



Cisco HX Data Platform Security Hardening Guide

Version 5.5.1 rev3

October 2023
OBSOLETE IF PRINTED



Document Information

Document Summary	Prepared for	Prepared by	
Cisco HX	v.5.5.1 rev3	Field	Aaron Kapacinskas
Last Modified	10 October 2023		
Previous Version	5.0.2 rev8		
Changes in this version:			
Expanded Intersight Management Security Section			
Added IS RBAC discussion			
Added API behavior for non-removal of PUT and DELETE methods			
Updated external CA certificate procedure in appendix F			
Updated STIG information to cover STIG for ESXi 7.0			
Added information on persistent root to the Admin Shell section			
Added Appendix J on removal of air-gapped persistent root			

Intended Use and Audience

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

Bias Statement

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

Legal Notices

All information in this document is provided in confidence and shall not be published or disclosed, wholly or in part to any other party without Cisco's written permission.



Contents

- Document Information 2**
 - Intended Use and Audience 2
 - Bias Statement 2
 - Legal Notices 2
- Prerequisites 9**
- Introduction 9**
- Secure Product and Development Components 9**
 - Development Milestones..... 9
 - CSDL Philosophy..... 9
 - CSDL Product Adherence Methodologies 11
 - Vulnerability Handling..... 12
 - Tenable IO Scanning..... 12
 - CERT Advisory 13
 - VMware ESX Patching..... 13
 - HXDP Patching 14
 - Additional Vulnerability Testing Measures 14
 - Secure Platform “Modules” 14
 - Control Plane 14
 - Data Security..... 14
 - Management Security 14
 - Certification Process 15
 - ACVP 15
 - Current Certifications 15
 - FIPS..... 15
 - Common Criteria 17
 - Other Certifications and Procedural Guidelines 17
 - CSfC 17
 - ISO 27001 17
 - ISO 27001:2013 18
 - FISMA 18



FedRAMP	18
SOC 2 Type 2	18
TEMPEST (.....	19
IAVA	19
HIPAA	19
UK TSR Compliance	19
NERC CIP	22
CNSA	24
DISA APL.....	24
Targeted Certifications	24
Counterfeit Prevention	24
HX Components and Environment.....	25
Solution Components	25
Cisco UCS	26
Cisco UCS Fabric Interconnects (FIs).....	27
HX Nodes	27
Management Interfaces: HX Connect and the VMware vCenter Plug-in	29
VMware vCenter	29
VMware ESX	29
VMs	29
Client Machines	29
HX Secure Network Environment and Component Requirements.....	29
Port Requirements for Communication	30
Scans Showing Undocumented Ports.....	30
Port Requirements and Logical Traffic Flow for Replication	31
Intersight and PVA/CVA Connectivity Requirements and Resilience	32
Unicast and Multicast Requirements	34
Datastore Access	34
Auto Support and Smart Call Home (SCH)	35
Installation and ESX Best Practices and Security Considerations	36
Cisco HX Installer (HX Installer)	36
Minimum Infrastructure and Port Requirements for Local Installation.....	37
Default Passwords.....	39



VLANs and vSwitches	39
FI Traffic and Architecture	41
UCSM Requirements	42
VNICs	42
East-West Traffic	42
North-South Traffic	42
Upstream Switch	42
VLANs	42
Disjoint L2 Networks	43
Cisco HyperFlex Edge (HX Edge)	43
HX Data Security	43
Encryption Services	44
HX Native Software Encryption	44
Key Management with HX Software Encryption	46
SEDs	46
Key Management with SEDs	48
Secure Erase	48
Certificate Signing Requests (CSRs)	49
Networking Considerations	50
Encryption Partners	50
VM Level Encryption with 3 rd Party Vendors	51
Secure Communications	51
Usage of NFS in HXDP	52
HXDP Runtime Defenses	53
HX Management	54
Management Interfaces	54
Intersight	54
Intersight RBAC specific to HX	58
HX Connect	59
vCenter Plug-in	60
STCLI and HXCLI	61
Secure Admin Shell Access (HXDP 4.5.1 (a) and above)	63
REST APIs	65



AAA Domains.....	67
vCenter	67
AD Integration.....	67
User Management	67
Cisco HyperFlex User Overview	68
Local Users	68
UI Users	68
CLI Users	70
Auditing, Logging, Support Bundles	70
Setting Up Remote Logging for HXDP 4.0.1.a and Later.....	73
Password Requirements.....	75
Password Guidelines	76
Session Timeouts.....	77
HX Platform Hardening	81
HyperFlex Security Advisor (SA)	81
US Federal STIG (Secure Technical Implementation Guide) Settings	82
SSL Certificate Replacement.....	84
Secure Boot.....	84
SSL Certificate Thumbprint (Hash) and Signatures	88
Dynamic Self-Signed Certificates in HX	88
UCSM Certificate Management	89
HX and Perfect Forward Secrecy (PFS)	89
TLS Weak Protocol Disable	91
TLS Weak Cipher Disable	92
SSH (ESX) Lockdown Mode and Root Logins	92
SSH, Lockdown, and HX Upgrades.....	93
Tech Support Mode	93
Third Party Software Execution on Fls and HXDP	93
Whitelisting and other STCLI Security Commands.....	94
HX Data Platform Firewalling: IP Tables.....	94
Replication.....	96
Specific ESX Environment Hardening Settings Relevant to HXDP	96
Specific UCS Environment Hardening Settings Relevant to HXDP	96



Control VM (SCVM) Customization	97
References	97
Intersight Security Guides and Certifications	97
ESX Hardening Guide.....	97
UCS Hardening Guide	98
Cisco CSDL.....	98
Syslog-ng Configuration	98
Secure Boot Assurance Via Attestation.....	98
Appendix A: Networking Ports	99
Intersight	105
Appendix B: URLs Needed for Smart Call Home, Post Install Scripts, Intersight.....	106
Appendix C: ESX Hardening Settings	108
Appendix D: Acronym Glossary	120
Appendix E: Sample Syslog-ng Configuration File.....	122
Appendix F: Certificate Management and Use Cases	124
SCVM: How to generate and replace External CA Certificate in HX 4.0.2a+	124
HX Certificate Management	126
vCenter: How to generate and replace External CA Certificate.....	126
ESX: How to generate and replace External CA Certificate	126
ESX: To re-generate self-signed certificate	127
HX Cluster: Re-registration	127
Hyperflex Use Cases.....	127
Observations and Notes.....	128
vCenter: self-signed cert with certMgmt mode = vmsa (Default Mode)	128
vCenter: CA signed cert with certMgmt mode = vmsa (Default Mode)	129
vCenter: CA signed cert with certMgmt mode = custom	129
Hyperflex Use Cases: HX 4.5.1a+ and NGINX Self-Signed Without CA Signed	129
What Happens When Your Self-Signed Certificate Expires	131
Appendix G: SCH Configuration and Proxy.....	137
Appendix H: Changing the SSO Timeout for vCenter.....	140
Appendix I: Password Recovery	141
Password recovery for 4.5(1a), 4.5(2a)	141
Password recovery for 4.5(2b) and 5.0(1x)	143



Admin password recovery for 4.5(2c) and 5.0(2a) onward	143
Appendix J: Persistent Root Air-Gapped Revocation	145



Prerequisites

We recommend reviewing the release notes, installation guide, and user guide before proceeding with any configuration. The Cisco HyperFlex Data Platform (HXDP) should be installed and functioning per the installation guide. Please contact your Cisco Representative or Cisco Support if assistance is needed. **Please note that this version of the Hardening Guide has pruned content prior to 4.0.2a for manageability. Future versions will truncate at 4.5.2a and 5.0.2a as the release model advances. If you need older versions, please contact your Cisco representative.**

Introduction

The Cisco HyperFlex Data Platform Security Hardening Guide provides guidance for HyperFlex (HX) users in ensuring that their product is deployed in a more robust and secure manner. It is necessary to understand the architecture and components of the solution in order to complete this properly. This document provides recommended configuration settings and deployment architectures for HXDP-based solutions. It is intended to be used in conjunction with product documentation for deployments where extra consideration for platform security is required. For product documentation, please contact your Cisco representative.

Secure Product and Development Components

Cisco HyperFlex product components are developed, integrated, and tested using the Cisco Secure Development Lifecycle (CSDL). Secure product development and deployment has several components ranging from inherent design and development practices, testing the implementation, and finally a set of recommendations for deployments that maximize the security of the system.

Development Milestones

Each iteration of the product's development addresses needs for ongoing security fixes and general feature enhancements that include security components (new deployment models, changes in management, partner on-boarding, etc.). At every stage of development, the Hardening Guide undergoes potential enhancements relative to findings and new features.

- The HX Hardening Guide has the following components:
 1. VMware ESXi settings
 2. Cisco UCS settings
 3. HX Hardening
- The system is configured in QA to accommodate the relevant settings identified above and run through a typical deployment test.
- The result is a validated set of best practices for security and is communicated through the CSDL process and exposed in the Hardening Guide.

CSDL Philosophy

A poor product design can open the way to vulnerabilities. The CSDL is designed to mitigate these potential issues.

At Cisco, our "secure design" approach requires two types of considerations:



- Design with security in mind
- Use threat modeling to validate the design's security



Designing with security in mind is an ongoing commitment to personal and professional improvement through:

- Training
- Applying the Product Security Baseline (PSB) design principles
- Consider other industry-standard secure design principles
- Be aware of common attack methods and design safeguards against them
- Take full advantage of designs and libraries that are known to be highly secure
- Consider all entry points

We also reduce design-based vulnerabilities by considering known threats and attacks. With threat modeling, we:

- Follow the flow of data through the system.
- Identify trust boundaries where data may be compromised.
- Based on the data flow diagram, generate a list of threats and mitigations from a database of known threats, tailored by product type.
- Prioritize and implement mitigations to the identified threats.

The goal of this effort is to ensure a security mind set at every stage of development:

- Secure Design
- Secure Coding
- Secure Analysis
- Vulnerability Testing
- Secure deployments

HX product development focuses on two areas to satisfy the CSDL model:

- Internal Requirements
 - Adhere to the secure development process
- Market based requirements
 - Complete and validate against certifications (Federal)
 - Document and educate (HX Hardening Guide)

CSDL Product Adherence Methodologies

Cisco CSDL adheres to Cisco Product Development Methodology (PDM), ISO27034 and ISO9000 compliance requirements. ISO 27034 standard provides an internationally recognized standard for application security. Details for ISO 27034 can be found [here](#). The ISO 9000 family of quality management systems standards is designed to help organizations ensure that they meet the needs of customers and other stakeholders while meeting statutory and regulatory requirements related to a product or service. ISO 9000 details are [here](#).

The CSDL process is not a one-time approach to product development. It is recursive, with vulnerability testing, penetration testing, and threat modelling plugging into subsequent development that feeds back into the process. This process follows ISO9000 and ISO27034 standards as part of an internationally recognized set of guidelines. The approaches involved often take a solution-wide methodology. For example, the use of our continually updated CiscoSSL crypto module to guarantee that HX (along with other elements in the Cisco offering) are always secure and meet FIPS certification requirements.



Vulnerability Handling

Tenable IO Scanning

Common Vulnerabilities and Exposures (CVE) scanning is a critical part of most deployments. Many industries and Federal organizations standardize on Tenable IO (formerly Nessus Scanner) to implement various DISA or CIS audits.

- CIS is [Center for Internet Security](#)
- DISA is [Defense Information Systems Agency](#)

In our CSDL efforts, we use Tenable IO, produced by Tenable, in our development process

- Tenable IO Scanner – <https://www.tenable.com/products/tenable-io>

The vulnerability scanning workflow is as follows:

- Choose a build to test against (based on dot release development timing)
- Update the scanner signatures and plug-ins for our test date
- Freeze the scanner – line in the sand
- Test, fix as needed, check-in the safe build
- Fixes are immediately scheduled for Critical and High
- CSDL may identify others in Medium and Low and Info that need remediation.

A typical Nessus scan configuration summary might look like this:

- HX 4.0(2a)
- Compliance checks:
 - DISA RHEL 5
 - CIS L2 Ubuntu 16.04 LTS
 - CIS Apache 2.2
- Plug-ins:
 - All plug-ins enabled, same day update
- Sample Report:
 - Output is color coded.
 - 5 Alert Levels: Critical, High, Medium, Low, Info.
 - Notes: System is clean, one low warning, rest are info only.

Scan Information					
Start time:	Thu Mar 30 14:28:32 2017				
End time:	Thu Mar 30 14:29:25 2017				
Host Information					
IP:	10.10.10.10				
MAC Address:	08:00:27:00:00:00				
OS:	Linux Kernel 3.13.0-110-generic on Ubuntu 16.04				
Results Summary					
Critical	High	Medium	Low	Info	Total
0	0	0	1	66	67
Results Details					
0/icmp					
10114 - ICMP Timestamp Request Remote Date Disclosure					
Synopsis					
It is possible to determine the exact time set on the remote host.					
Description					
The remote host answers to an ICMP timestamp request. This allows an attacker to know the date that is set on the targeted machine, which may assist an unauthenticated, remote attacker in defeating time-based authentication protocols. Timestamps returned from machines running Windows Vista / 7 / 2008 / 2008 R2 are deliberately incorrect, but usually within 1000 seconds of the actual system time.					
Solution					
Filter out the ICMP timestamp requests (13), and the outgoing ICMP timestamp replies (14).					

CERT Advisory

Computer Emergency Response Team (CERT) advisories come up as new vulnerabilities are identified. Cisco's internal CERT team monitors and alerts product groups to potential issues that might affect their respective components. When these items are identified by CERT or are otherwise indicated by vendor partners (VMware, etc.). Patches are either developed or acquired from the respective vendors.

VMware ESX Patching

Patches for VMware are immediately supported if they are within the regularly supported VMware dot release. There are no hard commitments for when support for new VMware dot releases will be available, but there are continuous release onboarding processes that occur within QA for each new VMware release.

We do not support cluster level remediation through VUM or vLCM: neither of aware of the HX cluster running on top.

You can try to remediate one host at a time in those tools, but this has diminishing value. The HX ESXi upgrading process is a simple one click as long as it is launched through HX Connect or Intersight. If you use Intersight, it is even simpler than VUM or vLCM as you have the HX customized image that is selected from a drop down and just proceed. No downloading of files into repos or customizing builds are required in that scenario.



HXDP Patching

Cisco scans development builds weekly for CVEs using Tenable's Nessus scanner. Based on these results, we begin developing the patch for critical CVEs related directly to HX as soon as they are discovered. The fixes are rolled into an immediate release or a regularly scheduled incremental with turnaround within 90 days, depending on severity.

The HyperFlex release model identifies Long Term Support releases and Feature Releases. Long Term Support releases are supported, maintained, and patched for 30 months from initial release. You can identify LTS releases because they have the X.Y(2x) designation. For example, 4.5(2a) is an LTS release.

Additional Vulnerability Testing Measures

Cisco also utilizes an internal tool for threat modeling called ThreatBuilder. This tool is used to explicitly map out application components and services and to identify potential attack surfaces and develop line items for direct evaluation. This information along with industry tools are used for vulnerability and exploit testing by Cisco's ASIG (Advanced Security Initiatives Group). ASIG also uses fuzzing and manual testing as part of their suite of tools.

Secure Platform "Modules"

At a high level, HX system security can be broken down into 3 broad categories. These are the Control Plane, Data Security and Management Security.

Control Plane

The control plane deals with system communication. This is the subsystem that implements FIPS compliant encrypted communication protocol engine for communication that may originate outside of the system, for example, from an administrator. It also deals with inter-component communication between nodes which happens on a trusted, internal, non-routed 10GB network.

Data Security

Securing data in the system is the job of the Secure Encrypted Disk (SED) subsystem. The HX nodes are SED capable, meaning that they can incorporate and function using encrypted disks. Key management for this can be handled locally or via remote KMIP servers in HA configurations.

Management Security

Managing the system through the UI or through the command line requires secure communication mechanisms. This is handled via HTTPS for the vCenter plug-in or for HX Connect (the native HTML 5 UI). SSH for encrypted command line access is also handled. Management security also entails role-based access control as well as auditing and logging of system activities and user input.



Certification Process

Federal compliance and audit-based certifications are a critical component of a standardized and predictable security posture. They are critical in most Federal deployments, especially those dealing with financial and defense arenas. The Cisco Global Certification Team (GCT) works to complete various certifications.

