:	UCS Management	Cohesity Application	
VLAN ID:	133	133	
Subnet:	10.29.133.0	10.29.133.0	
Subnet Mask:	255.255.255.0	255.255.255.0	
Gateway:	10.29.133.1	10.29.133.1	
Device	UCS Management Addresses	OS Interface	Cohesity VIP
Fabric Interconnect A	10.29.133.104		
Fabric Interconnect B	10.29.133.105		
UCS Manager	10.29.133.106		
Cohesity Node #1	10.29.133.133	10.29.133.143	10.29.133.152
Cohesity Node #2	10.29.133.134	10.29.133.144	10.29.133.153
Cohesity Node #3	10.29.133.135	10.29.133.145	10.29.133.154
Cohesity Node #4	10.29.133.136	10.29.133.146	10.29.133.155

DNS

DNS servers are required to be configured for querying Fully Qualified Domain Names (FQDN) in the Cohesity application group. DNS records need to be created prior to beginning the installation. At a minimum, it is required to create a single A record for the name of the Cohesity cluster, which answers with each of the virtual IP addresses used by the Cohesity nodes in round-robin fashion. Some DNS servers are not configured by default to return multiple addresses in round-robin fashion in response to a request for a single A record, please ensure your DNS server is properly configured for round-robin before continuing. The configuration can be tested by querying the DNS name of the Cohesity cluster from multiple clients, and verifying that all of the different IP addresses are given as answers in turn.

The following tables will assist with gathering the required DNS information for the installation, by listing the information required, and an example configuration:

Table 10 DNS Server Information

Item	Value	A Records
DNS Server #1		
DNS Server #2		
DNS Domain		
vCenter Server Name		
UCS Domain Name		
Cohesity Cluster Name		

Table 11 DNS Server Example Information

Item	Value	A Records
DNS Server #1	10.29.133.110	
DNS Server #2		
DNS Domain	hx.lab.cisco.com	
vCenter Server Name	vcenter.hx.lab.cisco.com	10.29.133.121
UCS Domain Name	HX1-FI	
Cohesity Cluster Name	chx-cluster01	10.29.133.152

Item	Value	A Records
		10.29.133.153
		10.29.133.154
		10.29.133.155

NTP

Consistent time clock synchronization is required across the components of the Cohesity cluster, provided by reliable NTP servers, accessible for querying in the Cisco UCS Management network group, and the Cohesity Application group. NTP is used by many components, such as Cisco UCS Manager, vCenter, the Cohesity cluster nodes, and the HyperFlex Storage Platform. The use of public NTP servers is highly discouraged, instead a reliable internal NTP server should be used.

The following tables will assist with gathering the required NTP information for the installation by listing the information required, and an example configuration:

Table 12 NTP Server Information

Item	Value
NTP Server #1	
NTP Server #2	
Timezone	

Table 13 NTP Server Example Information

Item	Value
NTP Server #1	ntp1.hx.lab.cisco.com
NTP Server #2	ntp2.hx.lab.cisco.com
Timezone	(UTC-8:00) Pacific Time

VLANs

Prior to the installation, the required VLAN IDs need to be documented, and created in the upstream network if necessary. There is one VLAN that needs to be trunked to the two Cisco UCS Fabric Interconnects which manage the Cohesity cluster; the VLAN for the Cohesity Application group. The VLAN IDs must be supplied during the Cisco UCS configuration steps, and the VLAN names should be customized to make them easily identifiable.

The following tables will assist with gathering the required VLAN information for the installation by listing the information required, and an example configuration:

Table 14 VLAN Information

Name	ID
<<cohesity_vlan>>	

Table 15 VLAN Example Information

Name	ID
cohesity-vlan-133	133

Network Uplinks

The Cisco UCS uplink connectivity design needs to be finalized prior to beginning the installation. Refer to the network uplink design possibilities in the Network Design section.

The following tables will assist with gathering the required network uplink information for the installation by listing the information required, and an example configuration:

Table 16 Network Uplink Configuration

Fabric Interconnect Port	Port Channel	Port Channel Type	Port Channel ID	Port Channel Name
A	☐ Yes ☐ No	☐ LACP ☐ vPC		
	☐ Yes ☐ No			
	☐ Yes ☐ No			
	☐ Yes ☐ No			
B	☐ Yes ☐ No	☐ LACP ☐ vPC		
	☐ Yes ☐ No			
	☐ Yes ☐ No			
	☐ Yes ☐ No			

Table 17 Network Uplink Example Configuration

Fabric Interconnect Port	Port Channel	Port Channel Type	Port Channel ID	Port Channel Name
A	1/53 ☒ Yes ☐ No	☐ LACP ☒ vPC	13	vpc13
	1/54 ☒ Yes ☐ No			
	☐ Yes ☐ No			

Fabric Interconnect Port	Port Channel		Port Channel Type	Port Channel ID	Port Channel Name
		☐ Yes ☐ No			
B	1/53	☒ Yes ☐ No	☐ LACP ☒ vPC	23	vpc23
	1/54	☒ Yes ☐ No			
		☐ Yes ☐ No			
		☐ Yes ☐ No			

Usernames and Passwords

Several usernames and passwords need to be defined or known as part of the Cohesity installation and configuration process. The following tables will assist with gathering the required username and password information by listing the information required and an example configuration:

Table 18 Usernames and Passwords

Account	Username	Password
UCS Administrator	admin	<<ucs_admin_pw>>
Cohesity Administrator	admin	<<cohesity_admin_pw>>
HyperFlex Administrator	admin	<<hx_admin_pw>>
vCenter Administrator	<<vcenter_administrator>>	<<vcenter_admin_pw>>

Table 19 Example Usernames and Passwords

Account	Username	Password
UCS Administrator	admin	Cisco123
Cohesity Administrator	admin	admin
HyperFlex Administrator	admin	CIsco123!!
vCenter Administrator	administrator@vsphere.local	!Q2w3e4r

Physical Installation

Install the Fabric Interconnects and the Cisco UCS C-Series rack-mount servers according to their corresponding hardware installation guides listed below.

Cisco UCS 6454 Fabric Interconnect:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/6454-install-guide/6454.pdf

Cisco UCS C240 M5 LFF rack-mount server:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/C240M5/install/C240M5.pdf

Cabling

The physical layout of the Cohesity system was previously described in section Physical Topology. The Fabric Interconnects and C-series rack-mount servers need to be cabled properly before beginning the installation activities.

Table 20provides an example cabling map for installation of a Cohesity system, using four C-Series Cohesity converged nodes as tested in this document.

Table 20 Example Cabling Map

Device	Port	Connected To	Port	Type	Length	Note
UCS6454-A	L1	UCS6454-B	L1	CAT5	1FT	
UCS6454-A	L2	UCS6454-B	L2	CAT5	1FT	
UCS6454-A	mgmt0	Customer LAN		CAT5		Management interface
UCS6454-A	1/1	Cohesity Server #1	mLOM port 1	Twinax	3M	Server 1
UCS6454-A	1/2	Cohesity Server #2	mLOM port 1	Twinax	3M	Server 2
UCS6454-A	1/3	Cohesity Server #3	mLOM port 1	Twinax	3M	Server 3
UCS6454-A	1/4	Cohesity Server #4	mLOM port 1	Twinax	3M	Server 4
UCS6454-A	1/53	Customer LAN				uplink
UCS6454-A	1/54	Customer LAN				uplink
UCS6454-B	L1	UCS6454-A	L1	CAT5	1FT	
UCS6454-B	L2	UCS6454-A	L2	CAT5	1FT	
UCS6454-B	mgmt0	Customer LAN		CAT5		Management interface
UCS6454-B	1/1	Cohesity Server #1	mLOM port 3	Twinax	3M	Server 1
UCS6454-B	1/2	Cohesity Server #2	mLOM port 3	Twinax	3M	Server 2
UCS6454-B	1/3	Cohesity Server #3	mLOM port 3	Twinax	3M	Server 3
UCS6454-B	1/4	Cohesity Server #4	mLOM port 3	Twinax	3M	Server 4
UCS6454-B	1/53	Customer LAN				uplink
UCS6454-B	1/54	Customer LAN				uplink

Warning: Do not connect port 1 of the VIC 1457 to Fabric Interconnect A, and then connect port 2 of the VIC 1457 to Fabric Interconnect B. Only use ports 1 and 3 on the VIC 1457 card. Using ports 1 and 2 only will lead to discovery and configuration failures.

Cisco UCS Installation

This section describes the steps to initialize and configure the Cisco UCS Fabric Interconnects, to prepare them for the Cohesity installation. For installations of Cohesity being integrated into an existing Cisco UCS domain, the following steps outlining the initial setup of the Fabric Interconnects, and their uplink port configuration can be skipped. In this situation, the steps beginning with the configuration of the server ports and server discovery onwards, including sub-organizations, policies, pools, templates and service profiles, must still be performed.

Cisco UCS Fabric Interconnect A

To configure Fabric Interconnect A, complete the following steps:

1. Make sure the Fabric Interconnect cabling is properly connected, including the L1 and L2 cluster links, and the management ports, then power the Fabric Interconnects on by inserting the power cords.

2. Connect to the console port on the first Fabric Interconnect, which will be designated as the A fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.

3. Start your terminal emulator software.

4. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, and 1 stop bit.

5. Open the connection which was just created. You may have to press ENTER to see the first prompt.

6. Configure the first Fabric Interconnect, using the following example as a guideline:

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

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

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


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

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

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

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

Enter the password for "admin":
Confirm the password for "admin":

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

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

Enter the system name:  HX1-FI

Physical Switch Mgmt0 IP address : 10.29.133.104
```

```
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.29.133.1

Cluster IPv4 address : 10.29.133.106

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

  DNS IP address : 10.29.133.110

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

  Default domain name : hx.lab.cisco.com

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

Following configurations will be applied:

  Switch Fabric=A
  System Name=HX1-FI
  Enforced Strong Password=no
  Physical Switch Mgmt0 IP Address=10.29.133.104
  Physical Switch Mgmt0 IP Netmask=255.255.255.0
  Default Gateway=10.29.133.1
  Ipv6 value=0
  DNS Server=10.29.133.110
  Domain Name=hx.lab.cisco.com

  Cluster Enabled=yes
  Cluster IP Address=10.29.133.106
  NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized

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

Configuration file - Ok
```

Cisco UCS Fabric Interconnect B

To configure Fabric Interconnect B, complete the following steps:

1. Connect to the console port on the first Fabric Interconnect, which will be designated as the B fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.

2. Start your terminal emulator software.

3. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, and 1 stop bit.

4. Open the connection which was just created. You may have to press ENTER to see the first prompt.

5. Configure the second Fabric Interconnect, using the following example as a guideline:

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

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

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


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

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

   Enter the admin password of the peer Fabric interconnect:
     Connecting to peer Fabric interconnect... done
     Retrieving config from peer Fabric interconnect... done
     Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.133.104
     Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
     Cluster IPv4 address          : 10.29.133.106

     Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address

   Physical Switch Mgmt0 IP address : 10.29.133.105


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

Configuration file – Ok
```

Cisco UCS Manager

Log into the Cisco UCS Manager environment by completing the following steps:

1. Open a web browser and navigate to the Cisco UCS Manager Cluster IP address, for example
   https://10.29.133.106
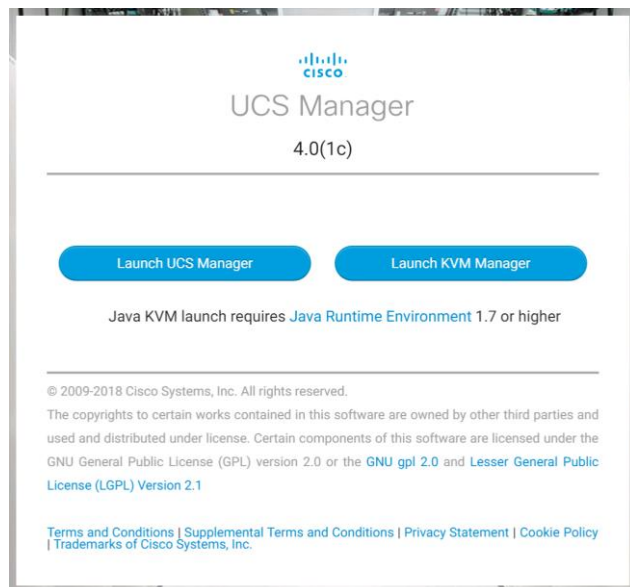
2. Click the "Launch UCS Manager" HTML link to open the Cisco UCS Manager web client.

3. At the login prompt, enter "admin" as the username, and enter the administrative password that was set during the initial console configuration.

4. Click No when prompted to enable Cisco Smart Call Home, this feature can be enabled at a later time.
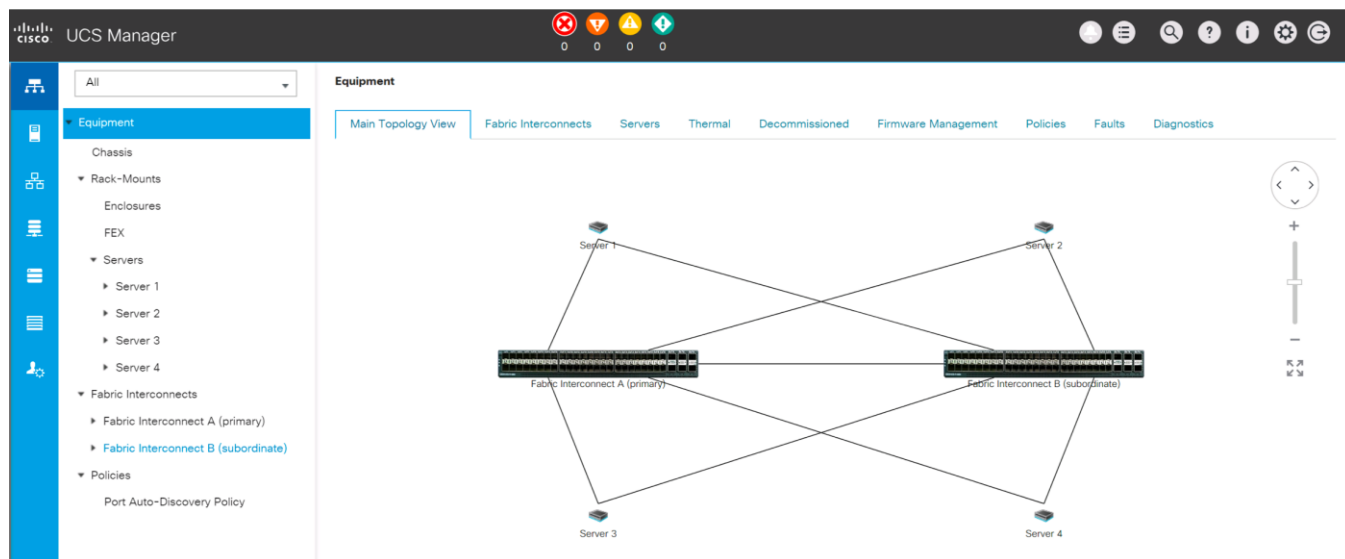
Cisco UCS Configuration

Configure the following ports, settings, and policies in the Cisco UCS Manager interface prior to beginning the Cohesity DataPlatform installation.

Cisco UCS Firmware

Your Cisco UCS firmware version should be current as shipped from the factory, as documented in the Software Components section. This document is based on Cisco UCS infrastructure, Cisco UCS B-series bundle, and Cisco UCS C-Series bundle software versions 4.0(1c). If the firmware version of the Fabric Interconnects is older than this version, the firmware must be upgraded to match the requirements prior to completing any further steps. To upgrade the Cisco UCS Manager version, the Fabric Interconnect firmware, and the server bundles, refer to these instructions:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Firmware-Mgmt/3-2/b_UCSM_GUI_Firmware_Management_Guide_3_2.html

NTP

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

1. In Cisco UCS Manager, click the Admin button on the left-hand side.

2. In the navigation pane, select All > Time Zone Management, and click the carat next to Time Zone Management to expand it.

3. Click Timezone.

4. In the Properties pane, select the appropriate time zone in the Time Zone menu.

5. Click Add NTP Server.

6. Enter the NTP server IP address and click OK.

7. Click OK.

8. Click Save Changes and then click OK.
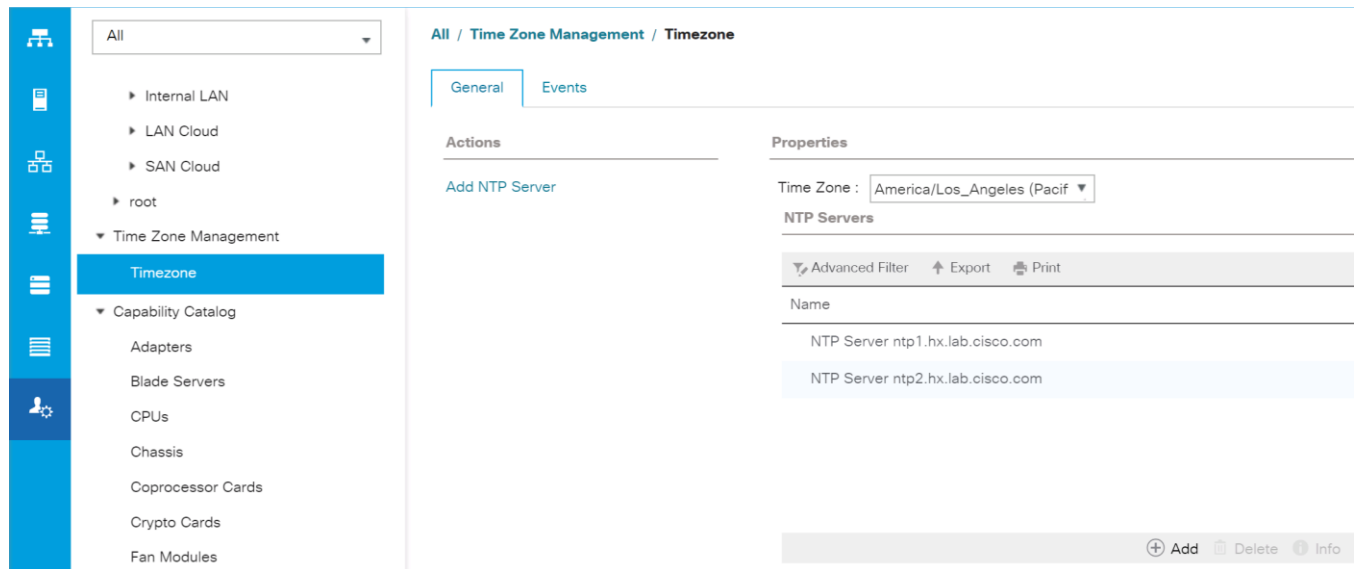
Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment button on the left-hand side.