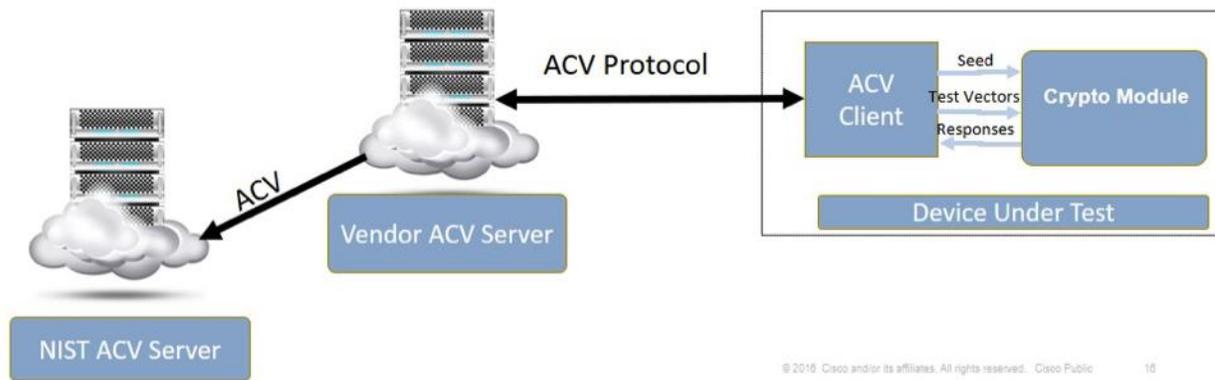
ACVP

The Automated Cryptographic Validation Protocol (ACVP) is part of a NIST program to automate FIPS and Common Criteria testing superseding the process used in the [Cryptographic Algorithm Validation Program \(CAVP\)](#) and the [Cryptographic Module Validation Program \(CMVP\)](#). Details can be found here:

<https://csrc.nist.gov/Projects/Automated-Cryptographic-Validation-Testing>

Beginning in CY2020, Cisco began using ACVP built into the CiscoSSL module for HyperFlex in order to process Federally accepted cryptographic certifications. The process will use a series of ACVP/NIST proxy infrastructure servers to complete the certifications using communication directly to NIST validation servers. The figure below shows the general product architecture used for ACVP.

Proxy/Validation Authority Architecture Automated Cryptographic Validation System



Current Certifications

FIPS -- The **F**ederal **I**nformation **P**rocessing **S**tandard (**FIPS**) Publication 140-2, is a U.S. government computer security standard used to approve cryptographic modules.

HyperFlex is compliant with FIPS140-2 level 1 via direct implementation of the FIPS compliant CiscoSSL crypto module. The module, once implemented, is vetted by a 3rd party that is federally certified to ascertain compliance status.

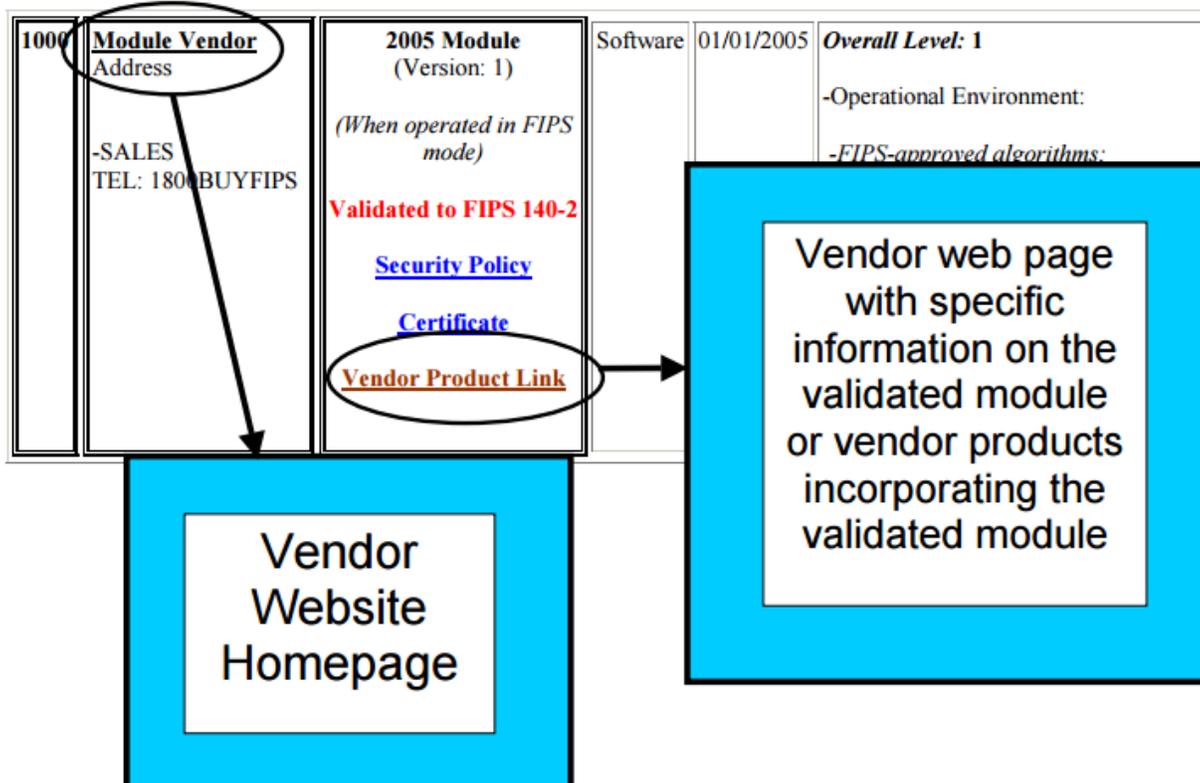
- Utilizes CiscoSSL module
 - Already FIPS compliant
 - SSH approved cipher list
 - SSL/TLS implementation
 - Eliminates weak or compromised components
 - Regularly updated
- Lab validates that the module is incorporated correctly
 - Build logs
 - Source access identifying calls to the module
- All admin access points to the cluster are covered here
 - SSH for CLI
 - HTTPS for UI

A comprehensive list of Cisco FIPS compliant products is listed here along with the corresponding reference with NIST.

- Cisco FIPS Certified Products: <http://www.cisco.com/c/en/us/solutions/industries/government/global-government-certifications/fips-140.html>
- Cryptographic Module Validation Program (CMVP) vendor list:
- <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401vend.htm>

CMVP Vendor Product Link

Help Customers Find Your Validated Module and Products





Common Criteria for Information Technology Security Evaluation (Common Criteria or CC) is an international standard (ISO/IEC 15408) for computer security certification, currently in v3.1 rev 5.

- System users specify their security functional and assurance requirements through the use of protection profiles, vendors can then make claims about the security attributes of their products, and testing laboratories can evaluate the products to determine if they actually meet the claims.
 - Customers have some security needs defined in a set of CC guidelines
 - This is my system; this is what I say it can do to meet those
 - Let us (vendor and lab) agree on a test, here is the procedure
 - Here are my results
 - You (lab) run it on your own and verify
 - Deliver certification

The second (first in CYQ42017) certification was completed in CYQ4 2019 for EAL2 on 3.5.2a and is currently available here: <https://www.commoncriteriaportal.org/files/epfiles/Certification%20Report%20NSCIB-CC-215885-CR.pdf>

The third certification was completed Q1 FY2022 for EAL2 on 4.5.2a and is available here: <https://www.tuv-nederland.nl/common-criteria/certificates.html>.

The 3.5.2, 4.0.2, and 4.5.2 code lines are “long term support” releases in the HX release model and will be available, maintained, and patched for 30 months from initial release. The follow-on release will be 5.5.2a.

Other Certifications and Procedural Guidelines

CSfC Commercial Solutions for Classified (CSfC) is a part of NSA's commercial cybersecurity strategy to deliver secure cybersecurity solutions leveraging commercial technologies and products to deliver cybersecurity solutions quickly. It is founded on the principle that properly configured, layered solutions can provide adequate protection of classified data in a variety of different applications. NSA has developed, approved, and published solution-level specifications called Capability Packages (CPs) and, through the National Information Assurance Partnership (NIAP), works with technical communities from across industry, government, and academia to develop, maintain, and publish product-level security requirements called Protection Profiles (PPs). CPs for Mobile Access, Multi-Site Connectivity, Campus Wireless LAN, and Data at Rest solutions.

CSfC is only achievable once a product has been Common Criteria certified against a US Scheme (NIAP) [NIAP approved Protection Profile](#) (PP/cPP) with an added PP-Module or extended Package (also on the NIAP Approved PP list). The Cisco product also must be listed on NIAPs [Product Compliant List](#).

Hyperflex has been Common Criteria certified using Evaluation Assurance Levels (EALs) but not against a NIAP approved cPP and PP-Module/extended Package. As a result, Hyperflex is not listed on the NIAP Product Component List. If you need CSfC compliance, please contact your Cisco representative.

ISO 27001 is not a certification for specific pieces of hardware as much as a dozen or so “Best Practices” in the form of checklists/guidelines for how organizations manage their security controls internally. It observes things like building access, password management, badging into a copier to make copies, etc. Training on a frequent basis is a part of the standard.



Cisco is ISO 27001 certified. This is a link to our ISO 27001 certificate: <https://www.cisco.com/c/en/us/about/approach-quality/iso-27001.html>

ISO 27001:2013 The Cisco Intersight Platform has completed its ISO 27001:2013 First Surveillance Audit from the external certification body/auditor Coalfire, and the certificate issued has been uploaded to [Trust Portal site](#).

The First Surveillance Audit included a review of the establishment and overall operating effectiveness of control areas that form Cisco Intersight's Information Security Management System.

FISMA (Federal Information Security Management Act) Cisco HyperFlex has not participated in a FISMA audit to date. For FISMA, Federal information systems must meet the minimum-security requirements. These requirements are defined in the second mandatory security standard required by the FISMA legislation, [FIPS 200](#) "Minimum Security Requirements for Federal Information and Information Systems". Organizations must meet the minimum security requirements by selecting the appropriate security controls and assurance requirements as described in [NIST Special Publication 800-53](#), "Recommended Security Controls for Federal Information Systems".

FedRAMP (Federal Risk and Authorization Management Program) Cisco HyperFlex is not FEDRAMP authorized for cloud-based security because Cisco is not a public cloud provider. For FEDRAMP in particular, the onus will fall on the cloud provider (Google/Azure/AWS etc.) unless you mean to include private cloud in a FEDRAMP assessment or authorization. If this is the case, then the private cloud will need to meet the FedRAMP standards and a POC is recommended.

As Intersight adoption increases, requests to host Intersight in different regions / clouds with the required certifications are desirable. Cisco has initiated a FedRAMP gap analysis effort to pursue this certification.

SOC 2 Type 2 (Service Organization Control 2) The Intersight team has completed the requirements for a Service Organization Control (SOC) 2, Type 2 External Audit covering Cisco's SaaS platform. Developed by the American Institute of CPAs ([AICPA](#)), SOC 2 is a common compliance framework specifically designed for service providers that store customer data in the cloud. It requires companies to establish long-term, ongoing internal practices regarding the security of customer data.

A SOC 2 Type 2 report is an industry-recognized report that provides reasonable assurance that platform controls are suitably designed, operating effectively as necessary, and meet the following criteria:

- **Security:** The service is protected against unauthorized access.
- **Availability:** The service is available for operation and use as committed or agreed upon.
- **Confidentiality:** The service adheres to privacy commitment.



A detailed report with right privileges can be viewed from [Trust Portal](#). There have been no exceptions reported during the certification process.

TEMPEST (Telecommunications Electronics Material Protected from Emanating Spurious Transmissions) TEMPEST is a certification that tests for electromagnetic pulse emanations, i.e, Emissions Security (EMSEC). While the actual standard remains classified, NSA's program information can be found here: <https://apps.nsa.gov/iad/programs/iad-initiatives/tempest.cfm>

TEMPEST certification is for electromagnetic emissions that can be monitored by outsiders (monitors, keyboards, server enclosures, etc.) to reconstitute the images or data rendered on these devices. UCS hardware (and by extension, HX) is not TEMPEST compliant. This compliance is typically only relevant to desktop class equipment in open workspaces – and if needed in certain environments agencies will place servers and other appliances in TEMPEST approved cabinets or rooms that shield everything contained inside.

Third party companies, however, often take Cisco gear, shield it, and get it TEMPEST certified. This is a presentation from a company that has TEMPEST certified a number of Cisco products: https://www.fbcinc.com/source/virtualhall_images/DOS_February/API/API_Product_Presentation.pdf

IAVA (Information Assurance Vulnerability Alert) patches are routine alerts. They are part of the IAVM (Information Assurance Vulnerability Management) Program and detail vulnerability fixes that are deemed critical for all systems in an environment by the DoD from the DoD CERT list. If a vulnerability is on IAVA's list, it will get sent to the admins that are signed up by an organization to receive them and they must be fixed to remain in compliance. Tenable IO (see the Vulnerability Scanning section) scans will pick these up and, based on severity, will be remedied in patch releases.

HIPAA (Health Insurance Portability and Accountability Act) requires healthcare organizations use data encryption technology to protect sensitive patient information. However, the law does not specify which types of encryption to use in order to accomplish this task. Key management mechanisms are not specifically called out either. In these respects, HXDP satisfies the HIPAA requirements. HXDP, however, is not officially certified with HIPAA because a fully compliant solution includes all elements of the ecosystem. HXDP would qualify as a compliant component.

HIPAA key management [does not require](#) separate KMS vendors for HX SWE and other environment components.

UK TSR Compliance

There is no mention specific to HCI being excluded by name in the legislation. The summary below is for a standalone cluster deployment a telecom might use in their own datacenters. Intersight requirements as they relate to the telecom services have not currently been evaluated. Large sections of the legislation and implementation is concerned with high-risk vendors. Cisco is not considered a high-risk vendor where HX is concerned. This designation is mostly centered on nation-states.

This [summary of the NCSC analysis](#) (pg 17-19) discusses virtualization items (7.4):

The software requirements detail the mitigations that the hypervisor operating system and software should implement. These mitigations focus on helping prevent known exploitation vectors from being available. The architectural requirements detail the mitigations that the operator should follow to securely architect their virtualised infrastructure. These mitigations focus on highlighting best practice architectural patterns around micro segmentation, secure administration, and patching. Combining these mitigations with a secure management plane (as defined in Section 2) will help an operator build and maintain a secure virtualisation fabric to support their network functions.

The management plane summary (pg. 18 section 7.2) is directly within HX capability (FIPS compliance, admin shell, isolated storage network) as well:

The primary intent of the TSRs relating to the management plane is to segregate critical management functionality from networks with direct access to the internet. Additionally, they contain principles to ensure that management is performed securely, e.g., by using well secured protocols and tightly controlling the network traffic permitted between management endpoints and equipment.

The [UK Telecoms Report](#) actually calls out the trend toward virtualization, service concentration, and supply chain resilience and security: pgs 12-13 (supply chain), pgs 22-23 (virtualization), 29-30 (threats, risks), 41, 43, 44 (requirements from the Telecom)

Throughout, there is a strong emphasis in the requirements around supply chain tracking and verification, which for HX (via USC) is well developed.

This legislation is summarized below (from [here](#)):

- **Specific security requirements:** This pillar requires all telecommunications providers to strengthen specific areas such as access controls and data confidentiality, which may be exploited to compromise the infrastructure, leading to downtime and/or security issues.

HX meets these requirements (especially with Software Based Encryption). Other Cisco components in the ecosystem can serve other needs and work with HX (e.g., on-wire encryption, firewalls (physical or virtual), etc.).

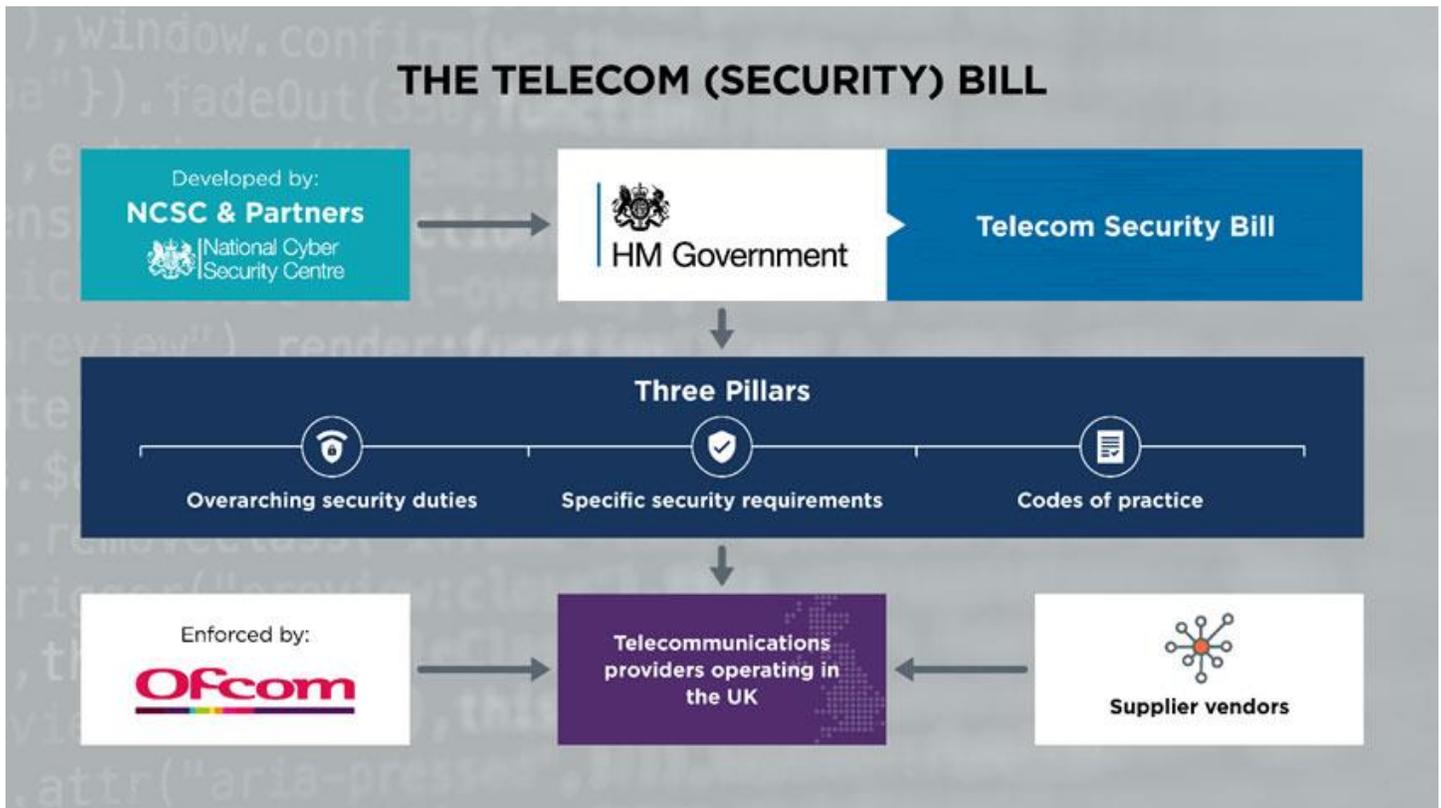
The rest of the requirements are around the telecom itself being responsible for choosing vendors that meet these requirements, implementing these measures, and reporting to users if there are any issues or security vulnerabilities.

OVERVIEW

What are the implications for UK Telecommunications providers?

The new regulations will require all telecommunications providers to demonstrate to Ofcom that they have maximized the cyber protection and resilience of their networks and optimized their security procedures. Ultimately, this will require telecommunications providers to:

- Understand and manage supply chain risk**
 Telecommunications providers delivering services to UK subscribers will now need to identify, document, report and respond to the threats posed by high-risk vendors to ensure the security of all software and hardware deployed in their estate across the entire system development life cycle (SDLC).
- Develop and sustain service resilience**
 Telecommunications providers will also need to ensure that critical aspects of their service are not reliant on international connectivity as part of their business continuity plan. To address this issue, telecommunications providers need to understand their current architectural deployment and which components need to be re-deployed in order to support these business continuity requirements.



The Telecom (Security) Bill was developed by the National Cyber Security Centre (NCSC) and key partners from government, academia, and industry. The bill places a new requirement on telecommunications providers to increase the security of their operations under three core pillars:

- Overarching security duties:** These will require all telecommunications providers to take appropriate and proportionate measures to identify and reduce all potential risks of systems being compromised.
- Specific security requirements:** This pillar requires all telecommunications providers to strengthen specific areas such as access controls and data confidentiality, which may be exploited to compromise the infrastructure, leading to downtime and/or security issues.



- **Codes of practice:** This pillar addresses the appropriate timeframe for compliance for different telecommunications providers and establishes metrics and measures to be enforced.

Telecommunications providers who are audited by Ofcom may be issued with enforcement penalties of up to 10% of a provider's turnover or £100,000 per day for non-compliance. Telecoms supplier vendors are also affected by the Bill as telecommunications providers will likely flow down responsibilities as part of the service they are procuring.

NERC CIP (North American Electric Reliability Corporation Critical Infrastructure Protection) is centered on the physical security and cybersecurity of assets deemed to be critical to the electricity infrastructure. [There are currently 11 CIP standards subject to enforcement](#), governing topics from system security management to recovery plans. NERC CIP compliance is more about policy and procedure than technology, and the responsibility of compliance is on the utility company not the technology provider. So, there is not a "FERC/NERC compliant HCI", per se. The idea is identifying capabilities that helps the customer facilitate compliance, and there are multiple HX security features, system configurations and hardening, as well as continuous security monitoring and advisories that are pertinent toward that goal.

A couple of examples (note: not exhaustive):

- CIP-007-6 R1 – Ports and Services
 - Requirement: CIP-007-6 Part 1.1 requires enabling only logical network accessible ports that have been determined to be needed by the Responsible Entity.
 - Mitigations: The HX Data Platform Hardening Guide provides guidance on port requirements, STCLI security commands for whitelisting, setting up IPtables on HX nodes to secure network traffic, etc.
- CIP-007-6 R2 – Security Patch Management
 - Requirement: CIP-007-6 Part 2.1 requires a patch management process for tracking, evaluating, and installing cyber security patches for applicable Cyber Assets.
 - Requirement: CIP-007-6 Part 2.2 requires, at least once every 35 calendar days, to evaluate security patches for applicability that have been released since the last evaluation from the source or sources identified in Part 2.1.
 - Mitigations: proactive PSIRT advisories publish guidance on current security vulnerabilities and mitigations that impact HX
- CIP-007-6 R4 – Security Event Monitoring
 - Requirement: CIP-007-6 Part 4.1 requires logging events for identification and investigation of cybersecurity incidents that includes minimally: detected successful logins, detected failed access and login attempts, detected malicious code
 - Mitigations: centralized audit logging in HXDP, position Next Gen Firewall for threat protection

The idea here is a comprehensive audit record and well-defined RBAC roles and division of user duties. Continuous monitoring would be a solution type of responsibility that would be handled with ecosystem components like Tetration and Splunk (analysis of syslog). Here is the overview of NERC information system compliance:

Energy producers and distributors that make up the bulk electric system for North America have multiple IT security and compliance challenges, which range from protecting consumers' payment card data and complying with the Payment Card Industry Data Security Standard, to adhering to the general internal audit control and disclosure requirements under Sarbanes-Oxley. In addition, utilities and firms that fall under the authority of the Federal Energy Regulatory Commission (FERC) must meet the cyber security standards of the FERC's certified Electric Reliability Operator (ERO), the North American Electric Reliability Corporation (NERC).

Just as physical surveillance tools such as video cameras are a critical part of physical security controls under NERC, the core technical requirements for cyber security as outlined in NERC CIP Standards 002-009 and other associated guidance from NERC require accountability throughout the authentication, access control, delegation, separation of duties, continuous monitoring and reporting of electronic access to critical infrastructure. And specific requirements from NERC CIP 005, 004, 007 and 008 taken together establish a clear obligation that all electronic access be audited, monitored, and archived in such a way that an organization can reproduce detailed privileged user sessions 24 hours per day, 7 days per week. This continuous monitoring requirement would be difficult to achieve with a combination of manual processes and system-level logs, which often do not tie actions to a unique identity.

Additional specific details are available from NERC itself:

<https://www.nerc.com/pa/comp/Pages/default.aspx>



<https://www.nerc.com/pa/comp/guidance/Pages/default.aspx>

CNSA (Commercial National Security Algorithm) is a schema that is called out by the NSA via this IETF memo:

<https://tools.ietf.org/id/draft-jenkins-cnsa-cmc-profile-00.html>

It describes which algorithms should be in use and what their profiles should look like. It is intended to give guidance for secure and interoperable communications for national security reasons:

"This document specifies a profile of the Certificate Management over CMS (CMC) protocol for managing X.509 public key certificates in applications using the CNSA Suite."

Cisco supports both elliptic cryptographic certificates (ECC) and RSA certificates, so this requirement is met:

"Elliptic Curve Digital Signature Algorithm (ECDSA) and Elliptic Curve Diffie-Hellman (ECDH) key pairs are on the curve P-384. FIPS 186-4 [\[DSS\]](#), Appendix B.4, provides useful guidance for elliptic curve key pair generation that SHOULD be followed by systems that conform to this document.

RSA key pairs (public, private) are identified by the modulus size expressed in bits; RSA-3072 and RSA-4096 are computed using moduli of 3072 bits and 4096 bits, respectively."

HyperFlex's FIPS certification via CiscoSSL implements Federally approved crypto modules to satisfy the complexity requirements as well. The fact sheet here lists the approved algorithms:

<https://apps.nsa.gov/iaarchive/customcf/openAttachment.cfm?FilePath=/iad/library/ia-guidance/ia-solutions-for-classified/algorithm-guidance/assets/public/upload/Commercial-National-Security-Algorithm-CNSA-Suite-Factsheet.pdf&WpKes=aF6woL7fQp3dJipHPerrFKTuHeUCZyCdxdcF3A>

CNSA compliance is just a matter of making sure to implement a cryptographic ecosystem according to the CNSA requirements since HX support all the documented methods.

DISA APL (Defense Information Security Agency Approved Product List) This certification is a multifaceted US Federal approval for products to operate in secure environments. It is currently under way with the Cisco Global Certification Team and the HX Business Unit. It is targeted for the HX 5.5.2a release.

Targeted Certifications

Future targeted certifications are always under evaluation with the Global Certification Team. We are in the planning phase for certifying HX with Common Criteria in the HX 5.5.2a release. DISA APL is expected to commence in the Fall of CY 2022.

Counterfeit Prevention

The Cisco Value Chain describes the development model used for all Cisco products, including HyperFlex. Cisco is a leader in industry and international standards on counterfeit reduction and has been engaged in decades-long efforts to prevent and detect the distribution of counterfeit products. Cisco incorporates tools

and processes to prevent counterfeiting—beginning with product development, through the manufacturing process, and in the marketplace.

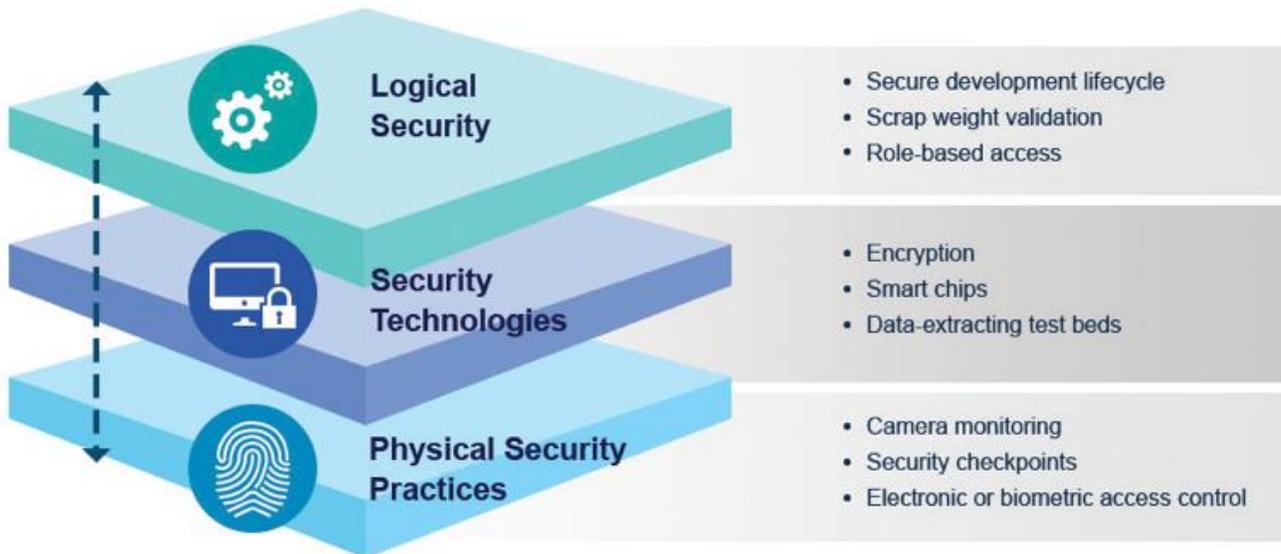
In collaboration with Cisco's Brand Protection, Legal, and other teams at Cisco, an end-user portal has been developed to aid customers in these efforts and can be accessed at: anticounterfeit.cisco.com

Cisco's Brand Protection has conducted numerous investigations into counterfeiting operations and worked with local law enforcement to disrupt those operations. The portal includes examples of the Brand Protection Team's work over the years, and the numerous resources that are available for Cisco customers and partners.

This Value Chain has the following characteristics:

- Comprehensive across all stages of solution's lifecycle
- Multi-layer approach, focused protection against:
 - Source code corruption
 - Hardware counterfeit
 - Misuse of intellectual property

This multilayered approach is shown below.