2. Select Fabric Interconnects > Fabric Interconnect A > Fixed Module or Expansion Module as appropriate > Ethernet Ports.

3. Select the ports that are to be uplink ports, right click them, and click Configure as Uplink Port.

4. Click Yes to confirm the configuration, and click OK.

5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module as appropriate > Ethernet Ports.

6. Select the ports that are to be uplink ports, right-click them, and click Configure as Uplink Port.

7. Click Yes to confirm the configuration and click OK.

8. Verify all the necessary ports are now configured as uplink ports, where their role is listed as "Network."

Uplink Port Channels

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

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. Under LAN > LAN Cloud, click the carat to expand the Fabric A tree.

3. Right-click Port Channels underneath Fabric A, then click Create Port Channel.

4. Enter the port channel ID number as the unique ID of the port channel (this does not have to match the port-channel ID on the upstream switch).

5. Enter the name of the port channel.

6. Click Next.

7. Click each port from Fabric Interconnect A that will participate in the port channel, and click the >> button to add them to the port channel.

8. Click Finish.

9. Click OK.

10. Under LAN > LAN Cloud, click the carat to expand the Fabric B tree.

11. Right-click Port Channels underneath Fabric B, then click Create Port Channel.

12. Enter the port channel ID number as the unique ID of the port channel (this does not have to match the port-channel ID on the upstream switch).

13. Enter the name of the port channel.

14. Click Next.

15. Click each port from Fabric Interconnect B that will participate in the port channel, and click the >> button to add them to the port channel.

16. Click Finish.

17. Click OK.

18. Verify the necessary port channels have been created. It can take a few minutes for the newly formed port channels to converge and come online.

Server Ports

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

Auto Configuration

A new feature in Cisco UCS Manager 3.1(3a) and later is Server Port Auto-Discovery, which automates the configuration of ports on the Fabric Interconnects as server ports when a Cisco UCS rack-mount server or blade chassis is connected to them. The firmware on the rack-mount servers or blade chassis Fabric Extenders must already be at version 3.1(3a) or later in order for this feature to function properly. Enabling this policy eliminates the manual steps of configuring each server port, however it does configure the servers in a somewhat random order. For example, the rack-mount server at the bottom of the stack, which you may refer to as server #1, and you may have plugged into port 1 of both Fabric Interconnects, could be discovered as server 2, or server 5, etc. In order to have fine control of the rack-mount server or chassis numbering and order, the manual configuration steps listed in the next section must be followed.

To configure automatic server port definition and discovery, complete the following steps:

1. In Cisco UCS Manager, click the Equipment button on the left-hand side.

2. In the navigation tree, under Policies, click Port Auto-Discovery Policy

3. In the properties pane, set Auto Configure Server Port option to Enabled.

4. Click Save Changes.

5. Click OK.

6. Wait for a brief period, until the rack-mount servers appear in the Equipment tab underneath Equipment > Rack Mounts > Servers, or the chassis appears underneath Equipment > Chassis.
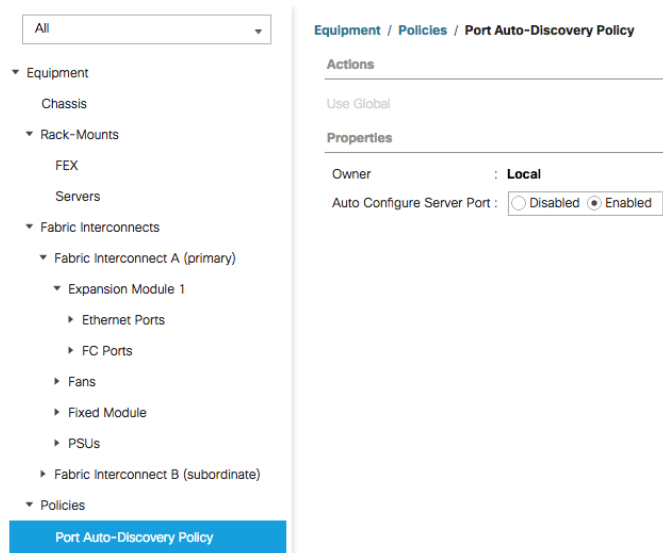
Manual Configuration

To manually define the specified ports to be used as server ports, and have control over the numbering of the servers, complete the following steps:

1. In Cisco UCS Manager, click the Equipment button on the left-hand side.

2. Select Fabric Interconnects > Fabric Interconnect A > Fixed Module or Expansion Module as appropriate > Ethernet Ports.

3. Select the first port that is to be a server port, right-click it, and click Configure as Server Port.

4. Click Yes to confirm the configuration and click OK.

5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module as appropriate > Ethernet Ports.

6. Select the matching port as chosen for Fabric Interconnect A that is to be a server port, right-click it, and click Configure as Server Port.

7. Click Yes to confirm the configuration and click OK.

37

8. Wait for a brief period, until the rack-mount server appears in the Equipment tab underneath Equipment > Rack Mounts > Servers, or the chassis appears underneath Equipment > Chassis.

9. Repeat Steps 1-8 for each pair of server ports, until all rack-mount servers and chassis appear in the order desired in the Equipment tab.

Server Discovery

As previously described, when the server ports of the Fabric Interconnects are configured and active, the servers connected to those ports will begin a discovery process. During discovery, the servers' internal hardware inventories are collected, along with their current firmware revisions. Before continuing with the process of associating the servers with their service profiles, wait for all of the servers to finish their discovery process and to show as unassociated servers that are powered off, with no errors.

To view the servers' discovery status, complete the following steps:

1. In Cisco UCS Manager, click the Equipment button on the left-hand side, and click Equipment in the top of the navigation tree on the left.

2. In the properties pane, click the Servers tab.

3. Click the Blade Servers or Rack-Mount Servers sub-tab as appropriate, and view the servers' status in the Overall Status column.
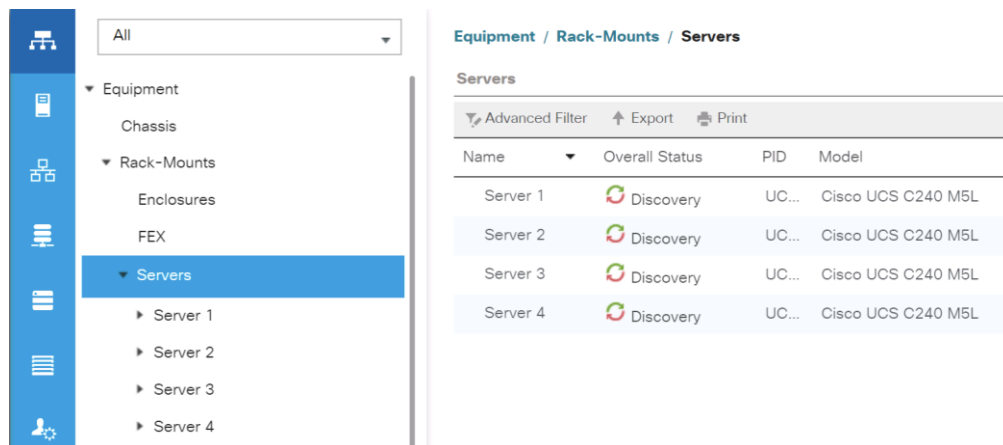
Cisco UCS Organization

Cisco UCS Manager sub-organizations are created underneath the root level of the Cisco UCS hierarchy, which are used to contain all policies, pools, templates and service profiles used by the connected servers. Creating a sub-organization specifically for the Cohesity cluster prevents problems from overlapping settings across policies and pools. This arrangement also allows for organizational control using Role-Based Access Control (RBAC) and administrative locales within Cisco UCS Manager at a later time if desired. In this way, control can be granted to administrators of only the Cohesity specific elements of the Cisco UCS domain, separate from control of root level elements or elements in other sub-organizations.

To create a sub-organization for the Cohesity cluster, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Servers > Service Profiles, right-click root, then click Create Organization.

3. Enter a name for the organization, for example "Cohesity" and optionally enter a description.

4. Click OK.
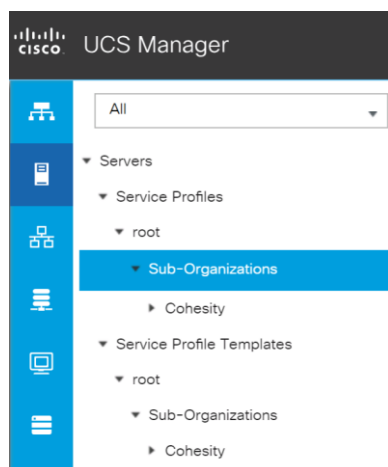
Cisco UCS LAN Policies

VLANs

Names and IDs for the required VLANs must be defined in the Cisco UCS configuration page prior to the Cohesity installation. The VLANs listed in Cisco UCS must already be present on the upstream network, and the Cisco UCS FIs do not participate in VLAN Trunk Protocol (VTP).

To configure the VLAN(s) required for the installation, complete the following steps:

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. In the tree hierarchy, underneath LAN > LAN Cloud, right-click VLANs, then click Create VLANs.

3. Enter a VLAN name which describes the VLAN purpose.

4. Leave the Multicast Policy Name as <not set>.

5. Choose the radio button for Common/Global.

6. Enter the VLAN ID for this VLAN as defined in the upstream switches.

7. Choose the radio button for Sharing Type: None.

8. Click OK.

Table 21and Figure 11 detail the VLANs configured for HyperFlex:

Table 21 Cisco UCS VLANs

Name	ID	Type	Transport	Native	VLAN Sharing	Multicast Policy
<<cohesity-vlan>>	<user_defined>	LAN	Ether	No	None	None

Figure 11   Cisco UCS VLANs
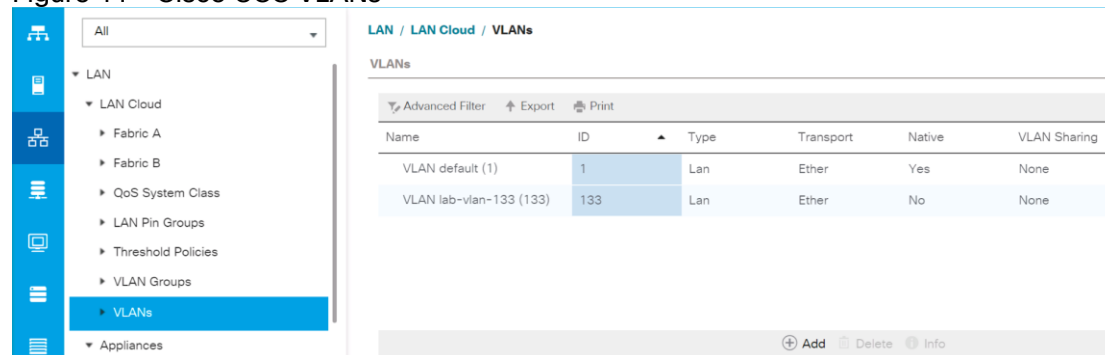
QoS System Classes

By default, Cohesity clusters do not utilize Quality of Service (QoS) policies in their service profiles, and instead place all network traffic into the default "Best-Effort" class. Notably, Cisco HyperFlex clusters are deployed using QoS and a specific configuration for the UCS QoS System Classes is set during installation. Changes to the UCS QoS System Classes require a reboot of both Fabric Interconnects. For this reason, if a single UCS domain is intended to contain both Cisco HyperFlex clusters and Cohesity, it is highly recommended to first deploy the

Cisco HyperFlex cluster(s). This allows the correct QoS system classes to be set without interrupting service to an existing workload, afterwards Cohesity and other systems can be deployed without any additional impacts.

Management IP Address Pool

A Cisco UCS Management IP Address Pool must be populated with a block of IP addresses. These IP addresses are assigned to the Cisco Integrated Management Controller (CIMC) interface of the rack mount and blade servers that are managed in the Cisco UCS domain. The IP addresses are the communication endpoints for various functions, such as remote KVM, virtual media, Serial over LAN (SoL), and Intelligent Platform Management Interface (IPMI) for each rack mount or blade server. Therefore, a minimum of one IP address per physical server in the domain must be provided. The IP addresses are considered to be an "out-of-band" address, meaning that the communication pathway uses the Fabric Interconnects' mgmt0 ports, which answer ARP requests for the management addresses. Because of this arrangement, the IP addresses in this pool must be in the same IP subnet as the IP addresses assigned to the Fabric Interconnects' mgmt0 ports.

To create the management IP address pool, complete the following steps:

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. In the tree hierarchy, underneath Pools > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click IP Pools, then click Create IP Pool.

4. Enter a name for the IP address pool, such as "Cohesity", and optionally enter a description.

5. Click the radio button for Assignment Order: Sequential in order to apply the addresses to the servers in sequence. Choosing Default will result in a random assignment order.

6. Click Next.

7. Click the Add button near the bottom to add a block of IPv4 addresses.

8. Enter the first IP address of the pool in the From: field.

9. Enter the size of the address pool in the Size: field.

10. Enter the correct values for the Subnet Mask, Default Gateway, and Primary and Secondary DNS servers.

11. Click OK.

12. Click Next.

13. In most cases, a pool of IPv6 addresses is not necessary, therefore click Finish.

Figure 12   Management IP Address Pool
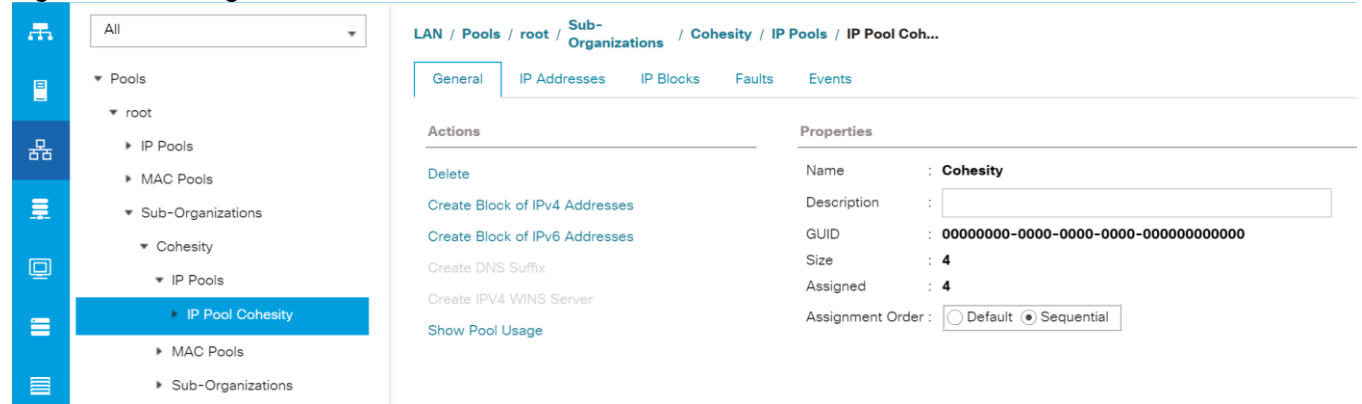
MAC Address Pools

One of the core benefits of the Cisco UCS and Virtual Interface Card (VIC) technology is the assignment of the personality of the card via Cisco UCS Service Profiles. The number of virtual NIC (vNIC) interfaces, their VLAN associations, MAC addresses, QoS policies and more are all applied dynamically as part of the service profile association process. Media Access Control (MAC) addresses use 6 bytes of data as a unique address to identify the interface on the layer 2 network. All devices are assigned a unique MAC address, which is ultimately used for all data transmission and reception. The Cisco UCS and VIC technology picks a MAC address from a pool of addresses, and assigns it to each vNIC defined in the service profile when that service profile is created.

Best practices mandate that MAC addresses used for Cisco UCS domains use 00:25:B5 as the first three bytes, which is one of the Organizationally Unique Identifiers (OUI) registered to Cisco Systems, Inc. The remaining 3 bytes can be manually set. The fourth byte (e.g. 00:25:B5:xx) is often used to identify a specific UCS domain, meanwhile the fifth byte is often set to correlate to the Cisco UCS fabric and the vNIC placement order. Finally, the last byte is incremented upward from the starting value defined, according to the number of MAC addresses created in the pool. To avoid overlaps, when you define these values you must ensure that the MAC address pools are unique for each UCS domain installed in the same layer 2 network.

Cohesity servers running inside the Cisco UCS domain require two vNICs, one in the A side fabric, and one in the B side fabric. To make identification and troubleshooting easier, it is recommended to create two MAC address pools; one for the A side fabric vNICs, and a second for the B side fabric vNICs, each with a unique identifier in the fifth byte.

To create the MAC address pools, complete the following steps:

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. In the tree hierarchy, underneath Pools > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click MAC Pools, then click Create MAC Pool.