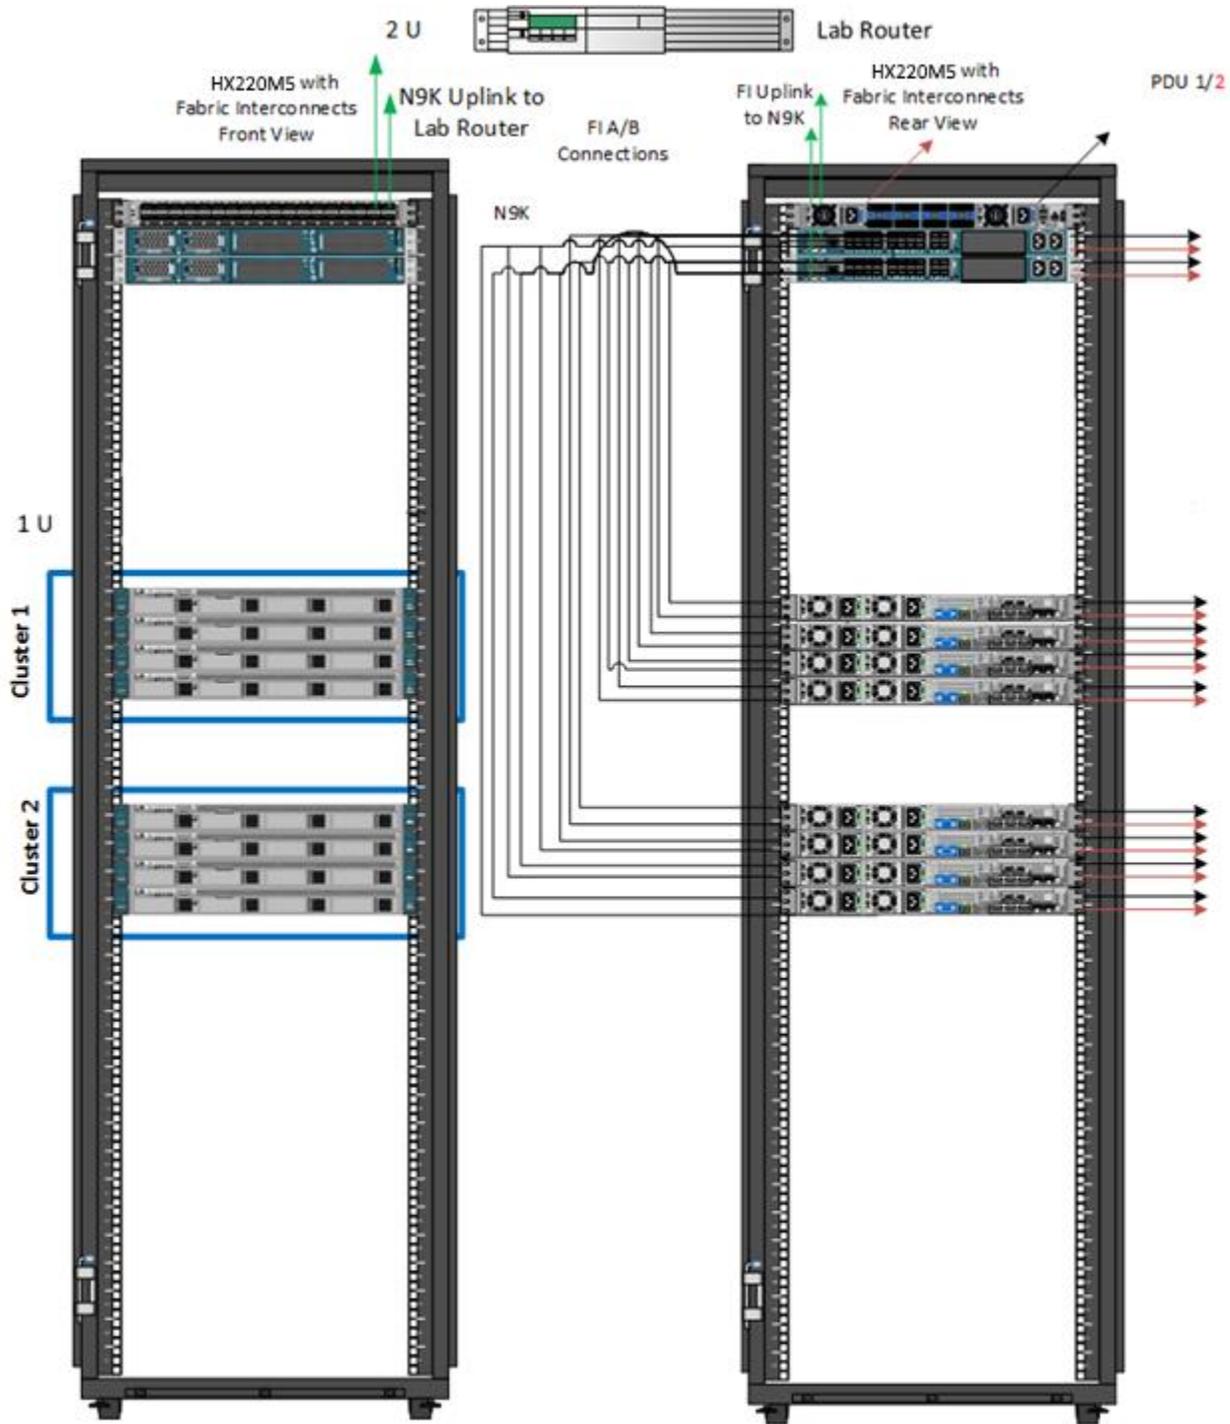


HX Components and Environment

This section details the different components in a typical HX deployment. It is critical to the secure environment that the various parts are hardened as needed.

Solution Components

An HX deployment consists of HX nodes on UCS connected to each other and the upstream switch via a pair of Fabric Interconnects (FIs). There may be one or more cluster and clusters can share the same FIs or be connected to their own, independent set. Clusters can be paired and use HXNR (Native Replication) for protection of VMs. Intervening optimizations appliances may also be deployed to aid with (or monitor or shape) cluster to cluster traffic. The following illustration shows a typical physical layout for this kind of deployment.



Cisco UCS

The physical HX node is deployed on a Cisco UCS 220 or 240 platform in either a hybrid or all flash configurations. A service profile is a software definition of a server and its LAN and SAN network connectivity, in other words, a service profile defines a single server and its storage and networking characteristics. Service profiles are stored in the Cisco 6248/6296, 6332/6332-16UP, and 64xx Series Fabric Interconnects and are



managed via specific versions of UCSM (the web interface for the FI) or via purpose written software using the API. When a service profile is deployed to a server, UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile. This automation of device configuration reduces the number of manual steps required to configure servers, network interface cards (NICs), host bus adapters (HBAs), and LAN and SAN switches.

The service profile for the HX nodes is created during cluster build at install time and is applied to the appropriate devices attached to the FI (identified by PID and associated hardware). These profiles should have their own, easily identifiable name and should not be edited after creation. They are preconfigured by the HX Installer with the settings required for HX to operate securely and efficiently (VLANs, MAC pools, management IPs, QoS profiles, etc.).

It is also worth noting that some larger UCS customer use custom MAC pool and UUID schema for all UCS domain deployments in the data center. Cisco does not support custom naming schemes. HX is an appliance, and to ensure consistent quality, user experience, and full TAC supportability these mundane details have been automated. For UUID, HXDP leverages the hardware derived UUID. For MAC, HXDP has a specific enumeration that cannot be changed.

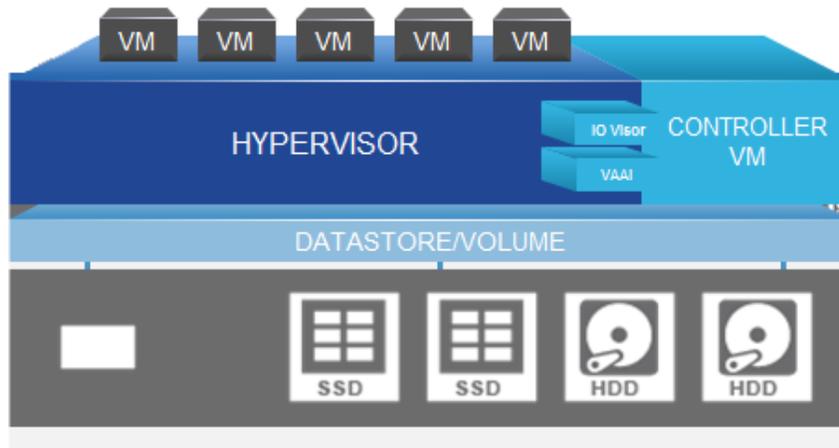
Cisco UCS Fabric Interconnects (FIs)

Cisco UCS FIs are a networking switch or head unit to which the UCS chassis connects. Fabric Interconnects are a core part of Cisco's Unified Computing System, which is designed to improve scalability and reduce the total cost of ownership of data centers by integrating all components into a single platform, which acts as a single unit. Access to networks and storage is then provided through the UCS fabric interconnect. Each HX node is dual connected, one SFP port to each FI for HA. This ensures that all vNICs on the UCS are dual connected as well, guaranteeing node availability. vNIC configuration is automated during HX installation and should not be altered.

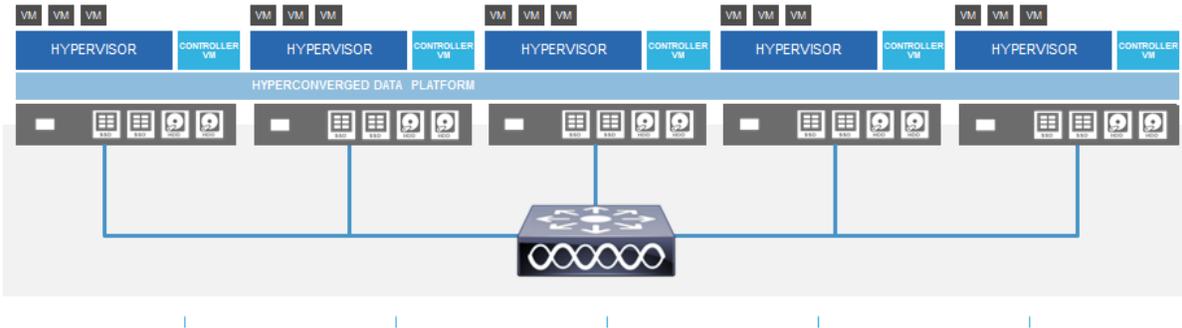
HX Nodes

The HX node itself is composed of the software components required to create the storage infrastructure for the system's hypervisor. This is done via the HX Data Platform (HXDP) that is deployed at installation on the node. The HX Data Platform utilizes PCI pass-through which removes storage (hardware) operations from the hypervisor making the system highly performant. The HX nodes use special plug-ins for VMware called VIBs that are used for redirection of NFS datastore traffic to the correct distributed resource, and for hardware offload of complex operations like snapshots and cloning.

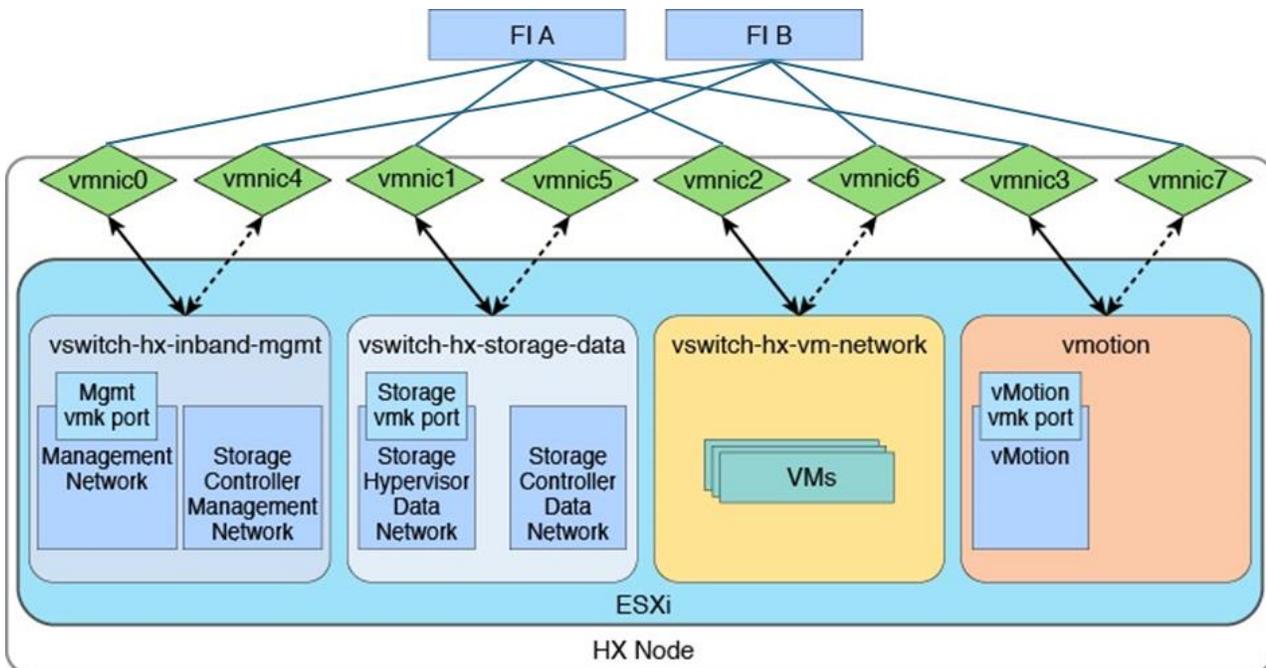
The following illustration shows a typical HX node architecture.



These nodes are incorporated into a distributed cluster as shown below.



The VMNIC ordering in HXDP 4.0.2 and above is shown below:





Management Interfaces: HX Connect and the VMware vCenter Plug-in

HX Connect is the native HTML 5.0 UI for the cluster. The HX vCenter plug-in is another management interface available in vCenter once the cluster is deployed. These are separate interfaces. Both are accessed via HTTPS in a web browser and are subject to the same user management (including RBAC) that is available for the CLI or the API

VMware vCenter

The Cisco HX Data Platform requires VMware vCenter to be deployed to manage certain aspects of cluster creation such as ESX clustering for HA and DRS, VM deployment, user authentication and various datastore operations. The HX vCenter plug-in is a management utility that integrates seamlessly within vCenter and allows comprehensive administration, management, and reporting of the cluster.

It is important to note that all compute and converged nodes must share a single vCenter cluster object for a given cluster. This 1:1 mapping is a requirement today.

VMware ESX

ESX is the hypervisor component in the solution. It abstracts node compute and memory hardware for the guest VMs. HXDP integrates closely with ESX to facilitate network and storage virtualization.

VMs

The HX environment provides storage for the guest VMs deployed in ESX using VLAN segmented networking. The VMs are available for external resources, typical of any elastic infrastructure deployment.

Client Machines

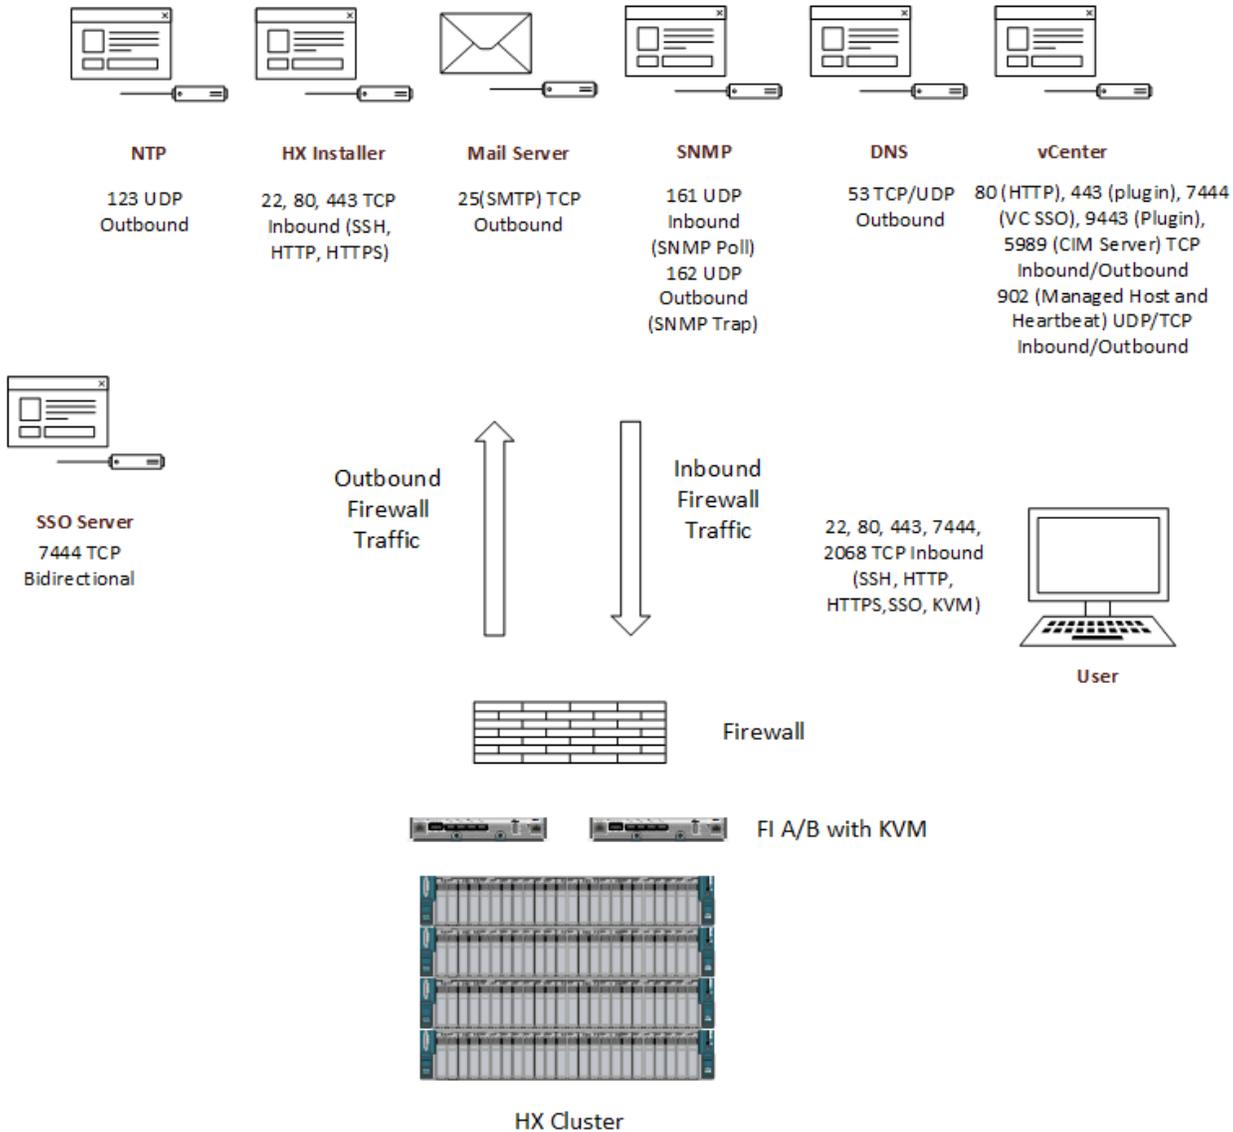
Client machines are defined here as external hosts that need to access resources deployed in HX. These can be anything from end users to other servers in a distributed application architecture. These client's access from external networks and are always isolated from any HX internal traffic by network segmentation, firewalling, and whitelisting rules.

HX Secure Network Environment and Component Requirements

The HX networking environment is segmented and isolated to provide out-of-the-box traffic security. This section identifies the networking communication (port) requirements and offers best practices for the Installer along with information regarding FI traffic and ESX networking (vSwitches). *ICMP is required for various aspects of cluster operation, including installation.* HyperFlex does not use Cisco Discovery Protocol (CDP).

Port Requirements for Communication

The diagram and table below indicate the various components, networking ports, and communication direction for HX.



Note that ICMP is required between CVM IPs and between CVM IPs and vCenter in order to conduct a cluster re-register command. See [Appendix A](#) for a comprehensive table on the port requirements.

Scans Showing Undocumented Ports

There are a few cases where users may scan an HX system and see undocumented or transient ports that appear to be open. This can happen when scanning externally on the management network and it can also happen if users place scanners on the closed data or replication networks. The ports you may see will be well outside of the normal well-known port range (0-1023), often with values between 30000-50000, and are ephemeral ports. An ephemeral port is a short-lived port number used by a transport protocol (TCP/UDP).



Ephemeral ports are allocated automatically from a predefined range by the IP stack software. These ports are not tied to any specific service and will be automatically filtered, deleted, and potentially re-used in the future. Filtered means that a firewall, filter, or other network obstacle is blocking the port so that Nmap, Nessus, Zenmap or similar scanners cannot tell whether the port it is open or closed. Closed ports have no application listening on them, though they could open up at any time. The default firewall rule in HX is to “block all” by default and only “allow selected”. An example of ports that may transiently appear in scans is listed below:

- 43387,45913,49775

These ports pose no security risk.

Port Requirements and Logical Traffic Flow for Replication

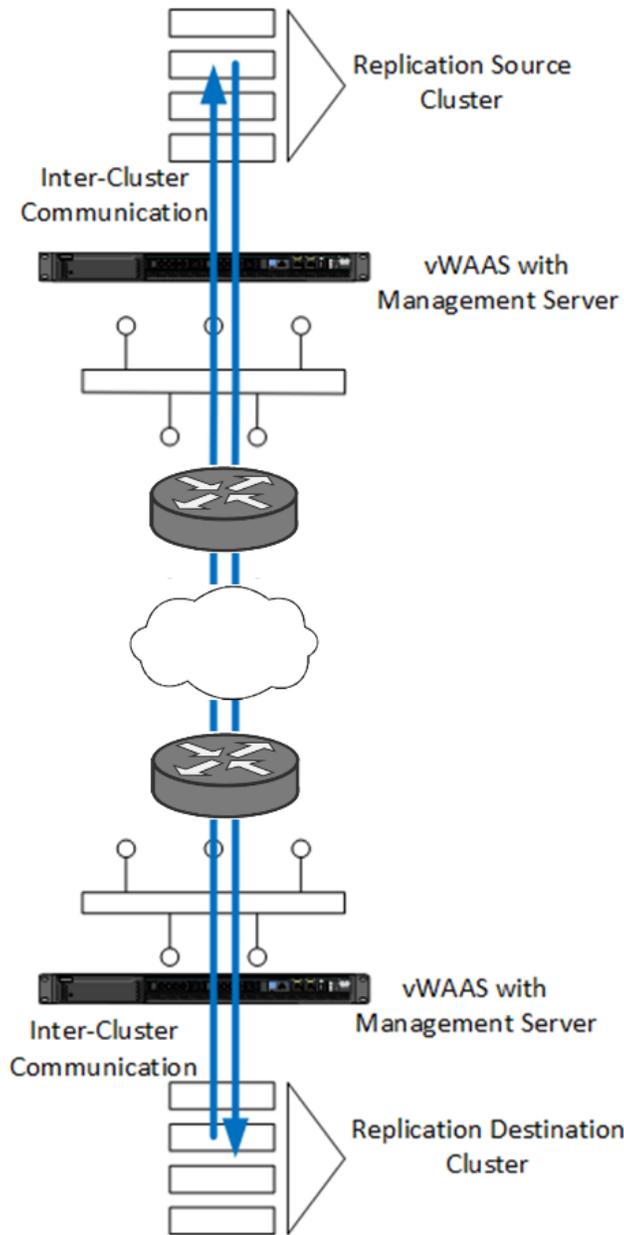
The following ports are opened for inter-cluster communications, during cluster-pairing: 9338, 3049, 9098, 4049, 4059, 8889.

These are the ports that are used in HX Replication:

- ICMP
- datasvcmgr_peer = 9338
- datasvcmgr_peer = 9339
- NRDR = 9350
- scvm (Storage Controller VM) = 3049
- cmap = 4049
- nnnfs = 4059
- replsvc = 9098
- nr (master for coordination) = 8889

Firewall entries are made on the source and destination machine during pairing to allow HX Data Platform access to the system(s) bi-directionally. This traffic needs to be allowed on WAN routers for each HXDP node IP address and cluster CIP-M address.

The following illustration shows the logical traffic flow for replication:



Intersight and PVA/CVA Connectivity Requirements and Resilience

Reference the HX Edge preinstall checklist for Intersight specifics.

Before installing the HX cluster on a set of HX servers, make sure that the device connector on the corresponding Cisco IMC instance is properly configured to connect to Cisco Intersight and be claimed.

- All device connectors must properly resolve *svc.intersight.com* and allow outbound initiated HTTPS connections on port 443. The current version of the HX Installer supports the use of an HTTP proxy.

- All controller VM management interfaces must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of HX Installer supports the use of an HTTP proxy if direct Internet connectivity is unavailable.
- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports (see appendix A).
- **WARNING:** use of [deep inspection \(layer 7 SSL decryption\)](#) by an intervening proxy firewall will prevent HX cluster deployment via Intersight. Disable this functionality for the cluster.

When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.

Post-cluster deployment, the new HX cluster is automatically claimed in Intersight for ongoing management.

Intersight installations require the UCSM or CIMC IP/Network to have SSH access to the ESX and SCVM IPs/network. Any firewalls in this path should be configured to allow the necessary ports

In summary:

Network Communication Requirements for CIMC:

- Communication between CIMC and vCenter via ports 80, 443 and 8089 during installation phase.
- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the Hyperflex Hardening Guide.
- This communication needs to be persistent. It is required for any and all upgrades (including firmware), monitoring, and UI cross-launch.
- CIMC to Intersight should only require 443. Per the preinstall guide:
- *All device connectors must properly resolve `svc.intersight.com` and allow outbound-initiated HTTPS connections on port 443. The current HX Installer supports the use of an HTTP proxy. The IP addresses of ESXi management must be reachable from Cisco UCS Manager over all the ports that are listed as being needed from installer to ESXi management, to ensure deployment of ESXi management from Cisco Intersight.*
- Allow port 22 between the UCSM (or CIMC) VLAN and the ESXi/SCVM management VLAN

Intersight Connectivity Consider the following prerequisites pertaining to Intersight connectivity:

- Before installing the HX cluster on a set of HX servers, make sure that the device connector on the corresponding Cisco IMC instance is properly configured to connect to Cisco Intersight and claimed.
- Communication between CIMC and vCenter via ports 80, 443 and 8089 during installation phase.
- All device connectors must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of the HX Installer supports the use of an HTTP proxy.
- All controller VM management interfaces must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of HX Installer supports the use of an HTTP proxy if direct Internet connectivity is unavailable.

- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the Hyperflex Hardening Guide.
- Starting with HXDP release 3.5(2a), the Intersight installer does not require a factory installed controller VM to be present on the HyperFlex servers. When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve svc.intersight.com and allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.
- Post-cluster deployment, the new HX cluster is automatically claimed in Intersight for ongoing management.

The Cisco Private Virtual Appliance (PVA) and Cisco Connected Virtual Appliance (CVA) are often deployed in air gapped environments and can be used to manage Cisco HX Software Based Encryption (SWE) keys in the same way that regular Intersight does (see the section on SWE). Since the PVA/CVA is local, unlike the cloud based Intersight solution, it is important to take regular backups of the appliance. Since the appliance manages the keys for SWE, a loss of the PVA and keys, will prevent the cluster from becoming available on reboot. To mitigate a PVA/CVA outage:

- DEKs can be exported out of the HX cluster as a last-resort recovery option.
- The PVA/CVA must be backed up regularly, as called out in Intersight documentation: https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/b_Cisco_Intersight_Appliance_Getting_Started_Guide/m_settings_dashboard.html#id_93773

Unicast and Multicast Requirements

Starting with version 3.0(1a), HXDP no longer uses the UCARP protocol and is 100% unicast traffic moving forward.

For previous versions that did use UCARP, since the well-known multicast address of 224.0.0.18 was used, there is no configuration needed on the switches to be able to support HX. This UCARP protocol falls under the IPv4 multicast link-local scope of 224.0.0.0/24. Link scoped multicast packets are flooded throughout the VLAN and IGMP snooping does not take effect on these multicast groups. Hence there is a very small amount of “management multicast” in use, but nothing that requires any network changes or specific infrastructure to support it.

Datstore Access

Access to the HX datstores by client machines is restricted to mounting by HX nodes only. This access is automatically granted during cluster install when the component nodes are identified. Access is also granted or revoked during expansion or removal respectively, when nodes are added or removed from the system. Access to the datstores for migration or backup purposes may be granted via the command line using the



STCLI whitelist command. HX nodes are not listed in the whitelist list because this is a manual, administrative setting for external machine access only. It should only be used during VM ingress/egress from the system as required and the list should be immediately purged once operations are complete.

The mount syntax needs to look like the following in order to work (mount ip:ip:/<datastore> <local dir>/) where the IP is the CIP (not the CIP-M). Here it is in action once the mounting host has been added to the whitelist:

```
kaptain@kaptain-vm:~/temp$ sudo mount 10.a.b.c:10.a.b.c:/ds01 mountpoint/
kaptain@kaptain-vm:~/temp$ su
Password:
root@kaptain-vm:/home/kaptain/temp# cd mountpoint/
root@kaptain-vm:/home/kaptain/temp/mountpoint# ls
auth.log          rhttpproxy.4.gz  vmkernel.4.gz    vprobed.log
clomd.log         rhttpproxy.5.gz  vmkernel.5.gz    vprobe.log
```

Auto Support and Smart Call Home (SCH)

You can configure the HX storage cluster to send automated email notifications regarding documented events. The data collected in the notifications can be used to help troubleshoot issues in your HX storage cluster.

Auto Support is the alert notification service provided through HX Data Platform. If you enable Auto Support, notifications are sent from HX Data Platform to the designated email addresses or email aliases that you want to receive the notifications. Typically, Auto Support is configured during HX storage cluster creation by configuring the SMTP mail server and adding email recipients. Only unauthenticated SMTP is supported for ASUP.

If the **Enable Auto Support** check box was not selected during configuration, Auto Support can be enabled post-cluster creation using the following methods:

Post-Cluster ASUP Configuration Method	Associated Topic
HX Connect user interface	Configuring Auto Support Using HX Connect
Command Line Interface (CLI)	Configuring Notification Settings Using CLI
REST APIs	Cisco HyperFlex Support REST APIs on Cisco DevNet .

Auto Support can also be used to connect your HX storage cluster to monitoring tools.

Smart Call Home is an automated support capability that monitors your HX storage clusters and then flags issues and initiates resolution before your business operations are affected. This results in higher network availability and increases operational efficiency.



Call Home is a product feature embedded in the operating system of Cisco devices that detects and notifies the user of a variety of fault conditions and critical system events. Smart Call Home adds automation and convenience features to enhance basic Call Home functionality. SCH supports a secure proxy for message transfer (see Appendix G). After Smart Call Home is enabled, Call Home messages and alerts are sent when triggered. This includes:

- Automated, around-the-clock device monitoring, proactive diagnostics, real-time email alerts, service ticket notifications, and remediation recommendations.
- Proactive messaging sent to your designated contacts by capturing and processing Call Home diagnostics and inventory alarms. These email messages contain links to the Smart Call Home portal and the TAC case if one was automatically created.
- Expedited support from the Cisco Technical Assistance Center (TAC). With Smart Call Home, if an alert is critical enough, a TAC case is automatically generated and routed to the appropriate support team through https, with debug and other CLI output attached.

See Appendix B and G for additional SCH information.

Installation and ESX Best Practices and Security Considerations

Before conducting any installation, review and complete the pre-installation checklist maintained here:

http://www.cisco.com/c/dam/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/HyperFlex_preinstall_checklist/Cisco_HX_Data_Platform_Preinstallation_Checklist_form.pdf

Cisco HX Installer (HX Installer)

During initial configuration, the cluster is installed on site using the HX installer. This installer can safely be removed from the environment immediately after cluster creation. It is typical for secure environments to isolate the deployment network during installation. In this scenario, the installer is never externally available during configuration. Since it is removed post deployment, installer threat exposure is minimized.

The following services in vSphere must be enabled after you create the HX Storage Cluster in vCenter:

- DRS
- vMotion
- High Availability

The installer verifies that the cluster components are correct and available as needed. This ensures that the deployment has no gaps that could jeopardize security.

- Ensure firmware and BOM compliance
- Deploy and cluster create (requires UCSM credentials for SED)
 - All nodes should be SED capable– no mixing of SED & non-SED drives
- Server Selection to shows for SED Capable nodes and validates non-SED node configurations
- Creates Service profiles



- VLANS
- IP addressing
- vNIC ordering
- QoS configuration
- MAC pools
- Creates ESX vSwitches with appropriate VLANs and address spaces
- Deploys HX Data Platform
- Deploys ESX Plug-ins
- Configures and starts the storage cluster.
- Sets default passwords and generates node-node communication secure certificates