4. Enter a name for the MAC address pool, such as "cohesity-mac-a", and optionally enter a description.

5. Click the radio button for Assignment Order: Sequential in order to apply the addresses to the servers in se-quence. Choosing Default will result in a random assignment order.

6. Click Next.

7. Click the Add button near the bottom to add a block of MAC addresses.

8. Modify the values in the 4th byte and 5th byte as necessary in the First MAC Address field. For example, change the field to read "00:25:B5:C0:A1:00"

9. Enter the size of the address pool in the Size: field.

10. Click OK.

11. Click Finish.

12. Repeat steps 1-11 to create any additional MAC address pools required, for example a second pool for the B side vNICs.

Table 22details an example of MAC Address Pools configured for Cohesity and their association to the vNIC templates created afterward:

Table 22 MAC Address Pools

Name	Block Start	Size	Assignment Order	Used by vNIC Template
cohesity-mac-a	00:25:B5:C0:A1:00	32	Sequential	cohesity-vnic-a
cohesity-mac-b	00:25:B5:C0:B2:00	32	Sequential	cohesity-vnic-b

Figure 13   MAC Address Pools
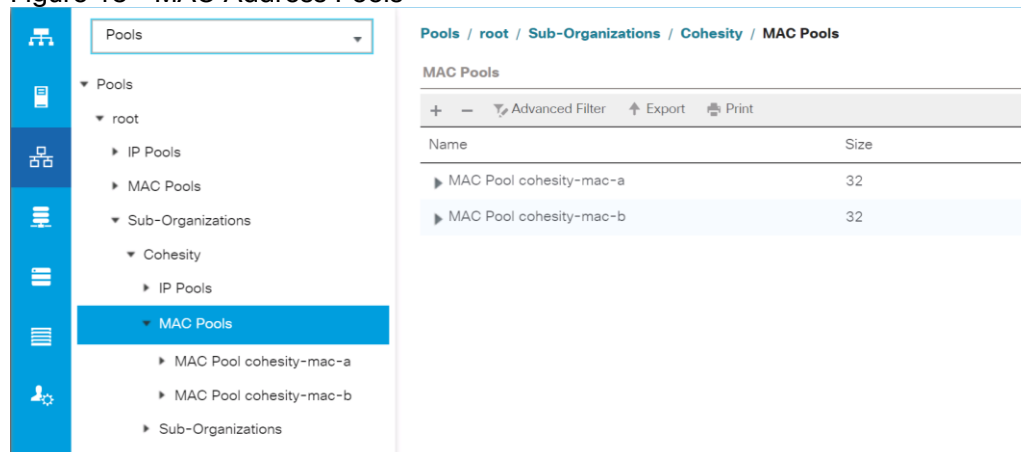
Network Control Policies

Cisco UCS Network Control Policies control various aspects of the behavior of vNICs defined in the Cisco UCS Service Profiles. Settings controlled include enablement of Cisco Discovery Protocol (CDP), MAC address registration, MAC address forging, and the action taken on the vNIC status if the Cisco UCS network uplinks are failed. Of these settings, the most important for the Cohesity DataPlatform is the setting to mark the vNICs as Link Down if there is a failure of all the uplinks from that Fabric Interconnect. This helps ensure that the OS level bonding in the Cohesity nodes will correctly fail over to the other fabric if all uplinks from one FI are lost.

To configure the Network Control Policy, complete the following steps:

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click Network Control Policies, then click Create Network Control Policy.

4. Enter a name for the policy, and optionally enter a description.

5. Click the radio button to set CDP: Enabled.

6. Make sure the setting for Action on Uplink Fail is set to Link Down.

7. All other settings can be left at their defaults.

8. Click OK.

Table 23details the Network Control Policy configured for Cohesity, and the assignment to the vNIC templates created:

Table 23 Network Control Policy

Name	CDP	MAC Register Mode	Action on Uplink Fail	MAC Security	Used by vNIC Template
Cohesity	Enabled	Only Native VLAN	Link-down	Forged: Allow	cohesity-vnic-a cohesity-vnic-b

Figure 14 Network Control Policy

vNIC Templates

Cisco UCS Manager has a feature to configure vNIC templates, which can be used to simplify and speed up configuration efforts. VNIC templates are referenced in service profiles and LAN connectivity policies, versus configuring the same vNICs individually in each service profile, or service profile template. VNIC templates contain all the configuration elements that make up a vNIC, including VLAN assignment, MAC address pool selection, fabric A or B assignment, fabric failover, MTU, QoS policy, Network Control Policy, and more. Templates are created as either initial templates, or updating templates. Updating templates retain a link between the parent template and the child object, therefore when changes are made to the template, the changes are propagated to all remaining linked child objects. An additional feature named "vNIC Redundancy" allows vNICs to be configured in pairs, so that the settings of one vNIC template, designated as a primary template, will automatically be applied

44

to a configured secondary template. For all Cohesity vNIC templates, the "A" side vNIC template is configured as a primary template, and the related "B" side vNIC template is a secondary. In each case, the only configuration difference between the two templates is which fabric they are configured to connect through.

To create the vNIC templates, complete the following steps:

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click vNIC Templates, then click Create vNIC Template.

4. Enter a name for the template, and optionally enter a description.

5. Click the radio button for Fabric ID: Fabric A, and ensure the checkbox for Enable Failover is left unchecked.

6. Click the radio button for Redundancy Type: Primary Template. Leave the Peer Redundancy Template as <not set>.

7. Leave the Target checkbox for Adapter as checked, and for VM as unchecked.

8. Click the radio button for Template Type: Updating Template.

9. In the list of VLANs, click the checkbox next to the VLAN which was created for Cohesity cluster traffic in order to select it, and click the radio button on the right for Native VLAN in order to pass the traffic without VLAN ID tags.

10. Scroll down in the window, ensure that the CDN source is left as vNIC Name, and the MTU is set to 1500.

11. Choose the MAC Address Pool created earlier for the A side fabric for this vNIC.

12. Choose the Network Control Policy created earlier for this Cohesity sub-organization.

13. Click OK.

14. Repeat steps 1-13, but do so for the B side vNIC template, which requires the following changes:

    a. Give the template a unique name for the B side template.
    b. Choose Fabric B for the Fabric ID.
    c. Choose Secondary Template for the Redundancy Type.
    d. Choose the vNIC template just created earlier as the Peer Redundancy Template.
    e. Choose the MAC Address Pool created earlier for the B side fabric.

The following tables detail the initial settings in each of the vNIC templates created for the Cohesity DataPlatform:

Table 24 vNIC Template cohesity-vnic-a

vNIC Template Name:	cohesity-vnic-a
Setting	Value
Fabric ID	A

Fabric Failover	Disabled
Target	Adapter
Type	Updating Template
MTU	1500
MAC Pool	cohesity-mac-a
QoS Policy	<none>
Network Control Policy	Cohesity
VLANs	<<cohesity-vlan>>	Native: Yes

Table 25 vNIC Template cohesity-vnic-b

vNIC Template Name:	cohesity-vnic-b	
Setting	Value	
Fabric ID	B	
Fabric Failover	Disabled	
Target	Adapter	
Type	Updating Template	
MTU	1500	
MAC Pool	cohesity-mac-b	
QoS Policy	<none>	
Network Control Policy	Cohesity	
VLANs	<<cohesity-vlan>>	Native: Yes

LAN Connectivity Policies

Cisco UCS Manager has a feature for LAN Connectivity Policies, which aggregates all of the vNICs or vNIC templates desired for a service profile configuration into a single policy definition. This simplifies configuration efforts by defining a collection of vNICs or vNIC templates once, and using that policy in the service profiles or service profile templates.

To create the LAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click the LAN button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click LAN Connectivity Policies, then click Create LAN Connectivity Policy.

4. Enter a name for the policy, and optionally enter a description.

5. Click the Add button near the bottom to add a VNIC.

6. Enter a name for the vNIC being added, for example vNIC0.

7. Click the Use vNIC Template checkbox.

8. In the vNIC Template drop-down box, choose the A side vNIC template created earlier.

9. Click the Redundancy Pair checkbox.

10. In the Peer Name field, enter a name for the redundant vNIC, for example vNIC1.

11. In the Adapter Policy drop-down box, choose the Linux policy.

12. Click OK.

13. Click OK.

Table 26 details the LAN Connectivity Policy configured for Cohesity:

Table 26 LAN Connectivity Policy

Policy Name	Use vNIC Template	vNIC Name	vNIC Template Used	Adapter Policy
Cohesity	Yes	vNIC0	cohestiy-vnic-a	Linux
		vNIC1	cohestiy-vnic-b	

Cisco UCS Server Policies

BIOS Policies

Cisco M5 generation servers no longer use pre-defined BIOS setting defaults derived from Cisco UCS Manager, instead the servers have default BIOS tokens set from the factory. The current default token settings can be viewed at the following website: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/Reference-Docs/Server-BIOS-Tokens/4-0/b_UCS_BIOS_Tokens_Guide_4_0.html

A BIOS policy must be created to modify the setting of M5 generation servers to enable Serial over LAN communication, which can be used during troubleshooting efforts.

To configure the BIOS policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click BIOS Policies, then click Create BIOS Policy.

4. Enter a name for the policy, and optionally enter a description.

5. Click OK.

6. Click the name of the BIOS Policy which was just created.

7. In the right-hand pane of the Cisco UCS Manager screen, click the Advanced tab.

8. Click the Serial Port sub-tab.

9. Change the Serial Port A enable Value to Enabled in the drop-down list.

10. Click the Server Management tab at the top of the pane.

11. Change the Console Redirection BIOS setting Value to Serial Port A in the drop-down list.

12. Click Save Changes.

Table 27 details the BIOS Policy configured for Cohesity:

Table 27 BIOS Policy

Policy Name	Setting	Value
Cohesity	Serial Port A	Enabled
	Console Redirection	Serial Port A

Boot Policies

Cisco UCS Boot Policies define the boot devices used by blade and rack-mount servers, and the order that they are attempted to boot from. Cisco UCS C-Series M5 generation rack-mount servers which run the Cohesity DataPlatform have their Linux operating system installed to a pair of internal M.2 SSD boot drives, therefore they require a unique boot policy defining that the servers should boot from that location. In addition, a local CD/DVD boot option is included to allow the server to search for the installation ISO media during the Cohesity installation steps.

To configure the Boot Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click Boot Policies, then click Create Boot Policy.

4. Enter a name for the template, and optionally enter a description.

5. Leave all settings at their defaults, ensuring the Boot Mode option is set to Legacy.

6. In the Boot Order area, click the + symbol next to Local Devices to expand the list.

7. Click the blue link for "Add CD/DVD", you will see this selection added to the boot order.

8. Click the blue link for "Add Embedded Local Disk".

9. In the pop-up window, click the radio button for Any, then click OK.

10. Click OK.

Figure 15 details the Cohesity Boot Policy:

Figure 15   Cisco UCS Cohesity Boot Policy

Host Firmware Packages

Cisco UCS Host Firmware Packages represent one of the most powerful features of the Cisco UCS platform; the ability to control the firmware revision of all the managed blades and rack-mount servers via a policy specified in the service profile. Host Firmware Packages are defined and referenced in the service profiles. Once a service profile is associated to a server, the firmware of all the components defined in the Host Firmware Package are automatically upgraded or downgraded to match the package. A Host Firmware Package is created for the Cohesity nodes, which uses the simple package definition method, applying firmware revisions to all components that matches a specific Cisco UCS firmware bundle, versus defining the firmware revisions part by part.

To configure the Host Firmware Package, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click Host Firmware Packages, then click Create Host Firmware Package.

4. Enter a name for the package, and optionally enter a description.

5. Click the radio button for Simple package selection.

6. In the Blade Package and Rack Package drop-down lists, choose the package version that matches the de-sires firmware version. In most cases, the version chosen would match the currently running Cisco UCS Manager and Fabric Interconnect versions, for example, 4.0(1c)B, and 4.0(1c)C.

7. Choose a Service Pack revision if applicable.

8. In the Excluded Components list, make sure all checkboxes are unchecked.

9. Click OK.

The following figure details the Host Firmware Package used for Cohesity:

Figure 16   Cohesity Host Firmware Package

Servers / Policies / root / Sub-Organizations / Cohesity / Host Firmw... / Cohesity

General	Events

Actions

Delete

Show Policy Usage

Use Global

Modify Package Versions

Modify Backup Package Versions

Properties

Name   : **Cohesity**

Description   :

Owner   : **Local**

Blade Package : **4.0(1c)B**   Blade Backup Package :

Rack Package : **4.0(1c)C**   Rack Backup Package :

Service Pack   :

Local Disk Configuration Policies

Cisco UCS Local Disk Configuration Policies are used to define the configuration of disks installed locally within each blade or rack-mount server, most often to configure Redundant Array of Independent/Inexpensive Disks (RAID levels) when multiple disks are present for data protection. Since Cohesity converged nodes providing storage resources utilize software defined storage, the nodes do not require a local disk configuration to be set. Therefore, a simple policy which allows any local disk configuration is all that is required.

To configure the Local Disk Configuration Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click Local Disk Config Policies, then click Create Local Disk Configuration Policy.

4. Enter a name for the policy, and optionally enter a description.

5. Leave all options at their default settings, ensuring the Mode drop-down list is set to "Any Configuration".

6. Click OK.

Figure 17 detail the Local Disk Configuration Policies configured by the HyperFlex installer:

Figure 17   Cohesity Local Disk Configuration Policy

Maintenance Policies

Cisco UCS Maintenance Policies define the behavior of the attached blades and rack-mount servers when changes are made to the associated service profiles. The default Cisco UCS Maintenance Policy setting is "Immediate" meaning that any change to a service profile that requires a reboot of the physical server will result in an immediate reboot of that server. The Cisco best practice is to use a Maintenance Policy set to "user-ack", which requires a secondary acknowledgement by a user with the appropriate rights within Cisco UCS Manager, before the server is rebooted to apply the changes. In addition, the On Next Boot setting is enabled, which will automatically apply changes the next time the server is rebooted, without any secondary acknowledgement.

To configure the Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click Maintenance Policies, then click Create Maintenance Policy.

4. Enter a name for the policy, and optionally enter a description.

5. Click the radio button for Reboot Policy: User Ack.

6. Check the checkbox for On Next Boot.

7. Click OK.

Figure 18 details the Maintenance Policy configured for Cohesity:

Figure 18   Cisco UCS Maintenance Policy

Serial over LAN Policies

Cisco UCS Serial over LAN (SoL) Policies enable console output which is sent to the serial port of the server, to be accessible via the LAN. For many Linux based operating systems, the local serial port can be configured as a local console, where users can watch the system boot, and communicate with the system command prompt interactively. Since many servers do not have physical serial ports, and often administrators are working remotely, the ability to send and receive that traffic via the LAN is very helpful. Interaction with SoL can be initiated by connecting to the CIMC IP address configured by UCS Manager using SSH, and entering valid Cisco UCS manager credentials.

To configure the Serial Over LAN Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click Serial Over LAN Policies, then click Create Serial Over LAN Policy.

4. Enter a name for the policy, and optionally enter a description.

5. Select the radio button for Serial Over Lan State: Enable

6. Select 115200 from the Speed drop-down list.

7. Click OK.

Figure 19 details the SoL Policy configured for Cohesity:

Figure 19   Cisco UCS Serial over LAN Policy

IPMI Access Profile

Cisco UCS Intelligent Platform Management Interface (IPMI) Policies allow for remote interactions with physical hardware resources via the LAN, such as querying power states or forcing servers to power on or off. The Cohesity DataPlatform requires IPMI access to each node, and asks for the IPMI addresses and credentials during the installation. Consequently, and IPMI policy is required to enable the functionality via the CIMC interfaces, and to define the username and password which has access to the IPMI commands.

To configure the IPMI Policy, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Policies > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click IPMI Access Profiles, then click Create IPMI Access Profile.

4. Enter a name for the policy, and optionally enter a description.

5. Click the radio button for IPMI Over LAN: Enable.

6. Click Add to create a user.

7. Enter the username.

8. Enter and confirm the desired password.

9. Click the radio button for Role: Admin.

10. Click OK.

11. Click OK.

Figure 19 details the IPMI configured for Cohesity:

Figure 20   Cisco UCS IPMI Policy

Cisco UCS Service Profile Templates

Cisco UCS Manager has a feature to configure service profile templates, which can be used to simplify and speed up configuration efforts when the same configuration needs to be applied to multiple servers. Service profile templates are used to spawn multiple service profile copies to associate with a group of servers, versus configuring the same service profile manually each time it is needed. Service profile templates contain all the configuration elements that make up a service profile, including vNICs, vHBAs, local disk configurations, boot policies, host firmware packages, BIOS policies and more. Templates are created as either initial templates, or updating templates. Updating templates retain a link between the parent template and the child object, therefore when changes are made to the template, the changes are propagated to all remaining linked child objects.

To configure the Service Profile Template, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Service Profile Templates > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click the sub-org name, then click Create Service Profile Template.

4. Enter a name for the template.

5. Click the radio button for Type: Updating Template.

6. In the UUID Assignment drop-down list, select Hardware Default.

7. Click Next.

8. In the Storage Provisioning section, click the Local Disk Configuration Policy tab, then in the drop-down list below, select the Local Disk Configuration Policy which was created for this template earlier.

9. Click Next.

10. In the Networking section, click the radio button for Use Connectivity Policy, then in the drop-down list below, select the LAN Connectivity Policy which was created for this template earlier.

11. Click Next.

12. In the SAN Connectivity section, click the radio button for No vHBAs, then click Next.

13. In the Zoning section, no changes are required, click Next.

14. In the vNIC/vHBA Placement section, no changes are required, click Next.

15. In the vMedia Policy section, no changes are required, click Next.

16. In the Server Boot Order section, in the Boot Policy drop-down list, select the Boot Policy which was created for this template earlier.