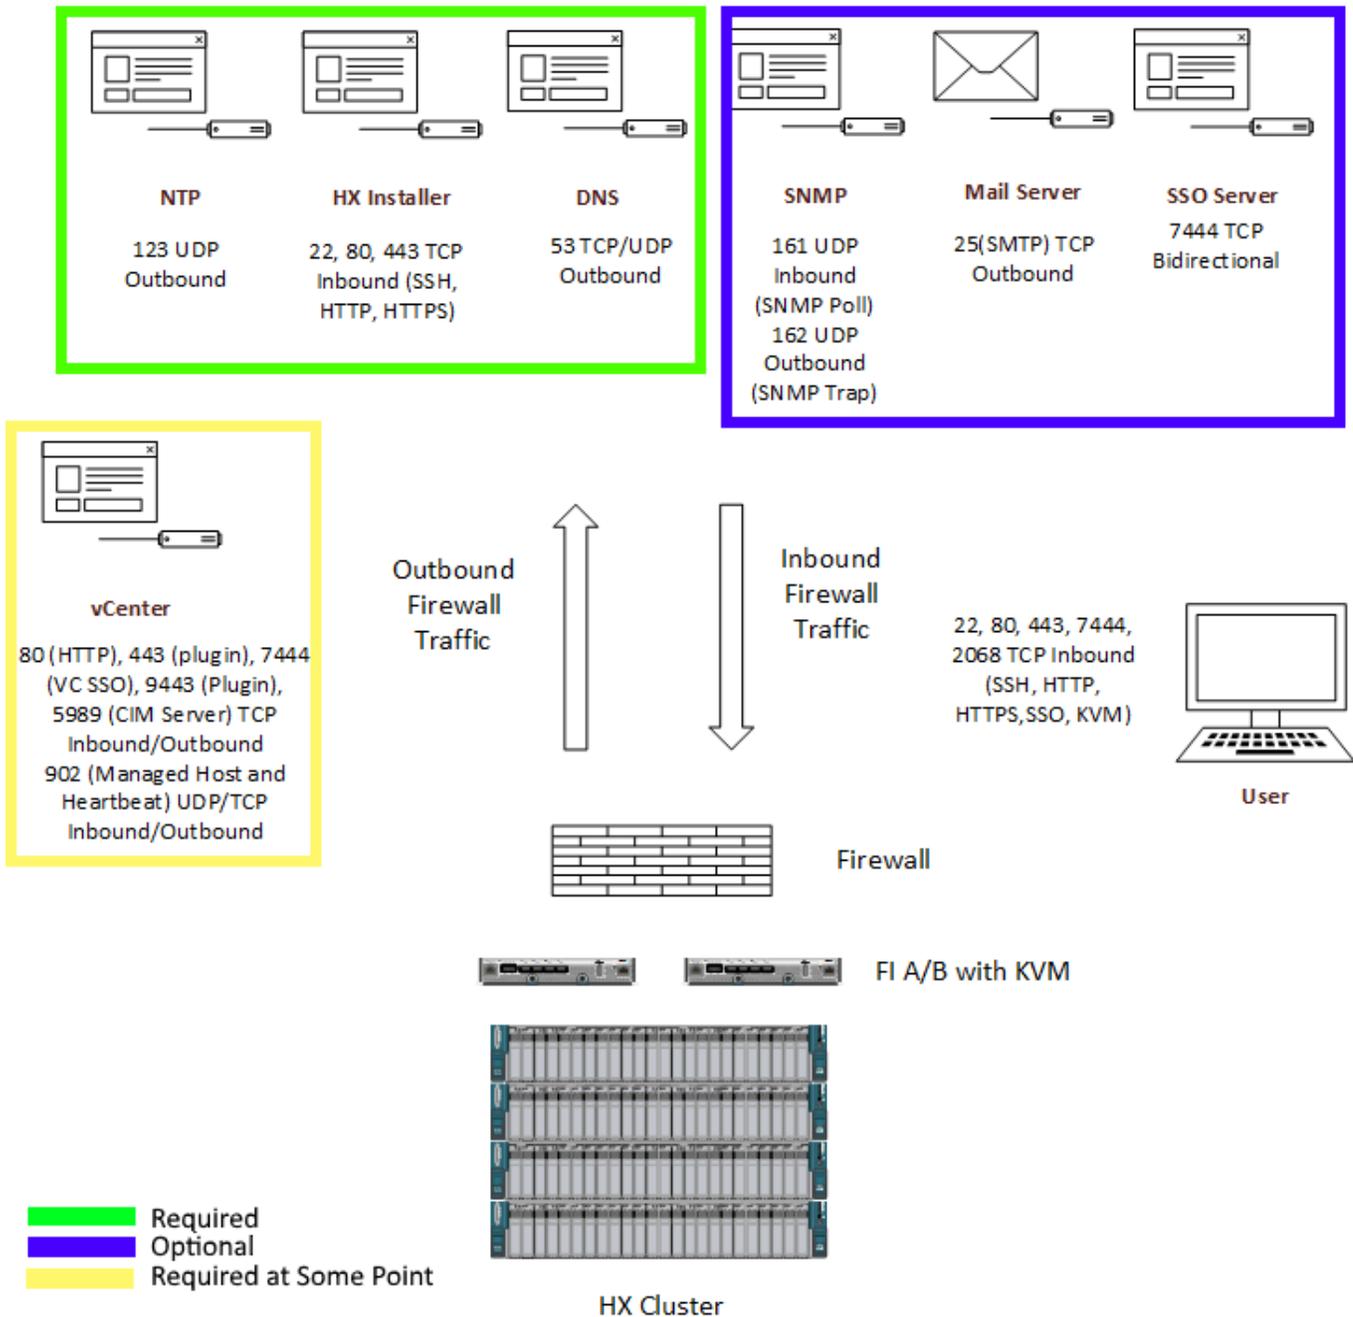
Hard passwords are enforced on HX UI interfaces and HX Data Platform settings during install. Additionally, admission control is disabled by default from the post_install script starting in 4.5(1a). HX does not need admission control to function properly, and it is disabled to reduce the feature-based attack surface. The user can enable it as needed using cluster settings in vCenter.

There is an option during installation to enable persistent root. See the section on Secure Admin Shell Access for details on the protected root account. It is not recommended to enable persistent root. Its intention is to make certain day 2 activities easier to complete in air-gapped environments since transport of the challenge-response can be difficult. If persistent root is enabled, you should disable it as soon as possible.

To revoke normal persistent root, simply su root from the admin shell command line and select the 4th option to revoke persistent root. To revoke persistent root on air-gapped (SLR/PLR license) environment, see Appendix J.

Minimum Infrastructure and Port Requirements for Local Installation

The following diagram shows the minimum required infrastructure needed to conduct a local installation of HX using the on-premises installer.



NTP, HX Installer and a DNS server address are the minimum required components. You will need to follow the nested vCenter procedure if you are leaving that component out at build time.

There are some things that can be eliminated from this diagram at the expense of some function:

- Remove SMTP if you are not using auto support phone home functionality
- Remove SNMP if you are not monitoring anything
- Remove the SSO server if you are only using vCenter credentials



- Remove DNS if you are only using IP addresses. DNS is a required field during install, but if you are not using it and do not intend to use it in the future, you can use a dummy address during installation.

Some things in this diagram can be removed after the deployment is complete to reduce attack surfaces:

- Remove DNS if not in use because you used only IP addresses
- Remove the HX Installer
- Migrate vCenter to HX for a nested deployment (this can be problematic if the cluster is experiencing problems and you need to access vCenter). If you used an external vCenter you will need to deploy a new vCenter on HX and use the `stcli cluster reregister` commands to move the instance to the new VC.

DHCP is only required if you intend to use it for the VM Network segment.

For a HyperV deployment, AD with AD integrated DNS, are both required. vCenter is not required at any point. SCVMM is not required either. The other mandatory components remain in place.

Default Passwords

Once the deployment using the installer is complete, make sure that any default passwords are changed or updated. The ESX hypervisor default password is Cisco123. There is no default set for the HXDP nodes since a hard password is enforced at install. Log in to each ESX node via CLI and update the root password as needed using `passwd root`.

VLANs and vSwitches

VLANs are created for each type of traffic and for each vSwitch. There are typically 4 vSwitches created during the install with associated VLANs for each. The vSwitches are for ESX management, HX management, ESX Data (vMotion), and HX Data (storage traffic between nodes for the datastores). HX Data Platform Installer creates the vSwitches automatically.

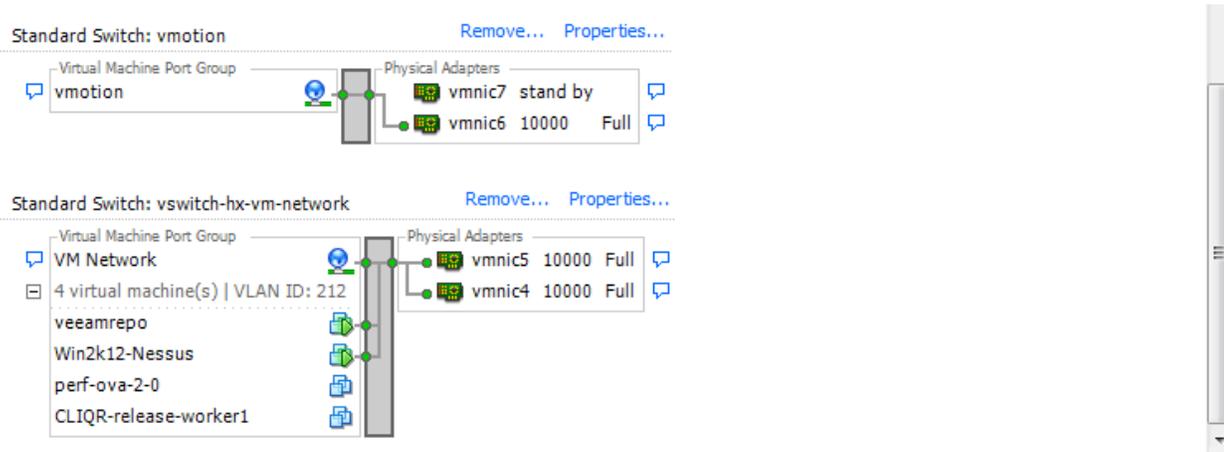
The zones that these switches handle are described below:

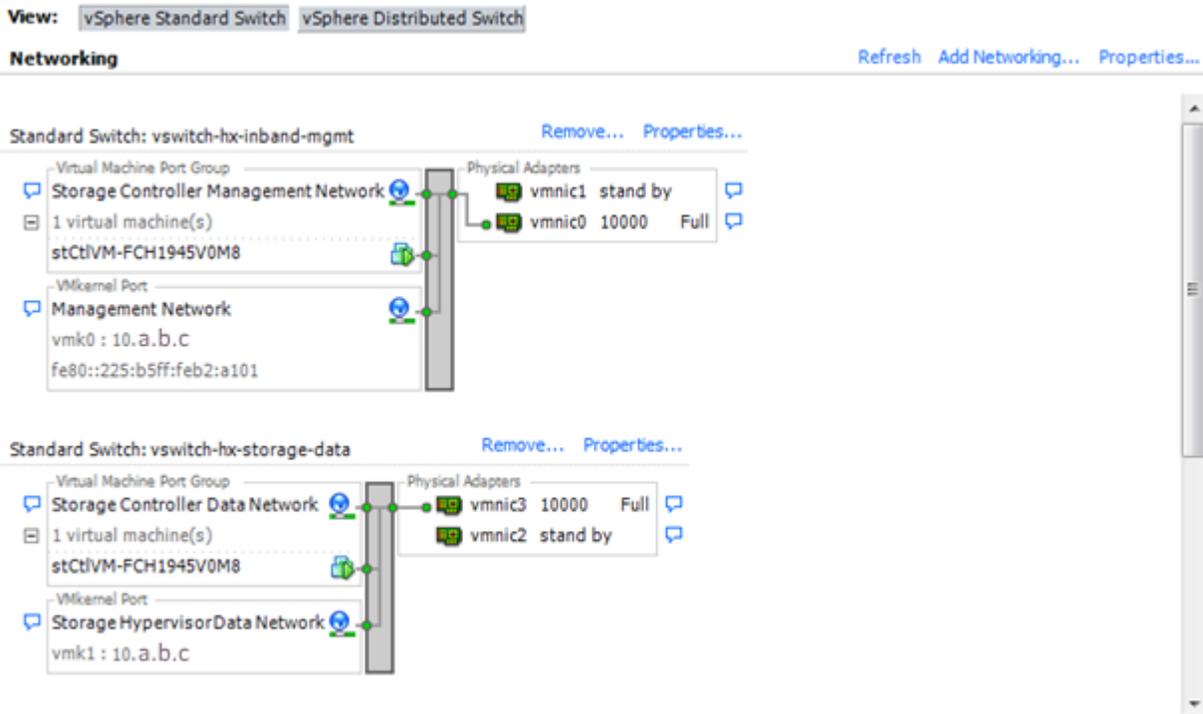
- **Management Zone:** This zone comprises the connections needed to manage the physical hardware, the hypervisor hosts, and the storage platform controller virtual machines (HXDP). These interfaces and IP addresses need to be available to staff responsible in administering the HX system, throughout the LAN/WAN. This zone must provide access to Domain Name System (DNS) and Network Time Protocol (NTP) services and allow Secure Shell (SSH) communication. The VLAN used for management traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching both the Primary Fabric Interconnect (FI-A) and Subordinate Fabric Interconnect (FI-B). In this zone are multiple physical and virtual components:
 - Fabric Interconnect management ports.
 - Cisco UCS external management interfaces used by the servers and blades, which answer via the FI management ports.
 - ESXi host management interfaces.
 - Storage Controller VM management interfaces.
 - A roaming HX cluster management interface.
- **VM Zone:** This zone comprises the connections needed to service network IO to the guest VMs that will run inside the HyperFlex hyperconverged system. This zone typically contains multiple VLANs that are trunked to the Cisco UCS Fabric Interconnects via the network uplinks and tagged with 802.1Q VLAN IDs. These

interfaces and IP addresses need to be available to all staff and other computer endpoints which need to communicate with the guest VMs in the HX system, throughout the LAN/WAN.

- **Storage Zone:** This zone comprises the connections used by the Cisco HX Data Platform software, ESXi hosts, and the storage controller VMs to service the HX Distributed File system. These interfaces and IP addresses need to be able to always communicate with each other for proper operation. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI-A from FI-B, and vice-versa. This zone is primarily jumbo frame traffic therefore jumbo frames must be enabled on the Cisco UCS uplinks. In this zone are multiple components:
 - A vmkernel interface on each ESXi host in the HX cluster, used for storage traffic.
 - Storage Controller VM storage interfaces.
 - A roaming HX cluster storage interface.
- **vMotion Zone:** This zone comprises the connections used by the ESXi hosts to enable vMotion of the guest VMs from host to host. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI-A from FI-B, and vice-versa.

These vSwitches and their associated port groups are tied to a pair of VNICs on each node in an active/standby mode for HA. Their typical networking configuration is shown below:





For an in-depth discussion of Virtual Distributed Switches (VDS) with HX, see the following resource:

<http://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/whitepaper-c11-737724.pdf>

The question often arises, “In a multi-cluster setup, does each HX cluster need to have separate VLAN/Subnets for the storage management and storage data interfaces?” In other words, would there be issues if the same VLAN/subnet is used for each cluster for storage management and storage data interfaces on the controller VM? It is recommended that a unique data VLAN per cluster is used as a best practice. This ensures data is secured within the cluster and there is not contention or broadcast traffic from other clusters on the same network.

However, this is not a hard requirement, and it is possible to put multiple clusters on the same storage VLAN, but you risk performance issues in heavily loaded environments. It is worth noting that early cluster deployments required that the cluster management IP (CIP) have a unique IP in the last octet. For example, if you had a /16 subnet, you should not use 172.16.100.10 and 172.16.101.10 as two cluster management IPs within the same VLAN. The installer has a check to detect this for modern releases, but you should avoid this situation as a best practice.

FI Traffic and Architecture

Traffic through the FIs comes in two general flavors. Intra-cluster traffic (between nodes), and extra-cluster traffic (client machine or replication related). All the FI configurations are managed, accessed, and modified through Cisco UCS Manager (UCSM).



UCSM Requirements

UCSM is the interface used to set up the FIs for Cisco UCS Service Profiles and for general hardware management. During installation, the HX Installer verifies that the appropriate UCSM build is in place for HX, and that the hardware is running a supported firmware version. You are given the option to upgrade these at installation if needed.

Cisco recommends disabling Serial over LAN (SoL) once the deployment is complete since it is no longer needed for ESX configuration. It is also recommended to change any default or simple passwords that were used. Be aware that if you disable SoL, cluster expansion will fail during the Hypervisor Configuration step. You will need to re-enable before continuing.

VNICs

For an in-depth discussion of vNIC see the following:

<https://supportforums.cisco.com/document/29931/what-concept-behind-vnic-and-vhba-ucs>

The VNICs for each vSwitch are in a predefined order and should not be altered in UCSM or ESX. Any changes to these (including active/standby status) could affect HX functionality.

East-West Traffic

East-West traffic on the FI is networking traffic that goes between HX nodes. This traffic is local to the system and does not travel out of the FI to the upstream switch. This has the advantage of being extremely fast by virtue of its low latency, low hop count, and high bandwidth. It also means that this traffic is not subject to external inspection since it never leaves the local system.

North-South Traffic

North-South traffic on the FI is networking traffic that goes outside the FI to an upstream switch and/or router. North-South traffic occurs during external client machine access to HX hosted VMs or for HX access to external services (NTP, vCenter, SNMP etc.). This traffic may be subject to VLAN settings upstream.

Upstream Switch

Configure the upstream switches to accommodate non-native VLANs. HX Installer sets the VLANs as non-native by default.

VLANs

The best practice is to use a separate subnet and VLAN for each of the networks. It is possible to use the same VLAN across multiple clusters connected to the same pair of Fabric Interconnects, for example: Management, vMotion, and VM guest VLANs. This is possible as long as you do not overlap IP addresses. It is, however, strongly recommended to keep the HX storage VLAN unique per cluster to ensure that storage traffic is secure and isolated. If deciding to reuse the storage VLAN against best practices, be extremely vigilant to avoid duplicate IP addresses. A duplicate IP can and will disrupt existing storage traffic on all clusters sharing the resource. Sharing storage VLANs is against best practices for the following reasons:

- Security. Storage traffic (not data at rest, but on-wire data) is unencrypted in this non-routed VLAN. This traffic should be segmented and isolated to prevent any eavesdropping.



- Isolation. Noisy or flooding ports can potentially disrupt traffic on the VLAN. Isolation prevents a problem from affecting more than one cluster.
- Debugging. Isolation dramatically aids in ease of troubleshooting.
- Risk. Humans make mistakes. Separate the traffic to avoid the human factor.