17. Click Next.

18. In the Maintenance Policy section, in the Maintenance Policy drop-down list, select the Maintenance Policy which was created for this template earlier.

19. Click Next.

20. In the Server Assignment section, leave the Pool Assignment set to Assign Later, and select the radio button for the desired power state to Up.

21. Click the + button next to Firmware Management to expand the section. In the Host Firmware Package drop-down list, select the Host Firmware Package which was created for this template earlier.

22. Click Next.

23. In the Operation Policies section, click the + button next to BIOS Configuration to expand the section. In the BIOS Policy drop-down list, select the BIOS Policy which was created for this template earlier.

24. Click the + button next to External IPMI Management Configuration to expand the section. In the IPMI Access Profile drop-down list, select the IPMI Access Profile which was created for this template earlier.

25. In the SoL Configuration Profile drop-down list, select the Serial Over LAN Policy which was created for this template earlier.

26. Click the + button next to Management IP Address to expand the section. Click the Outband IPv4 tab, then in the Management IP Address Policy drop-down list, select the Management IP Address Pool which was created for this template earlier.

27. Click Finish.

The following table details the service profile template configured for the Cohesity DataPlatform nodes:

Table 28 Cisco UCS Service Profile Template Settings and Values

Service Profile Template Name:	cohesity-nodes-m5
Setting	Value
UUID Pool	Hardware Default

Service Profile Template Name:	cohesity-nodes-m5
Setting	Value
Associated Server Pool	None
Maintenance Policy	Cohesity
Management IP Address Policy	Cohesity
Local Disk Configuration Policy	Cohesity
LAN Connectivity Policy	Cohesity
Boot Policy	Cohesity
BIOS Policy	Cohesity
Firmware Policy	Cohesity
Serial over LAN Policy	Cohesity
IPMI Policy	Cohesity

Create Service Profiles

When a Cisco UCS Service Profile Template has been created, individual Service Profiles for each Cohesity node can be created from the template. The unique identifying characteristics of the service profile, such as MAC addresses or IP addresses, are drawn from the pools and the configurations are set according the policies, when the service profile is created. By basing the service profiles on a template, all of the service profiles will have identical configurations. Because the service profiles are based on an updating template, if any errors are found, or changes need to be made to all of the servers, the changes can be made in the parent template, and all child profiles will inherit the change.

To configure the Service Profiles from the Template, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Service Profile Templates > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Right-click the Service Profile Template, then click Create Service Profiles From Template.

4. Enter a naming prefix, which will be applied to all of the spawned service profiles, for example "Cohesity-node-"

5. Enter the starting number for the number to be appended to the name prefix just entered.

6. Enter the number of service profiles to create from this template.

7. Click OK.

Create Service Profiles From Template **?** ✕

Naming Prefix : Cohesity-node-

Name Suffix Starting Number : 1

Number of Instances : 4

OK Cancel

Service Profile Association

When a Cisco UCS Service Profile has been created, it must be associated with a physical hardware asset in order to take effect. Service profile association requires the rack-mount servers or blade servers in the blade chassis to be present, fully discovered, and not currently associated with any other service profiles. Automatic assignment of service profiles can be done through the use of server pools and auto-discovery, but that configuration is not the recommended method for this paper, and therefore not covered in this document. Once the service profile association is initiated, all of the configuration elements and identities are applied to the server hardware, including storage, networking, policies and firmware upgrades. At the conclusion of the association process, the server will be ready for use, but with no operating system installed.

To associate the Service Profiles to the Cohesity node servers, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Service Profiles > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Click the first Service Profile you wish to associate, then in the right-hand pane, click the blue link for "Change Service Profile Association."

4. In the Server Assignment drop-down list, choose "Select existing Server."

5. Ensure the radio button for Available Servers is selected, in the list below you should see the connected and discovered Cisco C240 M5 LFF nodes which have not yet been associated with a service profile.

6. Select the radio button next to the first server to associate, then click OK.

7. Repeat steps 1-6 for each remaining service profile, choosing a subsequent C240 M5 LFF node to associate with.

Associate Service Profile

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

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

Server Assignment: Select existing Server ▼

◉ Available Servers ◯ All Servers

Select	Chassis ID	Slot	Rack ID	PID	Procs	Memory	Adapters
◉			1	UCSC-C2...	2	131072	1
◯			2	UCSC-C2...	2	131072	1
◯			3	UCSC-C2...	2	131072	1
◯			4	UCSC-C2...	2	131072	1

Restrict Migration : ☐

OK Cancel

As previously described, when the service profile association is started, there are many activities that must take place to finalize the configuration of the server. The process can take some time to complete, especially if there are significant firmware updates to perform in order to comply with the policy. Before continuing with the Cohesity installation processes, wait for all of the servers to finish their association process and to show an overall status of OK, with no errors.

To view the servers' discovery status, complete the following steps:

1. In Cisco UCS Manager, click the Equipment button on the left-hand side, and click Equipment in the top of the navigation tree on the left.

2. In the properties pane, click the Servers tab.

3. Click the Blade Servers or Rack-Mount Servers sub-tab as appropriate, and view the servers' status in the Overall Status column.

Cohesity Installation

Cohesity DataPlatform is installed in three phases; first is the initial software installation to all of the Cohesity nodes, followed by the initial network setup of a single node in order to access the Cohesity configuration webpage, and finally the initial Cohesity cluster configuration, which is done from the aforementioned webpage.

Cohesity Software Installation

The installation of Cohesity DataPlatform software is done via a bootable DVD ISO image file. Each node is booted from this image file, which will automate the process of installing the underlying Linux operating system, copy the Cohesity software packages, and prepare the nodes for the initial setup of the Cohesity cluster.

To install the Cohesity software on each Cisco UCS C240 M5 node, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Service Profiles > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Each Cohesity node will have its own service profile, for example: Cohesity-node-1. Right-click the first service profile and click KVM Console. The remote KVM console will open in a new browser tab. Accept any SSL errors or alerts, then click the link to launch the KVM console.
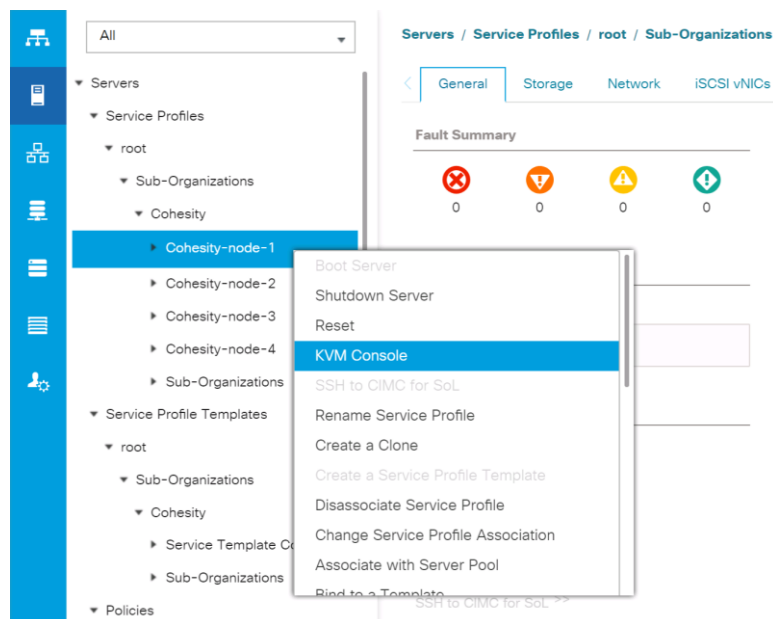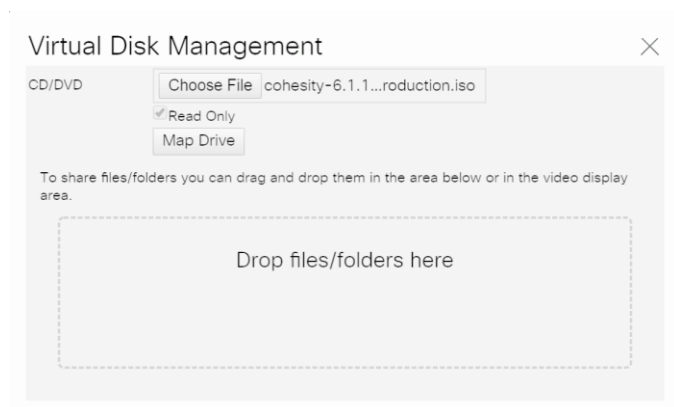

4. In the remote KVM tab, click the Virtual Media button in the top right-hand corner of the screen, then click Activate Virtual Devices.
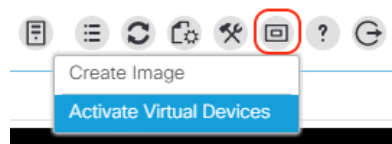
5. In the remote KVM tab, click the Virtual Media button in the top right-hand corner of the screen, then click the CD/DVD option.
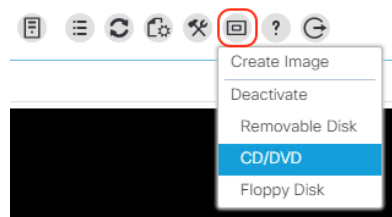
6. Click Choose File, browse for the Cohesity ISO installer file, and click Open.

7. Click Map Drive.

8. In the remote KVM tab, click the Server Actions button in the top right-hand corner of the screen, the click Reset.
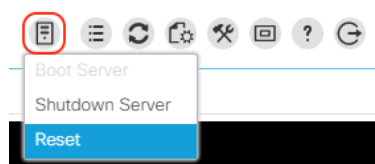
9. Click OK.

10. Choose the Power Cycle option, then click OK.

11. Click OK.

12. Observe the server going through the POST process until the following screen is seen. When it appears, press the F6 key to enter into the boot device selection menu.
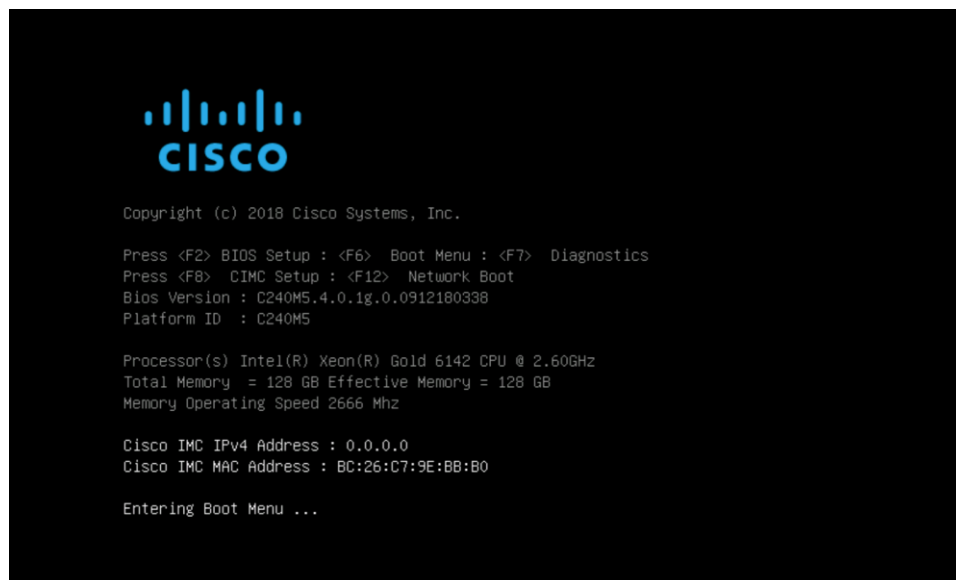
13. Select Cisco vKVM-mapped vDVD1.24, then press Enter.
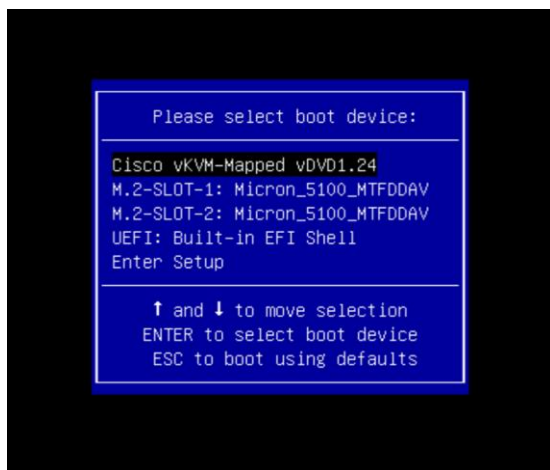
14. The server will boot from the remote KVM mapped Cohesity ISO installer and display the following screen:
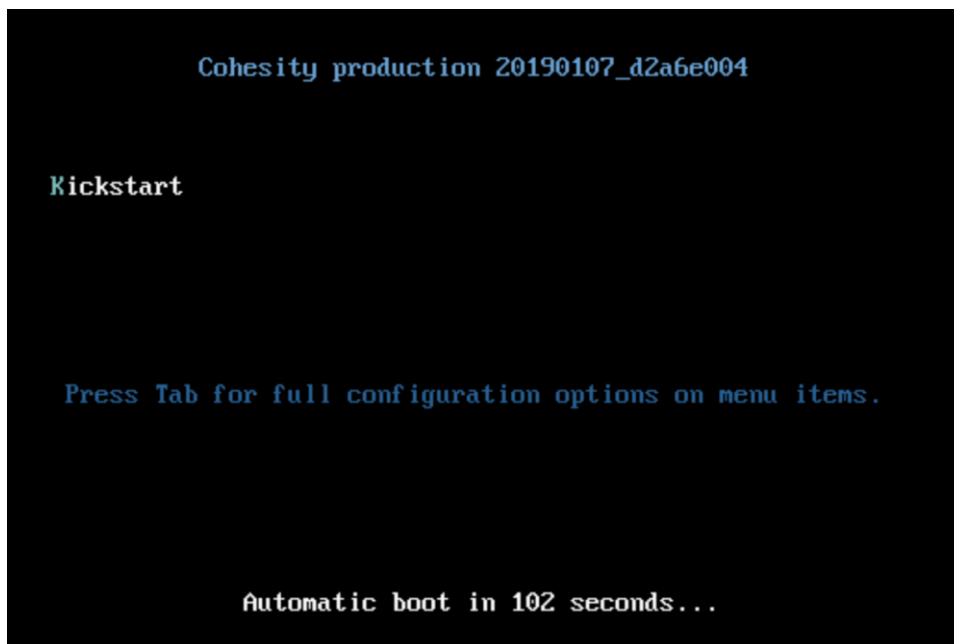
15. Allow the automatic timer to count down from 120 seconds, or press Enter.

16. The Cohesity installer will now automatically perform the installation to the boot media. Installation time takes approximately 70-75 minutes. Once the new server has completed the installation, the server will reboot and it will be waiting at the console login prompt screen seen below.

17. In the remote KVM tab, click the Exit button, then click Yes.

18. Repeat steps 3-17 for each additional Cohesity node being installed.

Cohesity First Node Configuration

In order to perform the initial cluster setup, the first node of the Cohesity cluster must be accessible via the network, so that the administrator performing the configuration can access the Cohesity configuration webpage running on that node. Cohesity nodes will automatically configure themselves with IPv6 link-local addresses, and use these addresses to discover each other on the same subnet. These IPv6 addresses can also be used to perform the initial configuration via the webpage. However, many environments are not configured to use IPv6, therefore it is more common to use IPv4 addresses to perform the initial configuration. To use IPv4 addresses, the first node must be manually configured with an IPv4 address, so that the webpage is accessible to the administrator's client computer.

To manually configure the first Cohesity node's IPv4 addressing, complete the following steps:

1. In Cisco UCS Manager, click the Servers button on the left-hand side.

2. In the tree hierarchy, underneath Service Profiles > root > Sub-Organizations, click the carat next to the name of the sub-organization created for Cohesity to expand it.

3. Each Cohesity node will have its own service profile, for example: Cohesity-node-1. Right-click the first service profile and click KVM Console. The remote KVM console will open in a new browser tab. Accept any SSL errors or alerts, then click the link to launch the KVM console.

4. When the node's local login prompt appears, login with the following credentials:

    a. Username: cohesity

    b. Password: Cohe$1ty

5. Edit the configuration file for the bond0 interface using a text editor, such as vi, entering the following lines, and their corresponding values:

    ◢ Note: Using sudo is required for root privileges.

    ```
    sudo vi /etc/sysconfig/network-scripts/ifcfg-bond0

    IPADDR=<ip>
    NETMASK=<mask>
    GATEWAY=<gw>
    ```

6. Restart the network services of the node:

    ```
    sudo systemctl restart network
    ```

7. Test the network is working properly by pinging the default gateway. You can also verify the IP address configuration by issuing the following command:

    ```
    ip addr
    ```

8.  Log out of the node:

    ```
    exit
    ```

9. In the remote KVM tab, click the Exit button, then click Yes.

Cohesity Cluster Setup

The initial setup of the Cohesity cluster is done via the configuration webpage, which is now accessible on the first node, at the IP address which was configured in the previous steps. Prior to beginning the initial cluster configuration, ensure that all of the Cohesity nodes which are to be included in the cluster have completed their initial software installation, and are fully booted. Additionally, ensure that all of the necessary IP addresses for all of the interfaces are known and assigned, and the DNS round-robin entries have been created.

To perform the Cohesity initial cluster configuration, complete the following steps:

1. In a web browser, navigate to the IP address of the first Cohesity node, which was just configured in the previous steps. For example: http://10.29.133.227

2. Accept any SSL warnings or errors due to the default self-signed certificate on the server, and proceed to the Cohesity Dashboard login screen.

3. Log into the Cohesity Dashboard webpage using the credentials below:

   a. Username: admin

   b. Password: admin
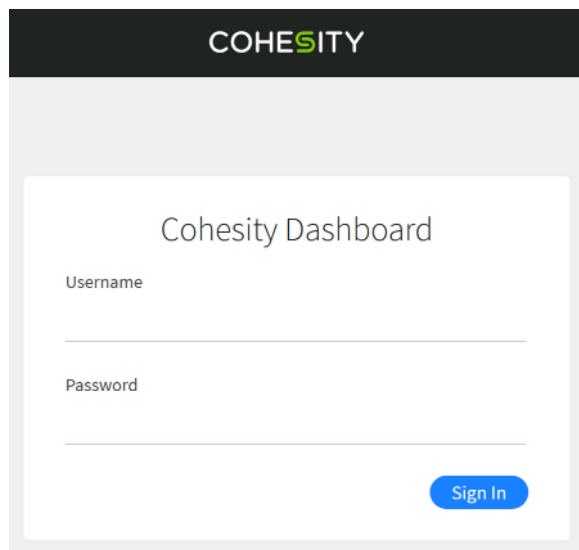
4. The Start Initial Cluster Setup screen appears, make sure that the number of nodes detected matches the number of servers you intend to install for this cluster. Click Get Started.

5. Select the nodes to add to this initial cluster, or click the link to Select All Available in the upper right-hand corner, then click Select Nodes.

6. For each server, enter the IP address which will be assigned to the Linux OS into the IP field.

> ⚠ Note: The servers may not be listed in order, please refer to Cisco UCS Manager to ensure that you are entering the IP addresses in an order that corresponds to the servers and service profiles. The Cohesity installation screen lists the serial numbers of the servers, which can be cross-referenced with the Equipment > Servers view in Cisco UCS Manager.

7. For each server, enter the IPMI address (CIMC Management IP Address) in the IPMI IP field.

> **Note:** The servers may not be listed in order, please refer to Cisco UCS Manager to ensure that you are entering the IP addresses in an order that corresponds to the servers and service profiles. The Cohesity installation screen lists the serial numbers of the servers, which can be cross-referenced with the Management IP of the Service Profile in Cisco UCS Manager.

8. Click Continue to Cluster Settings.

9. Enter the desired name of the cluster and the DNS domain suffix.

10. Enter the gateway IP address and subnet mask for the IP addresses being assigned to the OS and VIPs of the nodes.

11. Enter the subnet mask and gateway address of the subnet where the nodes' IPMI interfaces are configured.

> **Note:** This is the subnet mask and gateway for the IP subnet used by the CIMC interfaces, also called the external management IP addresses.

12. Enter the username and password for IPMI access, to match the username and password configured in the Cisco UCS Manager IPMI Access Profile, which was configured earlier.

13. Enter the required NTP server addresses, separated by commas.

14. Enter the hostname for the Cohesity cluster partition. This hostname typically matches the name of the cluster.

15. Enter the starting IP address for the VIP addresses that are being assigned to the Cohesity nodes. These IP addresses are the addresses which are resolved by DNS round-robin for the cluster, not the individual node IP addresses. For example: 10.29.133.231

16. Enter the last octet value for the end of the VIP range, for example: 234

17. Click Add VIP or VIP Range.

18. Optionally choose to enable systemwide encryption by toggling the switch. Encryption can be enabled at a later time for each separately configured storage domain. Because the latter option provides more flexibility, it is not recommended to enable systemwide encryption at this time, as this choice cannot be reversed.

19. Click Create Cluster.

20. Observe the cluster creation status. Additional details can be viewed by clicking Show Details.
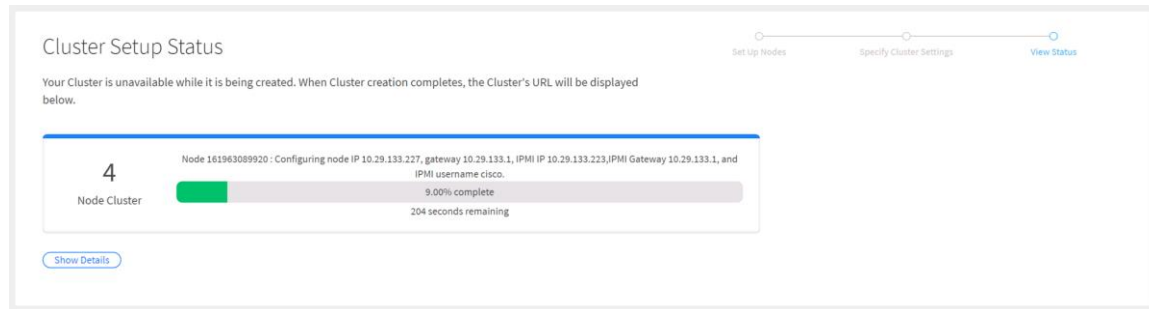
21. The status will appear to pause at 98-99% for a significant period of time while formatting the disks. The time to format nodes with 10 TB capacity disks will be longer than the time for nodes with 4 TB capacity disks. The time to create the cluster for a 4-node cluster with 10 TB disks is approximately 40 minutes.

22. After the setup completes, the web services will restart. After a few minutes, the Cohesity Dashboard webpage for the cluster will be available at the DNS round-robin address configured for the cluster. For example: https://chx-cluster01

Cohesity Software

Cohesity Dashboard

The primary management interface for the Cohesity DataPlatform is the embedded Cohesity Dashboard web interface. All cluster configurations, policies, jobs, and activities can be created, modified and monitored via the Cohesity Dashboard.

To log into the Cohesity Dashboard, complete the following steps:

1. In a web browser, navigate to the DNS round-robin name of the Cohesity cluster, which was just set up in the previous steps. For example: https://chx-cluster01

2. Accept any SSL warnings or errors due to the default self-signed certificate on the server, and proceed to the Cohesity Dashboard login screen.

3. Log into the Cohesity Dashboard webpage using the credentials below:

    a. Username: admin
    b. Password: admin
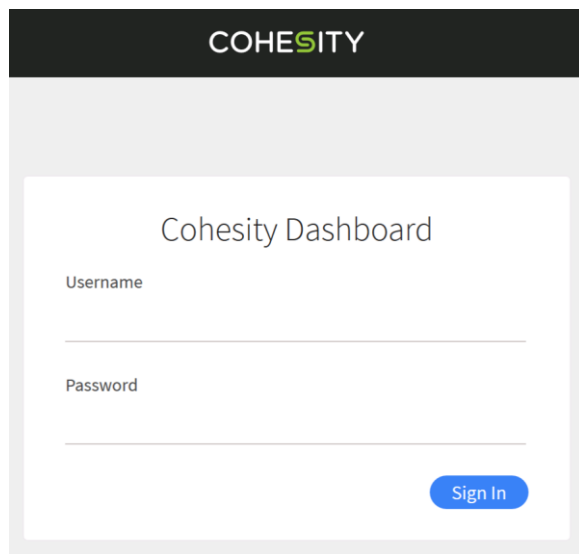
4. Accept the End User Licensing agreement by clicking Accept.

5. Enter a valid Cohesity license code.

6. Click Submit.

Cluster Configuration

After the initial Cohesity cluster setup has completed, the system is operating with a preset collection of default settings. In order to tailor the cluster to your individual needs, these various settings should be modified.

Partitions

Conceptually, a Cohesity cluster is a collection of nodes running the Cohesity software, which provide storage resources. A Cohesity cluster contains a single partition, therefore the partition is synonymous to the overall cluster. There are some settings which are modified as an overall cluster, and some advanced settings that can be configured at the partition level, such as custom host mappings, or VLANs. These settings are not required to be modified in the example configuration detailed in this document, and in most circumstances, the default partition, named "DefaultPartition" can be left unmodified.

Storage Domains

Storage Domains represent a subdivision of the default partition, and many settings can be modified at the Storage Domain level. In particular, settings for deduplication, compression, encryption, and data replication can all be controlled individually for each Storage Domain that is created. Protection Jobs and Views all target a specific Storage Domain in their configurations. This arrangement provides additional flexibility, because through the use of multiple domains, which are then targeted by different jobs and workloads, each of them can be tailored to meet the requirements of that job or workload.

To modify the default Storage Domain of the Cohesity cluster, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Platform menu at the top of the screen, click Cluster.

3. Click the Storage Domains tab.

4. Click the ellipses to the far right-hand side of the line item of the "DefaultStorageDomain", then click Edit.

5. Modify the name of the domain to one which helps to identify the usage and/or settings of this domain, for example, "domain_rf2_id_ic" which indicates a domain using a data replication factor of 2 for redundancy, plus inline deduplication, and inline compression.

6. Modify the settings of this domain for deduplication, compression, encryption, quotas, and failure tolerance as required.

7. Click Update Storage Domain.

To create additional Storage Domains with unique settings from the default domain, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Platform menu at the top of the screen, click Cluster.

3. Click the Storage Domains tab.

4. Click Add Storage Domain.

5. Enter the name of the domain, choosing one which helps to identify the usage and/or settings of this domain, for example, "domain_ec2_id_ic" which indicates a domain using erasure coding with 2 stripes for redundancy, plus inline deduplication, and inline compression.

6. Modify the settings of this domain for deduplication, compression, encryption, quotas, and failure tolerance as required.

7. Click Create Storage Domain.

Time zone

By default, a newly installed Cohesity cluster operates in the UMT (GMT+0) time zone. Organizations have many different standards for which time zone their hardware is configured to operate in. Some choose to have all hardware remain in UMT, others have all hardware operate in a single time zone regardless of physical location, meanwhile many always choose the local time zone of the hardware's physical location. Reliable NTP servers are required and defined during the cluster installation. Ensure that the NTP servers listed are accessible from the network where the Cohesity cluster is installed.

To modify the time zone of the Cohesity cluster, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Admin menu at the top of the screen, click Cluster Settings.

3. Click the link labeled "Change Time zone".

4. From the drop-down list, select the appropriate time zone for this cluster.

5. Click Save.

Security

A newly installed Cohesity cluster has a single administrative user named "admin" with a default password. No additional external authentication sources are configured during the installation. At a minimum, the default admin user's password should be modified away from the default.

Active Directory

Integration with Microsoft Active Directory allows for Active Directory user accounts to log into the Cohesity cluster to administer and use the system, plus it enables additional options to be selected and modified when creating Views that use the SMB protocol.

To join the Cohesity cluster to an Active Directory domain, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Admin menu at the top of the screen, click Access Management.

3. Click the Active Directory tab.

4. Click the link for Join Domain.

5. Enter the Active Directory domain name.

6. Enter a username and password for a user with administrative rights to join computers to the Active Directory domain.

7. Optionally, enter a specific Organizational Unit where the computer account should be located in the Active Directory hierarchy.

8. The Cohesity cluster name will automatically be listed as the machine account to be created. As long as the DNS round-robin records have been properly created prior to the Cohesity cluster installation, this is the only machine account name that is necessary.

9. Click the Add Active Directory button.

* Domain Name

| hx.lab.cisco.com |

* AD Admin

| administrator |

Specify in format **username** or **username@domain.com**

* Password

| •••••••• |

Note that the Active Directory Username and Password are not stored on the Cohesity Cluster.

Organizational Unit

| |

Specify in format **OUName** or **OUName/SubOUName**

Workgroup / NetBIOS Name

| |

* Machine Accounts

| chx-cluster01 ✕ |

Provide the unique name(s) to identify the Cluster on the domain. Separate multiple Machine Accounts with commas, e.g., machine1, machine2

[Add Active Directory]   (Cancel)

Users and Roles

Additional local or Active Directory users can be created to allow them to log into the Cohesity system, and perform tasks according to the roles assigned to their account, which are defined in the cluster.

To add an authorized user to the Cohesity cluster, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Admin menu at the top of the screen, click Access Management.

3. Click Add Users/Groups.

4. Select the radio button for either a Local User or an Active Directory User & Groups.

5. For a local user, enter the username, email address, password, and select a Role from the drop-down list.

6. For an Active Directory User or Group, select the Active Directory Domain from the drop-down list, search for and select either the AD user or group you wish to add, then select a Role from the drop-down list.

7. Click Add.

Type

○ Local User   ● Active Directory Users & Groups

* Active Directory Domain

    hx.lab.cisco.com   ▾

* User(s) or Group(s)

    👥 Domain Admins (hx.lab.cisco.com) ✕   ▾

* Roles

    Admin ✕   ▾

◯ Restrict access to specific Objects

Description

Add   (Cancel)

Passwords

Prior to changing the default System Admin account password, it is highly recommended to create at least one additional local, Active Directory, or LDAP user account with the Admin Role as outlined above. Using this secondary administrative account to make the password change ensures that administrators do not accidentally get locked out of the cluster due to a faulty password change.

To change the default System Admin password, complete the following steps:

1. Log into the Cohesity Dashboard web page as a user with the Admin Role.

2. From the Admin menu at the top of the screen, click Cluster Settings.

3. Toggle the radio switch for "Change System Admin Password".

4. Enter and confirm the new password for the local admin account.

5. Click Save.

Sources

Cohesity can provide protection for multiple platforms, including virtual machines running across a variety of hypervisors, cloud-based virtual machines, bare-metal servers, Oracle and Microsoft SQL databases, plus direct protection of storage array volumes. Cohesity protection jobs are each configured to target specific sources, therefore prior to configuring the protection jobs the sources must be registered with the Cohesity cluster.

Hypervisor Source

Protection of a VMware ESXi based cluster is conducted via the managing VMware vCenter Server. Cisco HyperFlex clusters running on VMware ESXi hypervisors also use vCenter for management of the cluster. To configure protection of a Cisco HyperFlex ESXi based cluster, the managing vCenter server must be registered as a source for the protection jobs.

To configure a hypervisor source, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Sources.

3. Click the Register button and from the drop-down list that appears, click Hypervisor.
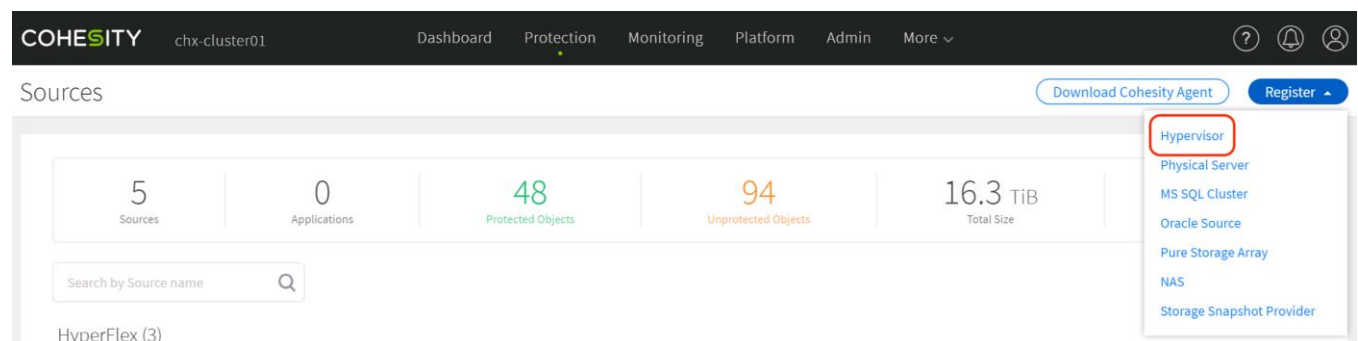
4. From the Hypervisor Source Type drop-down list, choose VMware: vCenter

5. Enter the hostname or IP address of the vCenter server managing the Cisco HyperFlex cluster being pro-
    tected, and an administrative username and password.

6. Toggle the radio button on for the "Auto Cancel Backups if Datastore is running low on space" option, and
    enter a minimum value of free space for the datastore. If the datastore housing the VMs being snapped drops
    below this amount of free space, the job will be automatically cancelled.

7. Click Register.

Storage Snapshot Provider

The Cohesity DataPlatform offers integration with storage-based snapshots, leveraging the native snapshot
technologies built directly into the storage arrays, versus using the standard VMware based virtual machine
snapshots. Cisco HyperFlex offers native storage-based snapshots, which provide space-efficient and crash-
consistent snapshots taken by the underlying Cisco HyperFlex Distributed Filesystem, instead of standard VMware
redo-log based snapshots. By using this integration via the Cisco HyperFlex API, the Cohesity protection jobs will
take Cisco HyperFlex native snapshots instead of VMware snapshots. Cohesity protection jobs will always fall
back to taking VMware native snapshots in case the HyperFlex native snapshot was not available. In order to use
the Cisco HyperFlex API to create native snapshots, the Cisco HyperFlex cluster(s) must be registered as a
Storage Snapshot Provider source.

In order for Cohesity Protection Jobs to always use native HX snapshots of the VMs running in the Cisco
HyperFlex cluster(s), it is important that the VMs to be protected not have any existing standard VMware redo-log
based snapshots. An existing VMware snapshot will prevent the creation of a subsequent HX native snapshot, and
instead all snapshots taken by the Cohesity system will continue to be VMware snapshots. In this situation, prior to
configuring Cohesity Protection Jobs it is recommended to delete all existing VMware snapshots from the virtual
machines running in the Cisco HyperFlex cluster(s), which will be protected by Cohesity using the Storage
Snapshot Provider integration. For VMs which already have existing HX native snapshots, no action is necessary,
because subsequent snapshots taken by the Cohesity system using the Storage Snapshot Provider integration
will continue to be HX native snapshots.

To configure Cisco HyperFlex as a Storage Snapshot Provider source, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Sources.