Do not use VLAN 1, the default VLAN, because it can cause networking issues, especially if Disjoint Layer 2 configuration is used. Use a different VLAN.

Disjoint L2 Networks

Please make sure to read and understand the following disjoint layer two document if this is a requirement in your environment:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-692008.html

You can just add new vNICs for your use case. We support the manual addition of vNICs and vHBAs to the configuration. Please see the HX VSI CVD:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/HX171_VSI_ESXi6U2.html for step by step instructions on how to do this safely. Follow the same procedures outlined in the CVD. Please do not use pin groups, as it may not properly prune the traffic and can cause connectivity issues as the designated receiver may not be set correctly.

Cisco HyperFlex Edge (HX Edge)

Typical HX Edge deployments use a trunk port configuration on the top of rack switch(es). VLAN trunking should limit the allowed VLANs to those required for the HyperFlex services and user VMs. By default, the switches will allow all VLANs to pass and could pose a security risk of allowing unfettered network access. See the Cisco HyperFlex Edge Deployment Guide for sample configurations that use "switchport trunk allowed VLAN" commands.

For HX Edge configurations with the add-in PCIe Quad port NIC, ensure any unused Ethernet ports remain disconnected from any virtual switches in ESXi. This will prevent unauthorized access to the virtual switching environment.

SED deployments are currently not supported with HyperFlex Edge. VM encryption by virtue of 3rd party encryption clients will work to encrypt the VMs deployed on Edge. Vormetric and Gemalto Safenet provide such clients. VMware vSphere 6.5 also incorporates a VM encryption capability that should work but has not yet been officially qualified.

HX Data Security

HX data-at-rest security is accomplished via Secure Encrypted Disks (SEDs) or with HX software-based encryption. SEDs are managed by Cisco HX Connect in conjunction with UCSM and local or remote key stores



using the Key Management Interoperability Protocol v1.1. HX SWE (Software Encryption) utilizes the built-in KMS within Intersight. See the Encryption FAQ for a comprehensive treatment of SEDs with respect to HX:

https://www.cisco.com/c/dam/en/us/support/docs/hyperconverged-infrastructure/hyperflex-hx-data-platform/HX_Encryption_FAQ.pdf

See the [Encryption Whitepaper](#) for an overview of HyperFlex Software Based Encryption:

Encryption Services

There are various encryption capabilities that have been developed with the CSDL guidelines in place. These include Self Encryption Drives (SEDs) and Cisco HyperFlex Software Based Encryption (SWE), which is a native feature of the HyperFlex Data Platform. Both are data-at-rest (DARE) implementations. Cisco has also qualified various Key Management solutions using VM level encryption from 3rd party partners like Gemalto and Vormetric (both parts of Thales as of this writing) for SED based encrypted clusters. These various key managers are only for SED based systems. Cisco's software encryption solutions use the Intersight integrated key manager.

Data-at-rest encryption can take place at any of several points in the write IO process and on various portions of the written data including the entire content or subsets therein. The types of DARE you might encounter are:

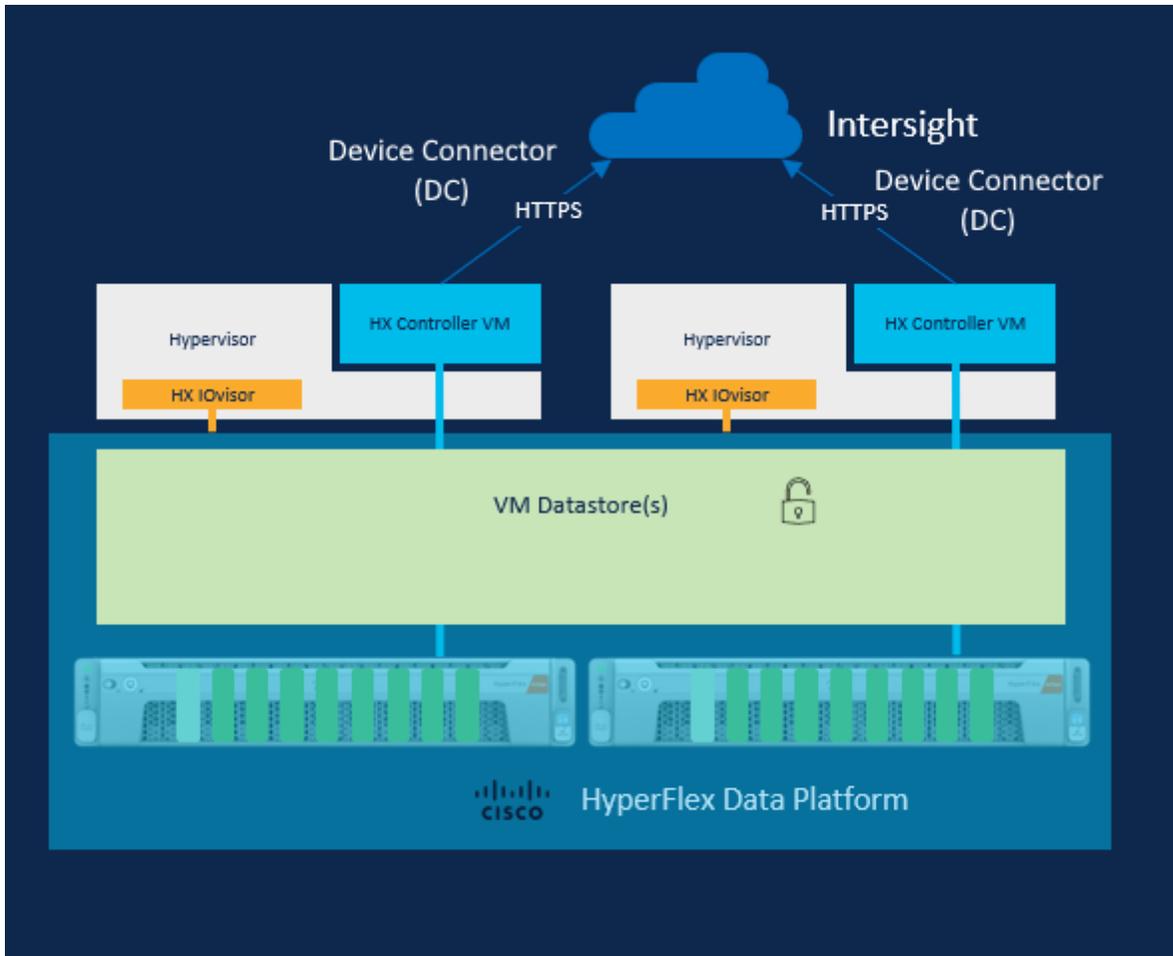
- Application-Level
 - Encryption occurs within the application before data is transmitted or stored
 - Encrypted content in this scenario can have fine-grained boundaries
 - Example: individual fields in a database
- Database-Level
 - Encryption occurs on a subset (tables or columns) or the entire database
 - Utilizes transparent data encryption from database vendors before data is stored
- File-Level
 - Encryption occurs at a file or volume level by agents of the operating system intercepting IO and applying encryption policies
- Disk-Level
 - Full disk encryption or self-encrypting drives (SEDs) enables encryption in hardware
 - Encryption occurs at the drive controller level when data is written to disk and decrypted when data is read from disk

Cisco HyperFlex native software encryption naturally takes place at the file level due to integration with the hypervisor. This allows the encryption package within HXDP to take full advantage of the storage optimizations that occur during ingest and destage to persistent storage, namely compression and deduplication, respectively

HX Native Software Encryption

HyperFlex Software Encryption (SWE) is enabled using Intersight, Cisco's hybrid cloud operations platform, and is available on all Intersight form factors: Software-as-a-Service, Connected Virtual Appliance, and Private Virtual Appliance. Regardless of whether a cluster is deployed by Intersight or imported/claimed in Intersight after deployment, HyperFlex SWE is enabled on the cluster with a simple one-line command after downloading the package. A few clicks under the cluster's "Operate" tab will activate the ability to create encrypted datastores. Key management for SWE takes place in Intersight and is transparent to the end user. The SWE option truly demonstrates "ease of use".

The figure below shows a 2-Node Edge cluster being claimed in Intersight and using the Cluster's Device Connector (DC) to interface with the Intersight Key Manager.



- Controller VMs request and receive keys from Intersight Key Manager through the Device Connector (DC)
- This is known as the “Key Encryption Key” (KEK)
- The KEK is used to encrypt the Data Encryption Key (DEK)
- The DEK is used to encrypt user data

HyperFlex Data Platform Software Encryption uses industry standard strong encryption algorithms and is compliant with US Federal certification requirements. It also takes advantage of Cisco HyperFlex's unique features and cloud technologies:

- Utilizes FIPS 140-2 compliant 256-bit AES-GCM inline encryption
- Leverages AES NI acceleration
- HXDP Boost Mode compatible
- Encrypts data end-to-end: encrypted at rest (caching and persistent tier) and on the wire during intra-cluster replication (i.e., Replication Factor)
- Manages keys natively with Intersight Key Manager
- Supports encrypted and unencrypted datastores coexisting in the same cluster
- Supports Standard Data Center and Edge clusters



- Supports All NVMe, All Flash and Hybrid drive systems
- Supports HyperFlex inline compression and deduplication optimizations
- Supports KEK rekey
- Supports secure drive erase

Cisco HyperFlex systems using SWE are drive agnostic. They can therefore take full advantage of the range of qualified drive capacities and drive technologies (i.e., All-NVMe, All Flash, Hybrid SFF and LFF). Since SWE takes place in the HyperFlex filesystem stack, there is a small performance impact due to additional CPU utilization in the HyperFlex Controller VMs, on the order of 5-10% depending on workload. If your use case is CPU bound by the Controller VMs for performance, you can easily mitigate the issue using HyperFlex "Boost Mode" whereby additional vCPUs can be assigned to the Control VMs.

SWE is supported by HyperFlex Datacenter (with or without Fabric Interconnects) and HyperFlex Edge cluster configurations. To enable SWE, each HyperFlex converged node system in the cluster must be licensed with either HX Data Platform Datacenter or Edge Premier, and with Intersight Essentials (or a higher tier license). In addition, SWE is an export-controlled commodity and is reviewed before being delivered.

Key Management with HX Software Encryption

HyperFlex Data Platform Software Encryption relies on the Cisco Intersight SaaS platform for key management. It affords the following additional advantages by virtue of the Intersight cloud based key management:

- Protects confidentiality of data at-rest from theft of storage media
- Theft of drives, servers, even clusters
 - Simply unclaim the cluster in Intersight. It cannot be re-claimed and re-gain access to the KEK stored in Intersight without your original passphrase.
- Drives disposed without adequate sanitization
 - Unreadable without access to the key encryption key (KEK) which only the HyperFlex system can access from Intersight

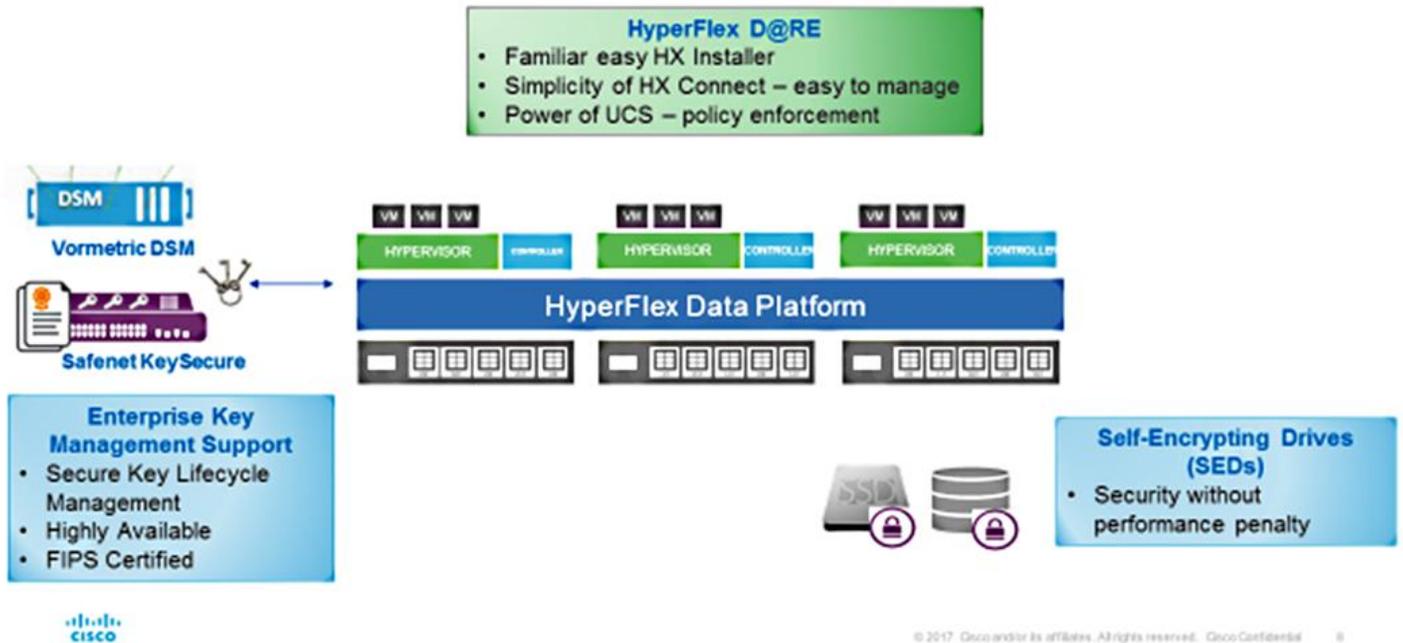
SEDs

A cluster is designated as SED capable or not at installation based on whether the cluster contains SED capable drives or not. After installation, this designation cannot be altered. SEDs provide native data-at-rest encryption, typically using AES 256. All qualified disks are FIPS 140-2 Level 2 validated components for data-at-rest encryption. The hardware encryption is built-in, thereby incurring no deployment overhead. The performance is comparable non-SED system and is transparent to data optimization functions (dedupe, compression).

Two encryption keys are associated with a SED implementation.

- Media Encryption Key – data is always stored in encrypted form
- Key Encryption Key secures the media encryption key

HyperFlex Data-At-Rest Encryption : SEDs



SEDs provide a mechanism for secure erase ensuring security during decommission:

Secure cluster Expansion

- Only SED capable node can be added to HX Cluster with SEDs
- Local key – seamless secure expansion
- Remote key – secure expansion requires lockstep with certificates/key management
- Certificates required to add new node securely
- Deployment will show warning and include steps to proceed and link to UI for certificate download
- User follows steps to upload certificate(s) and continue the deployment

SED on HX Edge is not currently supported (see the HX Edge section above). SEDs are not supported with Hyper-V.

Can access to the SEDs via the CVM be an attack vector if the CVM is compromised? In other words, since the control VM has direct ownership over the HX node disks (through VMDirectPath IO), and since the drives are self-encrypting, does this mean you effectively have unencrypted raw access to the disks through that login? Technically you can access data directly from the root shell, however, you would not be able to do much with this access. Since data is striped per disk, per node, and across file tree vnodes even, you would not be able to reconstruct the data into anything meaningful. You would only have small bits of information in various disks on various nodes. If this is still a concern, you can certainly encrypt via software (see Encryption Partners below) at the VM level, thereby mitigating any fractional data reads from even a compromised CVM.



The ESXi boot volume is not encrypted and there is not an option to encrypt this drive. There is no user data stored on this device. Theft of or raw access to this drive would only expose the ESXi operating system. The CVM root drive is separate from the data drives. The HX system/log drive is not encrypted and there is not an option to encrypt this drive. There is no user data stored on this drive. Theft of or raw access to this drive would only expose time stamps and block locations of the cache and capacity drives, both of which are encrypted.

Connectivity between a SED-enabled cluster and the KMS have a few requirements in environments that are firewall segmented. Access between FI-A/FI-B/VIP and the KMS is not required. Only the CIMC IPs of each node need to access the KMS. Allowing CIMC IPs along with KMS IP and port 5696 is sufficient.

Key Management with SEDs

Configuring encryption services with SEDs supports both local and remote key configurations. If you are not using local keys, then you need to configure a KMIP server. KMIP server key handling for SEDs is performed via encryption partners (Thales Vormetric, Gemalto Safenet, Entrust). The server specifics are entered using the Encryption workflow in HX Connect.