3. Click the Register button, and from the drop-down list that appears, click Storage Snapshot Provider.
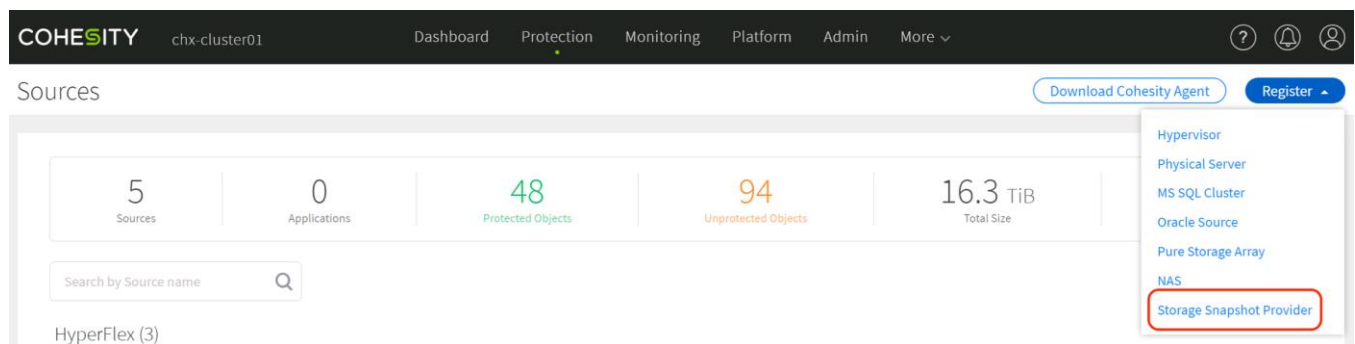
4. From the Snapshot Storage Provider Type drop-down list, choose Storage Snapshot Provider: Hyperflex

5. Enter the hostname or IP address of the roaming management interface of the Cisco HyperFlex cluster being protected, and an administrative username and password. This must be the roaming or floating management IP address, not the management IP address of any individual Cisco HyperFlex node.

6. Click Register.

Remote Clusters

When multiple Cohesity clusters are available across the landscape, the Cohesity clusters can be registered with one another for both remote management, and more importantly, replication of backed up snapshots across the network. Utilizing this feature provides a secondary copy of the snapshots in a different physical cluster, which can be located in a standby datacenter used for disaster recovery. This secondary Cohesity cluster can have a standby VMware vCenter server registered as a source, and backups can quickly be restored to this recovery system in case a disaster is declared, or a planned failover to the secondary system is required. In order to replicate snapshots, the originating cluster (i.e. the cluster which captures the snapshot) must register the receiving cluster, and in return, the receiving cluster must register the originating cluster. A pairing is established between Storage Domains in the two clusters. A single Storage Domain in the originating cluster is paired with a single Storage Domain in the receiving cluster. A many-to-one pairing can be done only across multiple originating clusters, each one pairing a single Storage Domain, but all of them paired with the same receiving Storage Domain. Replication frequency and retention is controlled as part of the Cohesity policies, which each Protection Job is then assigned to follow.

To register remote clusters, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Remote Clusters.

3. Click the Add Cluster button.

4. Enter one or more of the Virtual IP addresses of the remote cluster nodes.

5. Enter the username and password of a user with administrative rights in the remote cluster, then click Connect.

6. Toggle the switches to enable Replication, and Remote Access if desired.

7. Click the link to Add Storage Domain Pairing, and choose the local and remote Storage Domains to pair, then click Add.

8. Optionally, select to enable load distribution, outbound compression, data encryption, transfer speed limits and overrides as needed.

9. Click Create.

10. Repeat steps 1 through 9 on the second Cohesity cluster, registering the cluster in the opposite direction of the first registration.

External Targets

In addition to replication of snapshots between multiple Cohesity clusters, Cohesity can be configured to copy snapshots to non-Cohesity locations, which is referred to as Archiving. Archival is controlled via the Cohesity policies in the same manner as Replication, and as with Replication, the External Target must be configured as a possible location for the archival task. Multiple types of External Targets are available, including Google, Amazon Web Services, Microsoft Azure, Network Attached Storage (NAS), and more. As such, this document will not describe all of the options available for configuration of these External Targets.

Policies

Cohesity policies define the backup types, frequency, retention periods, replication and archival options for protection jobs. Three standardized policies are included by default during the installation, however in many cases these options will need to be customized for your specific use. The default policies can be edited; however, it is recommended to make a copy of one of the default policies to use as a starting point. Alternatively, a new policy can be created from scratch.

To configure Cohesity policies, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Policy Manager.

3. Click the Create Policy button.

    a. Alternatively, click the ellipses next to an existing policy and click Edit Policy.

      b.   In addition, you may click the ellipses next to an existing policy and click Copy Policy. The subsequent policy may then be edited.

4.   Enter a name for the Policy being created or edited, and optionally enter a description.

5.   Toggle the DataLock radio button option if desired. DataLock will prevent snapshots from being removed from the cluster, even if the protection job is deleted, until the snapshot has exceeded its retention period. This allows for additional policy compliance and data protection against unintended deletion of snapshots.

6.   Configure the backup schedule to define the frequency of the job, along with the retention period for the snapshots to be kept.

7.   The default option is for every backup job since the first to be run as an incremental job, saving space and network bandwidth. The option can be modified to perform a full backup at regular intervals if desired.

8.   Extended retention can be configured to keep specific regular snapshots for longer retention periods than the standard retention schedule, if desired.

9.   Blackout periods can be configured, to prevent new runs of a job configured with this policy from starting if the current time is within the blackout period.

10. Optional: If there are multiple Cohesity clusters available, the clusters can be registered with each other and then configured to replicate the snapshots being backed up. Configure replication of the backups to the remote Cohesity cluster if applicable.

11. Optional: If External Targets are available, such as cloud providers, storage arrays and object-based S3 storage systems, these external targets can act as archival locations for off-site storage of Cohesity backups. Configure archival of the backups to the remote target if applicable.

12. Click Create or Save as applicable.

Protection

Protection Jobs are configured in the Cohesity Dashboard to back up the configured sources, according to the Policies defined in the system. Each protection job obtains data from a single Source, operates according to the settings in a single Policy, and targets a single Storage Domain to store the snapshots. Because of this operational method, in order to back up virtual machines according to different schedules, or to target a different Storage Domain, a unique Protection Job must be created for each case. In the same way, backing up virtual machines from different sources, such as multiple Cisco HyperFlex clusters, or combinations of other sources, must be done in a distinct Protection Job per unique source.

During a Cohesity Protection Job, a new snapshot of the virtual machine is taken, and that snapshot is transferred via the network to the Storage Domain configured in the job. This constitutes a new incremental backup of that virtual machine. Once the snapshot is transferred, the snapshot of the virtual machine is deleted in the source hypervisor node. If the virtual machine being backed up was already running with an active snapshot, the new snapshot taken by Cohesity will be a child of the existing snap, then it will be deleted, coalescing the changes back into the existing snapshot level where the VM was already running. If the Storage Snapshot Provider integration with Cisco HyperFlex is enabled, then all of these snapshots will be taken as HX native snapshots. If the HX native snapshot attempt should fail, for example when an existing VMware standard redo-log snapshot exists, then the Protection Job will fall back to taking a standard VMware snapshot.

Of special note is a circumstance where a virtual machine has multiple snapshots, but the virtual machine has been reverted to a previous snapshot and is therefore not running as the most recent snapshot. Cohesity Protection Jobs will take the snapshot at the level where the VM is currently running, therefore, any changes contained in a child snapshot that is not the current running snapshot of the VM, will not be captured in the Cohesity backup. The existence of unused child snapshots will cause warnings during the execution of a Cohesity Protection Job, and such unused snapshots should be removed.

⚠ **Warning: When configuring Protection Jobs of Cisco HyperFlex clusters, it is critical that the HyperFlex Storage Controller Virtual Machines (SCVMs, which start with VM name stCtlVM-*) are not configured to be protected. Taking snapshots of the SCVMs and attempting to restore them is not a supported operation in Cisco HyperFlex clusters.**

To create a Protection Job, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Protection Jobs.

3. Click the Protect button, from the drop-down list that appears, click Virtual Server.

4. Enter a name for the job, and optionally enter a description.

5. Select a source for the job from the drop-down list.

6. From the pop-up window that appears, choose the virtual machine(s) that you wish to protect in this job.

   a. The three buttons on the top right of the window can be used to switch between a hierarchical inventory view, a folder view, or a list view of the virtual machines.

   b. The list of virtual machines may be searched for a specific VM name or names using the wildcard character (*) by clicking in the field next to magnifying glass button.

   c. The list of virtual machines can also be filtered using the drop-down list, for example choosing to show only VMs which are currently unprotected.

d. Individual VMs can be chosen by clicking the checkbox next to their name or it is possible to protect entire clusters, hosts, or folders by clicking the checkbox next to them.
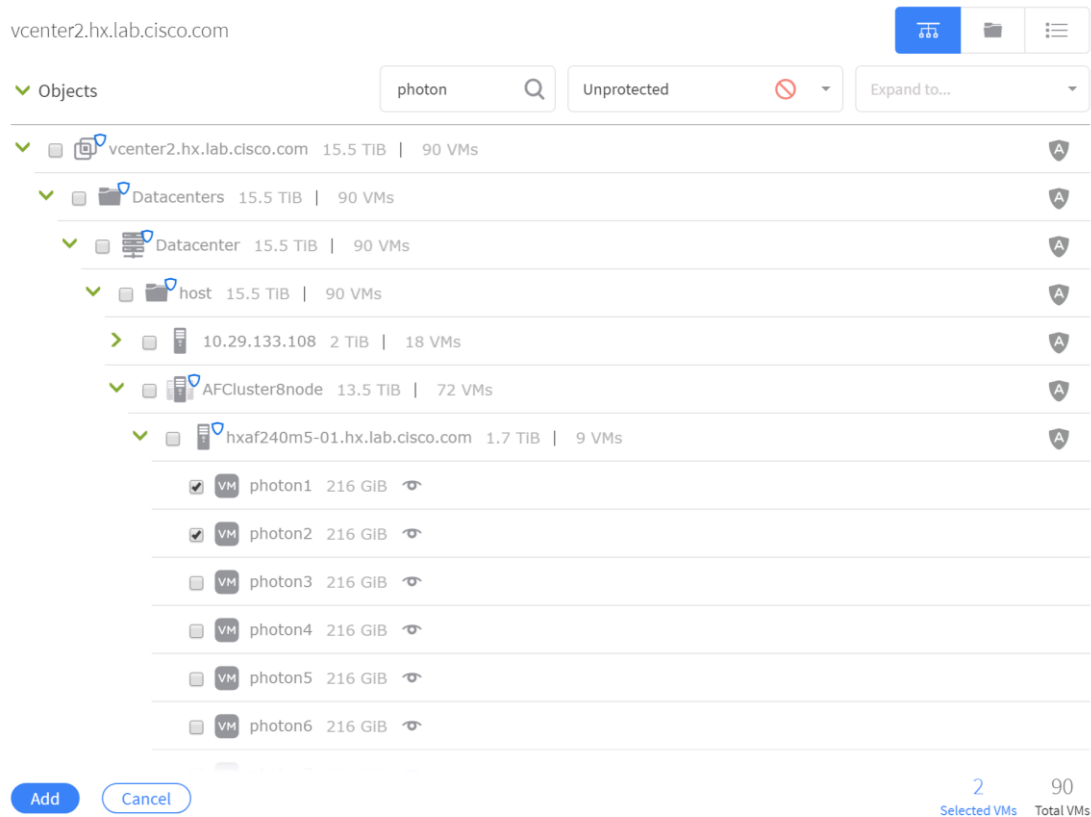

7. Click Add.

8. Select a policy from the drop-down list.

9. Select a storage domain from the drop-down list.

10. In order to take advantage of the Storage Snapshot integration with Cisco HyperFlex clusters, in the Advanced section, click the Edit link next to "Leverage Storage Snapshots for Data Protection." Toggle the radio button on, and select Hyperflex from the drop-down list that appears.

11. Enable App Consistent Backups if necessary, for example if the VM is running transactional software or databases which require the VM's filesystem to be quiesced as part of the backup. In the Advanced section, click the Edit link next to "App Consistent Backups." Toggle the radio button on, and optionally choose to fall back to a standard crash consistent snapshot, without quiescing the filesystem, should the quiesce operation fail.

> ⚠ **Warning: Guest VMs must be running the most current version of VMTools in order for quiesced snapshots to be taken properly, and application consistent backups to be performed.**

12. Optionally, modify the remaining job parameters as required. For example, modify the End Date to stop the job from running after a certain date, or change the QoS Policy in order to force the backup job to use the SSDs in the nodes instead of the HDDs.

13. Click Protect.
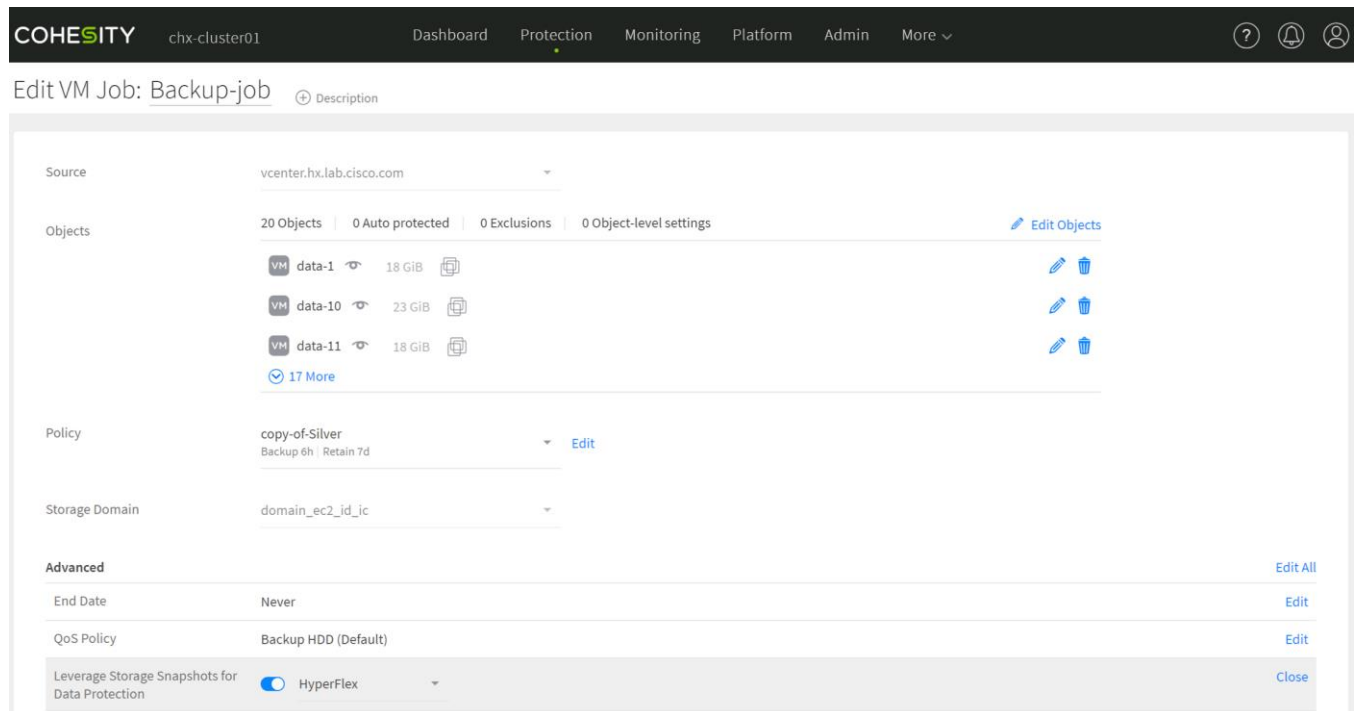
The newly configured protection job will perform its initial run immediately, unless it is currently within a blackout period, and the job will repeat itself according to the schedule set forth in the policy.

Recovery

Recovery jobs can be initiated to restore a virtual machine from the backed-up snapshots and return the virtual machine to service. A unique aspect of the Cohesity software is the sequence of the recovery process. When a recovery job is started, the Cohesity system will present an NFS-based datastore from itself, which is mounted to the ESXi host, inside of which are the virtual machine files that have been bloomed from the snapshots. The virtual machine will then be registered in vCenter from this location, and the VM will be powered on. This process returns the recovered VM to service much faster than typical recovery processes will, because the VM will immediately run with its virtual files sourced from the Cohesity NFS datastore. After the VM is powered on, a storage vMotion will relocate the virtual machine files to their original location. The benefit of this recovery workflow is amplified when multiple simultaneous VM recoveries are needed, because the time to return the VMs to service is very low, and the remaining process of relocating the VMs via storage vMotion happens in the background while the VMs are already online. A recovered VM will have no snapshots, even if the VM originally had snapshots at the time of the backup which is being restored.

Recovery jobs can be used to restore multiple virtual machines at one time, however there are two notable limitations to the restoration of multiple VMs. First, in order to restore multiple VMs in a single job, all of the VMs must have originated from the same source. Second, all of the VMs must have been backed up to the same Cohesity storage domain. For example, if some VMs are protected and targeted a storage domain using replication factor 2, meanwhile others were backup up to a storage domain using erasure coding, a single recovery job could only restore the VMs in the storage domain using replication factor 2, and could not recover the VMs in the storage domain using erasure coding. In order to recover both sets of VMs, two recovery jobs would need to be started. Similarly, if VMs originated from multiple sources then multiple jobs must be created to restore them.

Multiple recovery jobs can be created and run simultaneously when the above scenarios apply.

To initiate a recovery operation, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Recovery.

3. Click the Recover button, from the drop-down list that appears, click VMs.

4. In the search field, search for the name of the VM or VMs that need to be recovered. Wildcard characters can be used, and additional filters can be applied.

5. Check the checkbox next to the name(s) of the VMs you wish to recover.

6. Steps 4 and 5 can be repeated multiple times to select different VMs, for example searching for the name of one VM, selecting it, then clearing the search field, searching for the name of another VM, and then selecting that one as well.

7. Once all the desired VMs have been selected, click Continue.

8. The pencil icon next to each VM can be clicked to select the snapshot used as the recovery point. By default, the latest snapshot will be chosen. Choose the desired recovery point in the pop-up window and click Save.

Selected Objects				Recover as	
data-1 VM OS Linux	Storage Domain domain_ec2_id_ic	Job Name Backup-job		**data-1** Snapshot: Nov 1, 2018 8:28am, 18 GiB (Latest Snapshot)	✕
photon8 VM OS Linux	Storage Domain domain_ec2_id_ic	Job Name AF-backup		**photon8** Snapshot: Nov 1, 2018 5:29am, 216 GiB (Latest Snapshot)	✕

37 Recover Points for **data-1**

Recover Point ▼	Snapshot Type	Backup Type	Stored
○ Nov 1, 2018 8:28am	Crash-Consistent	Incremental Backup	
○ Nov 1, 2018 2:28am	Crash-Consistent	Incremental Backup	
◉ Oct 31, 2018 8:28pm	Crash-Consistent	Incremental Backup	
○ Oct 31, 2018 2:28pm	Crash-Consistent	Incremental Backup	
○ Oct 31, 2018 11:29am	Crash-Consistent	Incremental Backup	
○ Oct 31, 2018 8:51am	Crash-Consistent	Incremental Backup	
○ Oct 31, 2018 2:51am	Crash-Consistent	Incremental Backup	
○ Oct 30, 2018 8:51pm	Crash-Consistent	Incremental Backup	
○ Oct 30, 2018 2:51pm	Crash-Consistent	Incremental Backup	
○ Oct 30, 2018 8:51am	Crash-Consistent	Incremental Backup	

Save   Cancel   1 2 3 4 ≡

9. Optionally, toggle the radio button to choose to rename the recovered VMs.

10. Choose the option to recover the VMs to their original location, or to a new one. Recovery to a new location can only be done to a source already known by the Cohesity cluster, for example, a different vCenter server that is already configured as a source.

11. Choose the option to keep the networking configuration as it was originally configured, to leave the configuration but leave disconnected, or to leave the network detached.

12. Choose whether to power the recovered VM on, or to leave it powered off.

13. Click Finish.

The recovery job will immediately start the recovery process and restore the VMs using the settings specified in the job.

Views

A Cohesity View provides network accessible storage distributed across the Cohesity cluster, as either NFS volumes, SMB/CIFS mount paths, or S3 compliant object-based storage. A view targets a specific Cohesity storage domain, taking advantage of the settings in that domain regarding compression, deduplication, encryption, and the efficiency derived from the choice between data replication or erasure coding. In order to mount a view, the client computer must reside in a whitelisted subnet. The views created support the following protocol specific settings and capabilities:

- NFS version 3.0 is supported, however NLM locking is not supported.

- NFS mounts and filenames only support ASCII and UTF-8 filenames.

- SMB versions 3.0 and 2.x are supported.

- DFS links are not supported.

- Only NTLMv2 authentication is supported. (Windows 2008 R2 and earlier by default only use LM and NTLM, and therefore must be modified)

- SMB shares are not automatically discoverable, however they will be found when browsing directly to the Cohesity cluster, for example: \\<Cohesity_cluster_name> or \\<Cohesity_cluster_VIP>

- In order to define SMB share file and folder level ownership and permissions, the Cohesity cluster must be added to a Microsoft Active Directory domain. Without Active Directory integration, all shares give full control to everyone.

- Creating a view that contains all 3 protocols results in an S3 view that is read only. In order to create an S3 compliant view that has read/write capabilities, create a view that uses only the S3 protocol.

To create a view, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Platform menu at the top of the screen, click Views.

3. Click Create View.

4. Enter a name for the new View, and optionally enter a description.

5. Select the desired Storage Domain from the dropdown list.

6. Select the QoS Policy for this view from the dropdown list. Options include:

   a. Backup Target SSD: The Cluster sends sequential and random I/Os to SSD and the I/Os are not treated as high priority.

   b. TestAndDev High: The Cluster sends sequential and random writes to a distributed journal, which writes to two separate SSDs on different Nodes and acknowledges the I/O. The I/Os with this QoS policy are given higher priority compared to I/Os with other QoS policies except TestAndDev Low.

   c. TestAndDev Low: The same as TestAndDev High, except that the I/Os with this QoS policy are given higher priority compared to I/Os with other QoS policies.

   d. Backup Target High: The Cluster generally sends sequential data to HDD and random writes to SSD.

   e. Backup Target Low: The same as Backup Target High except that the priority for processing workloads with this policy is lower than workloads with Backup Target High.

7. Click the "Show Advanced Settings" link to expand the list of options available.

8. Select the protocol(s) to be used for this View by clicking the appropriate radio button.

9. Modify the additional options for the view being created as necessary. Some options are not available by default, for example setting specific SMB share ownership and default permissions is not available unless the Cohesity cluster has been joined to an Active Directory domain.

10. Click the Create View button.

Shares

When a new View is created, the name of the View is the default mount path for the clients, and this mount path will establish the top of the file tree. Additional Shares can be created which directly target subfolders within the file tree for ease of navigation. For example, a View can be created for a company division, and then subfolders for each department can be created, each with their own share. End users could navigate to the top of the tree by using the division share, or navigate and map directly to their department share.

To create an additional share within a View, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Platform menu at the top of the screen, click Views.

3. Click the ellipses next to the View you wish to modify, then click View Details.

4. Underneath the Shares & Mount Paths section, click Create Share.

5. Navigate the file tree to the folder where you wish for the new Share to be.

6. Enter a name for the new Share then click Create Share.
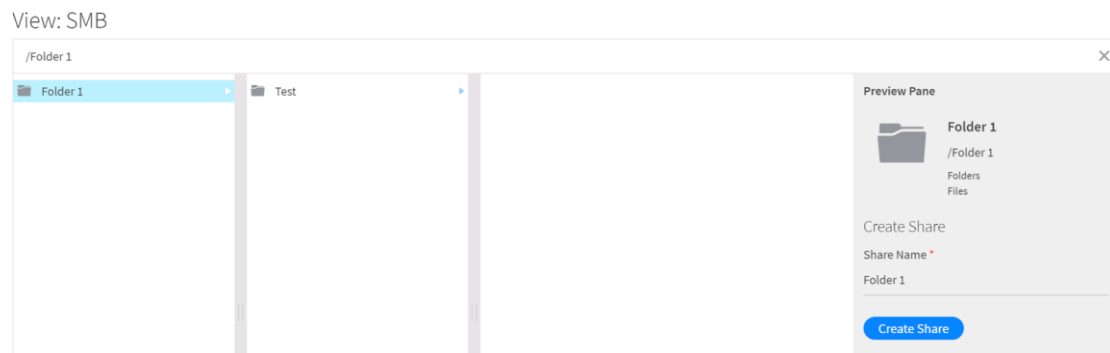
Global Whitelist

A Global Whitelist is configured to allow all clients which match the IP subnet to access all Views created in the Cohesity cluster. In addition to the Global Whitelist, the whitelist settings can be modified individually in the Advanced Settings of each View.

To modify the Global Whitelist, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Platform menu at the top of the screen, click Views.

3. Click Global Whitelist.

4. Enter an IP subnet and a subnet mask to add to the Global Whitelist and optionally enter a description.

5. Click Add.

View Protection

Views can also be protected by View Jobs similar to the way virtual machines are protected via Protection Jobs. View protection targets the same Storage Domain configured for the Cohesity View, and will operate according to the same configured Policies as do Protection Jobs.

To configure protection of a Cohesity View, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Platform menu at the top of the screen, click Views.

3. Click the ellipses next to the View you wish to back up then click Protect View.

4. Enter a name for the View Job, and optionally enter a description.

5. Choose a Policy from the drop-down list.

6. Click Protect.

View Recovery

Unlike a Recovery Job for a VM, which regenerates the VM from the backed-up snapshots, recovery of files from within a View involves creating a Recovery Job, which actually allows you to browse the file tree to locate the file(s) you need to recover. The files can then be downloaded and manually put back into their original locations by the administrative staff.

To recover files from a View snapshot, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Protection menu at the top of the screen, click Recovery.

3. Click the Recover button, from the drop-down list that appears, click Files or Folders.

4. Click the radio button for Browse or Specify Path.

5. Enter the name of the View for which you need to recover files in the search field.

6. When the protected View appears, click the name of the View.

7. From the pop-up window that appears, navigate the file tree until you locate the file(s) you need to recover. Click the file, then click Download File.

8. After all the necessary files have been downloaded locally to your computer, click Close.

9. Manually copy the downloaded file(s) to their original location the Cohesity View's file tree.
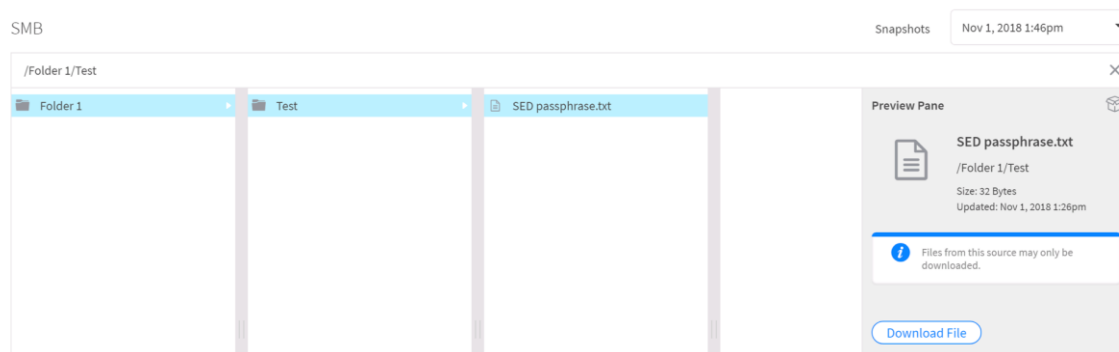
Test/Dev

An additional feature of the Cohesity software is the ability to rapidly clone virtual machines for testing or development purposes. A cloned VM is functionally a restored copy of a snapshot point in time of that VM, and not a clone of the currently running VM. A virtual machine which is cloned using the Cohesity dashboard runs with its virtual machine files and virtual disk files stored in an NFS datastore, which is temporarily mounted by two of the VMware hosts from the Cohesity cluster nodes. Because the VM runs directly from the Cohesity cluster, this use case is appropriate for functional testing, end-user acceptance, and software development or debugging activities, where performance is a secondary consideration. The Clone VMs task gives the administrator the opportunity to rename the cloned VM, clone it to another registered source, and change the network in which the VM is attached. In most cases, manual reconfiguration of the VMs network configuration would be necessary after the clone is created, assuming the VMs use static IP addressing, in order to avoid address conflicts.

To clone a VM for Test/Dev use, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the More menu at the top of the screen, click Test & Dev.

3. Click the Clone button, from the drop-down list that appears, click VMs.

4. In the search field, search for the name of the VM or VMs that need to be recovered. Wildcard characters can be used, and additional filters can be applied.

5. Check the checkbox next to the name(s) of the VMs you wish to recover.

6. Steps 4 and 5 can be repeated multiple times to select different VMs, for example searching for the name of one VM, selecting it, then clearing the search field, searching for the name of another VM, and then selecting that one as well.

7. When all the desired VMs have been selected, click Continue.

8. To rename the cloned VM, toggle the switch for Rename Cloned VMs, then enter a prefix or suffix to add to the name of the VM.

9. Select the Source from the drop-down list in order to pick the location to clone the VM. The Resource Pool drop-down list will list the hosts or clusters available. The VM Folder will list the available folders to place the VM into. Finally, the View field represents the name of the NFS datastore from the Cohesity cluster which will be mounted to the VMware hosts.

10. Choose to either leave the networking detached in order to perform manual reconfiguration, or click the radio button for Attach to a new network, and choose the VM port group to attach the cloned VM to.

11. Choose to leave the VM powered off, or to power it on after the task completes.

12. Click Finish.

To tear down and delete cloned VMs when they are no longer needed, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the More menu at the top of the screen, click Test & Dev.

3. From the list of clone jobs presented, click the name of the clone job that created the clone that you wish to remove.

4. Click the Tear Down Clone button.

5. Click the Yes, tear down button in the pop-up window that appears.

Monitoring

The Cohesity software offers numerous options for passive and proactive monitoring of the cluster, including job status, performance, hardware status, storage capacity and more.

Dashboard

The Dashboard screen in the Cohesity HTML management webpage provides a useful overview of the status of the overall system health, backup job runs, storage efficiency and performance over the past 24 hours. The dashboard allows the Cohesity administrator to see at a quick glance if any items need attention.

Figure 21   Cohesity Dashboard
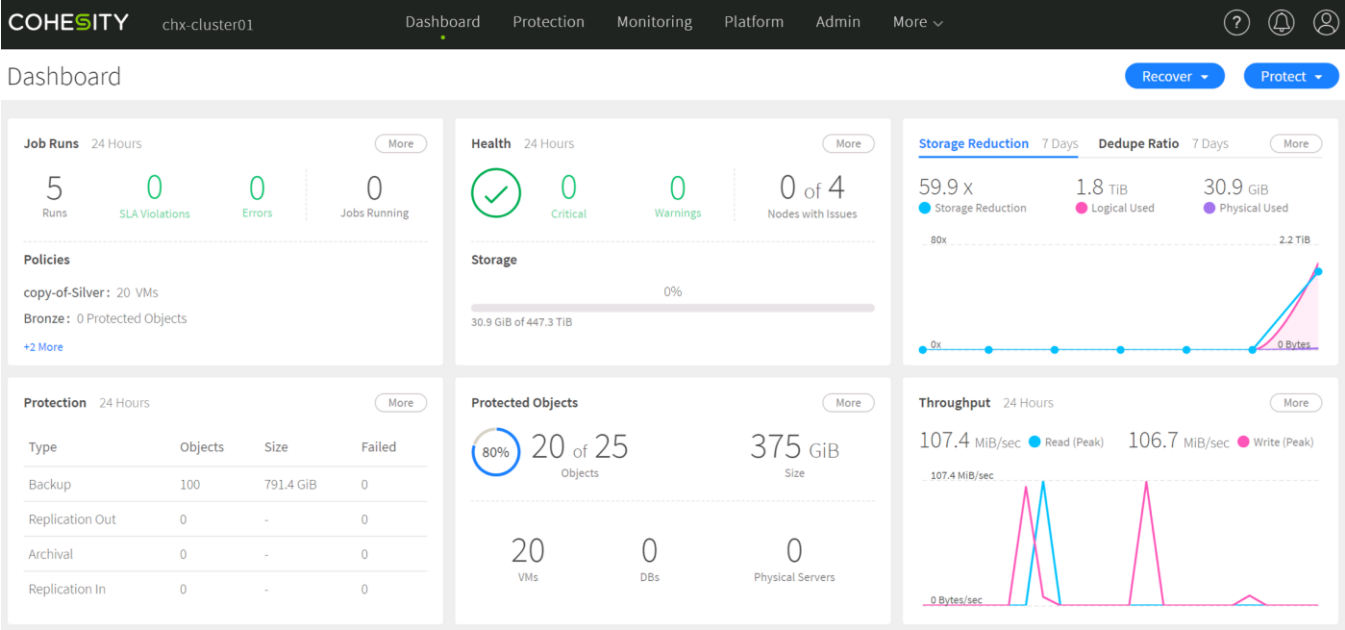
Performance

Under the Monitoring menu, the Performance screen can be used to view the storage I/O per second (IOPS), latency and throughput, plus the CPU and memory usage of the nodes in the Cohesity cluster. Views can be modified to see figures for the entire cluster, individual nodes, or individual Storage Domains. The view timeframe can be modified and the view can be zoomed in for greater detail.

Figure 22   Cohesity Performance

89
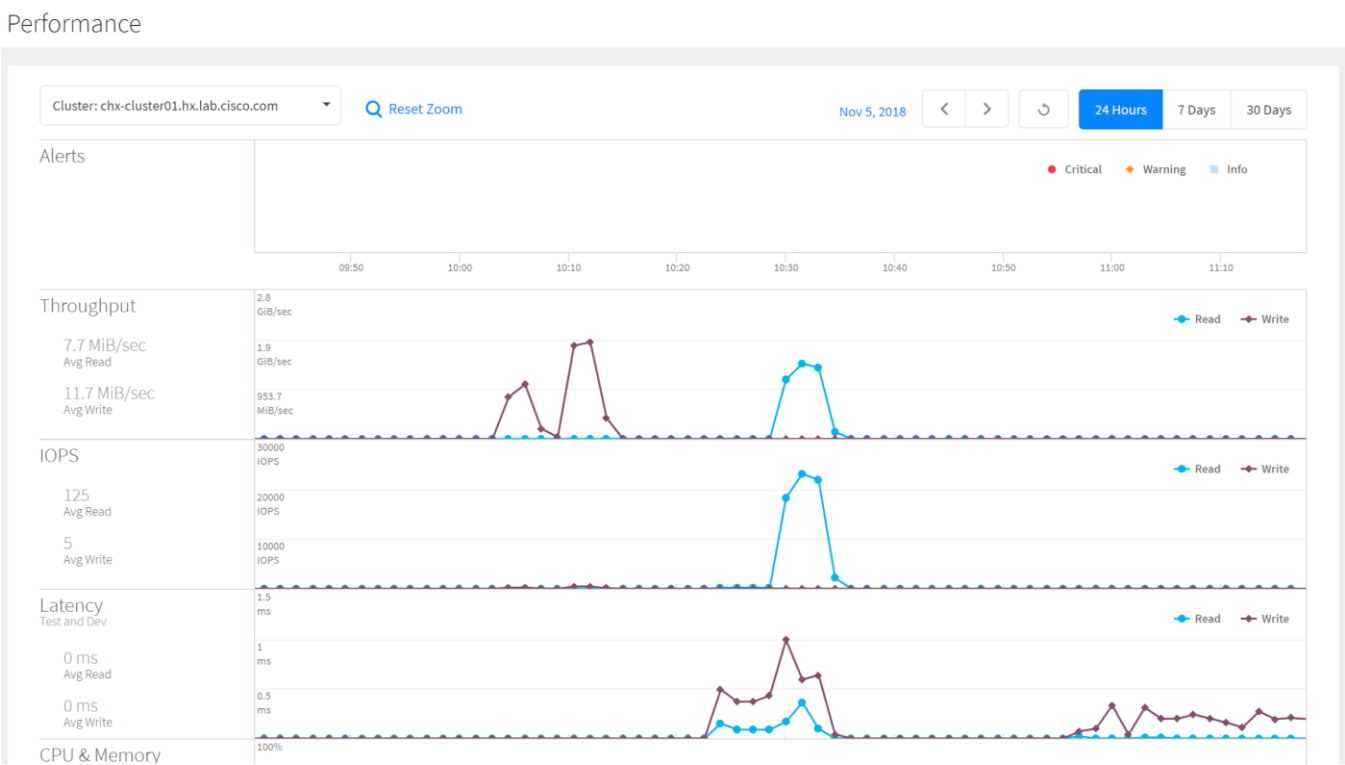
Alerts

Alerts in the Cohesity cluster can be viewed under the Monitoring menu by clicking Alerts. Alerts can be configured to automatically send a notice to an email recipient as well.

To configure alerts to be sent to an email recipient, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Monitoring menu at the top of the screen, click Alerts.

3. Click the Edit Alert Settings button.

4. Enter the email address of the recipient then click Add to List.

Figure 23   Cohesity Alerts
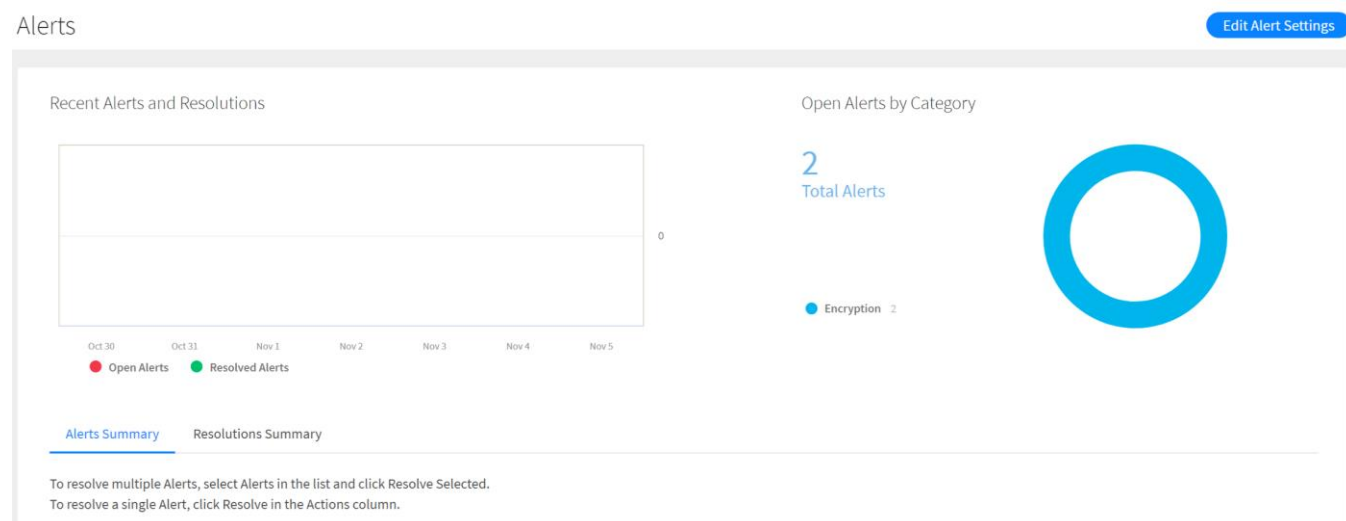
Storage

Under the Monitoring menu, clicking Storage will show a page detailing the storage space consumption in the Cohesity cluster, plus the data reduction efficiency figures due to deduplication and compression. The views in the charts can target the entire cluster or individual Storage Domains, and the timeframe for the charts can be customized.

Figure 24   Cohesity Storage Monitor

Reports

From the Monitoring menu, click Reports to view dynamic reports which can be generated for viewing, exporting, and also as regular emails. There are several report types available, such as backup snapshot summaries, backup job status, cluster health and storage, and many more. Reports can be tailored to show specific date ranges, sources, and statuses, and then configured to be regularly sent via email by clicking the email clock link. Note that not all reports can be configured to send automatically via email.

Figure 25   Cohesity Reports
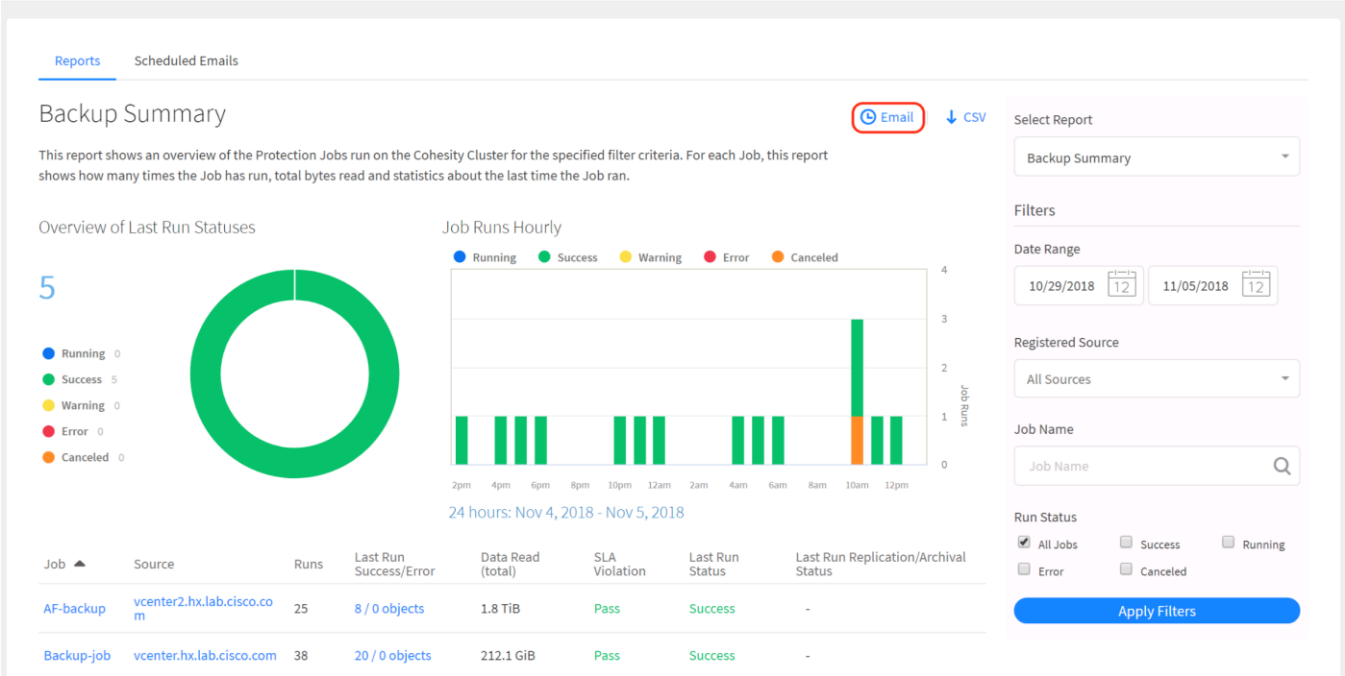
SMTP

In order to send alerts and/or scheduled reports to recipients via email, the Cohesity cluster's SMTP settings must be enabled.

To enable SMTP email alerts and reports, complete the following steps:

1. Log into the Cohesity Dashboard web page.

2. From the Admin menu at the top of the screen, click Cluster Settings.

3. Enter a system administrator email address which represents this Cohesity cluster, this will be the "From:" address for outgoing emails.

4. Toggle the switch to Enable SMTP Server.

5. Enter the SMTP server information as appropriate for your environment.

6. Toggle the switch to Test Email on Save.

7. Click Save.

SNMP

If necessary, SNMP traps can be sent to an SNMP receiver for parsing by a network management system. In the Monitoring > SNMP menu, click Edit to enable and configure SNMP traps, their destination, and the trap user.

Remote Support

Cohesity offers a remote support service named Support Channel, which is enabled by default. Support Channel initiates outgoing sessions to register the Cohesity system with Cohesity's Support Channel server and technical support team. Cohesity support staff can securely connect and log into the cluster remotely for on-demand technical support and troubleshooting using SSH and secure keys. In some cases, Support Channel connections may require the use of a proxy server to function properly.

Validation

Test Plan

Numerous scenarios were developed to rigorously test the integration between the Cohesity systems and Cisco HyperFlex, and also to test the redundancy and durability of the Cohesity system running within the Cisco UCS domain. All of the tests below were executed in Cisco's labs, using complaint hardware and software as listed previously, and configured according to the instructions in this document. Tests executed included, but were not limited to the following:

Installation

- Creation and application of Cohesity specific UCS policies and service profiles.

- Successful installation and initial configuration of Cohesity 6.1.0a software on Cisco UCS managed C240 M5L server hardware.

Core Functional Testing

- Enable Cohesity Storage Snapshot Provider features.

- Configure multiple Cisco HyperFlex clusters as Storage Snapshot Provider sources.

- Configure multiple VMware vCenter systems as sources which manage both single and multiple Cisco HyperFlex clusters.

- Configure and run Protection Jobs backing up VMs from multiple Cisco HyperFlex clusters.

- Configure and run a large-scale Protection Job, protecting ~500 VMs which generate random data.

- Configure and run Restore Jobs restoring VMs across multiple Cisco HyperFlex clusters.

- Configure Cohesity Views providing NFS, SMB and S3 compliant file services.

- Configure Test/Dev clones of VMs for temporary use.

Extended Functional Testing

- Backup and restore VMs to/from multiple Cisco HyperFlex deployment types, including traditional Hybrid clusters, All-Flash clusters, clusters with self-encrypting drives having encryption enabled, extended clusters with compute only nodes, and stretched multi-site clusters.

- Backup and restoration of VMs running on Cisco HyperFlex version 3.5 and 3.0 clusters.

- Configure Remote Cluster pairing and replication of snapshots between multiple Cohesity systems.

- Configure a NetApp FAS array as an External Target and archive backups to the array.

- Backup VMs using a mixture of jobs with the Storage Snapshot integration both enabled, and disabled.

- Backup of VMs without existing snapshots and also with existing VMware standard snapshots or existing HyperFlex native snapshots.

- Restore VMs that had no previously existing snapshots, and also with previously existing VMware standard snapshots and HyperFlex native snapshots.

- Restore VMs to Cisco HyperFlex clusters that were not the original location of the VMs.

- Configure Active Directory integration, use AD credentials for logins, and also for setting View ownership and permissions.

- Configure and test SMTP alerts and reporting.

Failover and Redundancy Testing

All failover and redundancy tests were conducted while at least one active Cohesity Protection Job was running.

- Fail the active network path for one Cohesity node.

- Fail the active network path for two Cohesity nodes.

- Fail all the network uplinks from a single Fabric Interconnect.

- Fail the active side Fabric Interconnect.

- Ungraceful shut down the Storage Controller VM (SCVM) of one HyperFlex node.

- Ungraceful shut down of one HyperFlex node, causing VMs to restart via VMware High Availability.

Bill of Materials

Below is an example Bill of Materials used to order four (4) of the Cisco UCS C240 M5L servers, which are compliant with the requirements to run the Cohesity DataPlatform software, along with the pair of Cisco Fabric Interconnects, and the 10 GbE cables to connect them, as used in the testing and reference design outlined in this document.

Table 29 Cohesity on Cisco UCS Sample Bill of Materials

Line Number	Item Part Number	Item Description	Quantity
1.0	UCSC-C240-M5L	UCS C240 M5 12 LFF + 2 rear drives w/o CPU,mem,HD,PCIe,PS	4
1.1	UCS-CPU-6142	2.6 GHz 6142/150W 16C/22MB Cache/DDR4 2666MHz	8
1.2	UCS-MR-X32G2RS-H	32GB DDR4-2666-MHz RDIMM/PC4-21300/dual rank/x4/1.2v	16
1.3	UCSC-PCI-1B-240M5	Riser 1B incl 3 PCIe slots (3 x8) all slots from CPU1	4
1.4	UCSC-PCI-2C-240M5	Riser 2C incl 3 PCIe slots (3 x8) supports front+rear NVMe	4
1.5	UCS-M2-240GB	240GB SATA M.2	8
1.6	UCSC-PSU1-1050W	Cisco UCS 1050W AC Power Supply for Rack Server	8
1.7	UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	4
1.8	CIMC-LATEST	IMC SW (Recommended) latest release for C-Series Servers.	4
1.9	UCSC-HS-C240M5	Heat sink for UCS C240 M5 rack servers 150W CPUs & below	8
1.10	UCSC-RNVME-240M5	C240 M5 Rear NVMe CBL(1) kit, Rear NVMe CBL, backplane	4

Line Number	Item Part Number	Item Description	Quantity
1.11	UCS-MSTOR-M2	Mini Storage carrier for M.2 SATA/NVME (holds up to 2)	4
1.12	UCSC-SAS-M5	Cisco 12G Modular SAS HBA (max 16 drives)	4
1.13	UCS-HD10T7KL4KN	10 TB 12G SAS 7.2K RPM LFF HDD (4K)	48
1.14	UCSC-NVMEHW-H3200	3.2TB 2.5in U.2 HGST SN200 NVMe High Perf. High Endurance	8
1.15	UCSC-MLOM-C25Q-04	Cisco UCS VIC 1457 Quad Port 10/25G SFP28 CNA MLOM	4
1.16	CAB-C13-CBN	Cabinet Jumper Power Cord, 250 VAC 10A, C14-C13 Connectors	8
2.0	UCS-FI-6454-U	UCS Fabric Interconnect 6454	2
2.1	N10-MGT016	UCS Manager v4.0	2
2.2	UCS-FAN-6332	UCS 6332/ 6454 Fan Module	8
2.3	UCS-ACC-6332	UCS 6332/ 6454 Chassis Accessory Kit	2
2.4	UCS-PSU-6332-AC	UCS 6332 Power Supply/100-240VAC	4
2.5	CAB-C13-CBN	Cabinet Jumper Power Cord, 250 VAC 10A, C14-C13 Connectors	4
3.0	SFP-10G-AOC3M=	10GBASE Active Optical SFP+ Cable, 3M	8

Summary

Primary data can be loosely defined as the mission-critical information required for companies to do business. Secondary data includes backup and recovery, files and objects, test and dev, archive and analytics. Both types of data are growing exponentially, are becoming harder to store, protect and access. Much of it, particularly secondary data, is buried or stored in silos. And, both types of data are becoming more critical for meeting business objectives. To meet compliance and regulatory requirements and to more cost-effectively store, protect and provide business applications with the right data to run them, companies need to modernize their data center operations. Hyperconvergence has been key in replacing slow, siloed data center systems with scale-out software-defined infrastructure on standard hardware with a pay-as-you-grow model.

Cisco HyperFlex and the Cohesity DataPlatform on Cisco UCS are modern, hyperconverged platforms for managing both primary and secondary workloads. Now, these two industry-leading platforms are integrated to deliver an agile web-scale solution, for on premise and cloud, that easily scales with customer's business needs, consolidates IT operations for all workloads, brings frictionless mobility to data and applications from edge to core to cloud, and put data to work. With the Cisco-Cohesity integrated solution, hyperconvergence meets hyperconvergence. This means customers can easily unify, protect, access and control their data across clouds, data centers or remote and branch offices at a significant cost saving compared to traditional alternatives.

For buyers of servers and storage solutions, who need radically simplified management, predictable pay-as-you-grow pricing, and greater productivity from data, Cisco HyperFlex and Cohesity on Cisco UCS provide unlimited scale, flexibility, and unified management for both primary and secondary data to reduce costs and improve business agility. Unlike antiquated point products that are slow to deploy, costly to purchase, and require expensive over provisioning or disruptive downtime, only Cisco and Cohesity provide a complete, integrated, and flexible hyperconverged infrastructure to help customers modernize and better manage their entire data center. Now, both primary and secondary data and applications can operate at web-scale and empower IT and business with agility to seamlessly support data growth, mobility, and insights – from edge to core to cloud.

For more information on Cisco-Cohesity integrated solution, please see:
https://www.cohesity.com/products/cisco.

About the Authors

Brian Everitt, Technical Marketing Engineer, Cisco UCS Data Center Engineering Group, Cisco Systems, Inc.

Brian is an IT industry veteran with over 20 years of experience deploying server, network, and storage infrastructures for companies around the world. During his tenure at Cisco, he has been a lead Advanced Services Solutions Architect for Microsoft solutions, virtualization, and SAP Hana on Cisco UCS. Currently his focus is on Cisco's portfolio of Software Defined Storage (SDS) and Hyperconverged Infrastructure solutions. Brian has earned multiple certifications from Microsoft, Cisco, and VMware.

Acknowledgements

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

- Damien Philip, Principal Solutions Architect, Cohesity

- Sanjeev Desai, Senior Director, Solutions Marketing, Cohesity