- Data is only as secure as the encryption keys.
- Key management is the tasks involved with protecting, storing, backing up and organizing keys.
- Specialized vendors provide enterprise key management offerings for SEDs
- Intersight provides key management services for HX native software encryption

Key Management best practices:

- Always deploy at least two KMIP servers, clustered for high availability
- No agents or software to deploy for key management
- Configure key backup and recovery
- Self-signed and CA signed certificates can be used

Workflows supported:

- Disable/Enable
- Re-key
- Secure Erase

Secure Erase

HyperFlex provides two secure erase mechanisms. One is specific to SED deployments, the other is designed for non-SED environments. In the SED version, individual disks can be securely erased on a per-disk basis. This functionality is available in HX Connect via the System InformationàDisks UI. This is mentioned in the SED section of this paper. For non-SED environments, secure erase is available using the `secure_disk_erase` command from the Admin Shell.

```
hxshell:~$ secure_disk_erase ?
usage: secure_disk_erase [-h] -d DISK_PATH [-m {0,1,2}] [-p ERASE_PATTERN]
                        [-o OVERWRITE_COUNT] [-s] [-r]
secure_disk_erase: error: the following arguments are required: -d/--disk-path
hxshell:~$
```

Note that the command takes erase patterns and overwrite counts as arguments for spinning disks. The backing tool for this command is `sg_sanitize`. For NVMe/SSD drives the tool used is `NVMe`. Each tool is invoked depending on the device type that is currently being erased.

https://docs.oracle.com/cd/E88353_01/html/E72487/sg-sanitize-8.html

<https://nvmexpress.org/open-source-nvme-management-utility-nvme-command-line-interface-nvme-cli/>

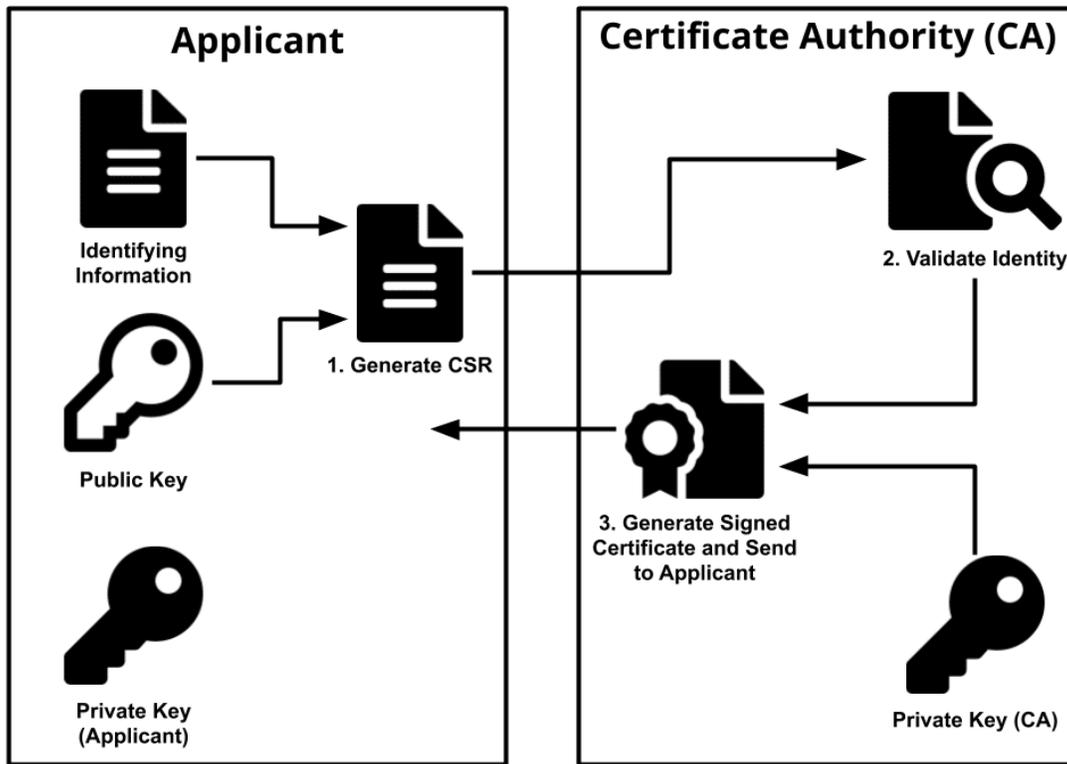
The functionality present here is in compliance with NIST standards. NIST standard and guidance:

<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-88r1.pdf>

Certificate Signing Requests (CSRs)

A component of the remote encryption workflow generates CSRs. CSRs need to be downloaded and signed. Signing can be "self" which refers to signing the CSR with a key you have generated yourself and installed on your KMIP infrastructure. If you are using a Certificate Authority (CA) then you will need to get the CSRs signed with your validated key from the CA.

The diagram below shows the CA/CSR/signing relationship.



Source: <https://rusvpn.com/en/blog/what-is-a-ca-certificate-and-how-does-it-work/>

HyperFlex supports RSA certificates and, by virtue of use of the CiscoSSL module, also supports ECC (Elliptic Curve Cryptography) certificates. RSA is currently the industry standard for public-key cryptography and is used by most SSL/TLS certificates.

A popular alternative first proposed in 1985, is Elliptic Curve Cryptography using a different formulaic approach to encryption. While RSA is based on the difficulty of factoring large integers, ECC relies on discovering the discrete logarithm of a random elliptic curve.

Networking Considerations

When using a KMS (Key Management Server) for remote key management, some additional networking ports may need to be opened. Port 443 is required for policy configuration between the control VMs and UCSM. Additionally, port 5696 is required for TLS communication between the CMIC of each node and the KMS server itself for secure information exchange. See Appendix A.

The HXDP deployment currently requires IPv4. IPv6 may be used for VMs and applications as needed.

Encryption Partners

Cisco HX partners with two industry-leading encryption and KMIP service providers. Cisco HyperFlex systems are KMIP 1.1 compliant. HyperFlex supports several Key Management Server vendors. Recent industry consolidations have seen vendors such as Thales Vormetric (DSM), Gemalto (KeySecure), and Hytrust (KeyControl) subsumed by Entrust. Entrust's KeyControl (formerly Hytrust) and any server vendor that supports KMIP 1.1 should function as expected, but qualifications for newer vendors are always under way.

Gemalto Safenet:

- Enterprise Key Management (EKM) solution

- Single, centralized platform for managing cryptographic keys and applications
- Simultaneously manage multiple, disparate encryption appliances and associated keys through a single, centralized key management platform
- Also provides a high performance encrypt/decrypt engine when combined with SafeNet's Data Protection portfolio

Thales Vormetric:

- Data Security Manager solution
- Single, centralized platform for managing cryptographic keys and applications
- Simultaneously manage multiple, disparate encryption appliances and associated keys through a single, centralized key management platform.
- Also, provides a transparent encryption client for guest VMs.

Note: KMIP 1.1 compliant key managers not explicitly listed as supported require qualification.

Two Versions of Vormetric Data Security Manager



VM Level Encryption with 3rd Party Vendors

Third party VM software encryption works above the HXDP storage layer. Encryption at a VM level of granularity is available with these partner solutions. Note that you can expect there will be no deduplication space savings, since encryption at this level necessarily “makes unique” all data sent to the storage subsystem.

Vormetric Transparent Client

<https://www.thalesecurity.com/products/data-encryption/vormetric-transparent-encryption>

Gemalto:

<https://safenet.gemalto.com/data-encryption/data-center-security/protect-file-encryption-software/>

ESX 6.5 VM encryption:

<https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.security.doc/GUID-B3DA9865-A28F-4EFD-ACF4-CBC8813ED110.html>

Secure Communications

All communication occurring with the HX platform management interfaces is FIPS compliant using SSH or HTTPS. See the section on Management Security above.

Note that when accessing the CLI using SSH and checking versions, the following will show up (depending on HX version):



ssh -V
CiscoSSH 1.6.20, OpenSSH_7.6p1, CiscoSSL 1.0.2u.6.2.374-fips

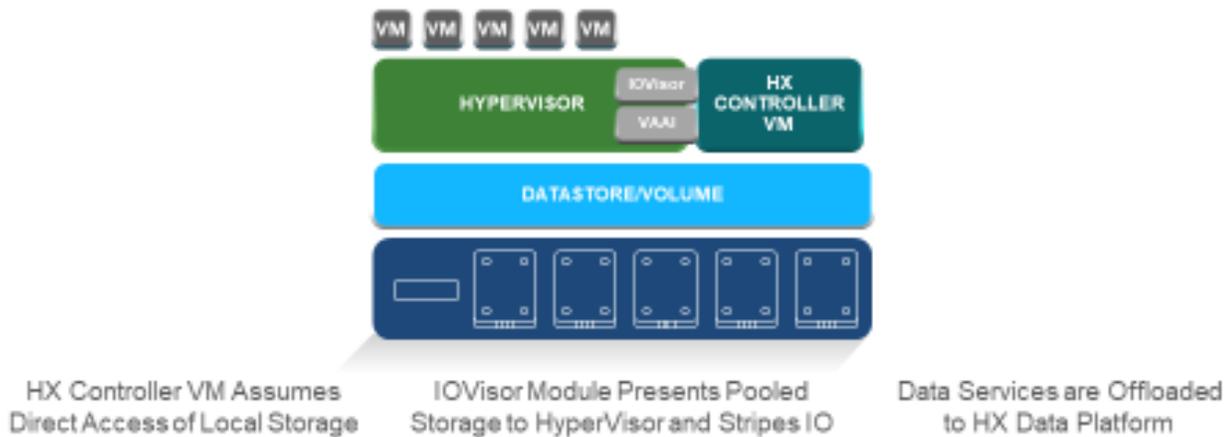
OpenSSH appears here because CiscoSSH is based on OpenSSH, and this versioning information is put in place for reference.

Usage of NFS in HXDP

The HX Data Platform uses a proprietary variant of NFSv3 to present an HX controlled files system to the hypervisor using a plugin called IOvisor. Each node runs an IOvisor instance in order to properly allocate the correct read and write resources in the distributed architecture. This communication path is completely internal, and available only between the hypervisor and the HX CVM present in each node that manages access to the underlying distributed storage. All other components in the system are disallowed from mounting the resource presented by the IOvisor.

Each CVM physically controls the underlying storage hardware and participates in allocating this resource to the distributed file system within which the datastores are created. The diagram below describes the logical placement of the HX IOvisor within the node architecture.

Inside HX Data Platform Node



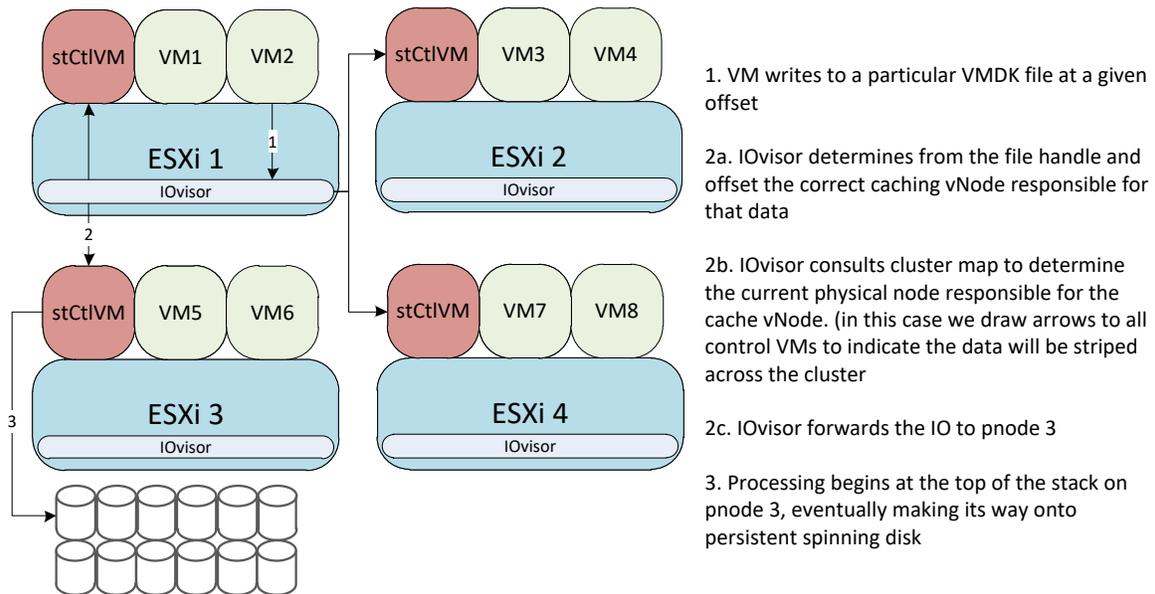
© 2017 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

VM IO is destined for the configured datastore, which is an abstracted container for the storage subsystem. This container is an NFS datastore created on the HX platform. The VMware storage stack utilizes an NFSv3 client to mount the datastore, presented up from the IOvisor.

The IOvisor, sometimes referred to as the SCVM client, lives as a process in user space inside ESXi and can be thought of as a simple NFS proxy. It behaves as a server for the VMware NFS client, while looking like a client to the controller VMs.

The IOvisor is very thin code, on the order of 2000 lines of code. It is designed to be a stateless router for IO and has a very small footprint. It is installed into ESXi as a VMware installation bundle (VIB) that is auto deployed during cluster installation. As such, it is set up to only allow mounts from cluster nodes. This allowed mapping is only updated during node failure or expansion events.

The IOvisor looks at the incoming NFS request and determines which distributed resource it belongs to. The IO routing process can be visualized in the figure below. Notice that NFS is maintained internal to the ESXi-HX CVM communication path only and is never exposed to the guest VMs or any other outside resource.



THE IOVISOR SERVES AS AN "IO ROUTER" IN THE DISTRIBUTED CLUSTER. THE CODE IS A THIN LAYER THAT SITS IN ESXi USER SPACE.

The IO must be directed via the IOvisor to the correct physical node (pNode). Each controller VM queries the Cluster Resource Manager (CRM) via the cluster IP to retrieve the pNode mappings. The communication occurs via a special NFS procedure that has been added as an extension to the base NFS protocol. The mapping table is then cached locally so there is minimal traffic to the cluster IP. The deterministic process to route to the correct node leaves the IOvisor as a stateless IO proxy.

HXDP Runtime Defenses

There are a host of runtime defenses that a system may employ. The HX CVM uses the following:

Secure Storage: using the Trust Anchor Module (TAM) for storing critical data like certificates, anti-counterfeit etc. (see the section on Secure Boot).

Secure JTAG: hardware debugging port. This has been removed to prevent potential access to the port.

FPGA Bitstream Security: encrypted or authenticated FPGA config image. M6 and forward servers utilizes an FPGA for secure firmware authentication.



Boot Integrity Visibility: able to show the boot integrity measurement. HXDP uses secure boot. See the section on Secure Boot.

ASLR: Address Space Layout Randomization to defend code injection attack. This is implemented in the SCVM kernel.

HX Management

There are four relevant management interfaces to consider with HX. There are the UI interfaces (native and vCenter plug-in) and there is the CLI and the REST API.

Management Interfaces

Intersight

Cisco Intersight combines the benefits of cloud-based management with security similar to on-premises systems. This management and automation platform is enhanced by analytics and machine learning techniques to increase efficiency and continuously evolve, so you can manage the growing complexity of your IT infrastructure. The software monitors the health and relationships of infrastructure components that use Cisco UCS or HyperFlex management. Telemetry and configuration information is collected and stored in accordance with Cisco information security requirements. Your data is isolated and displayed to you through an intuitive user interface. Because the software scales easily and frequent updates are implemented without impact, this simplified and consistent infrastructure management approach removes the difficulties of supporting typical tools and appliances.

The Cisco Intersight platform uses layered security. It encrypts data, complies with strict Cisco security and data handling standards, and separates management and IT production network traffic for additional isolation. As a result, you can have confidence that your cloud-based systems management platform offers the strong security you require.

The Intersight platform is developed, integrated, and tested using the Cisco Secure Development Lifecycle guidelines. This secure product development and deployment practice has several components ranging from inherent design and development practices, testing the implementation, and creating a set of recommendations for deploying with maximum security. Cisco development processes are ISO 27001 certified.

Single sign-on Single sign-on (SSO) authentication enables you to use a single set of credentials to log in to multiple applications. With SSO, you can log in to Intersight with your corporate credentials instead of your Cisco ID. Intersight supports SSO through SAML 2.0, acts as a service provider (SP), and enables integration with identity providers (IdPs) for SSO authentication.

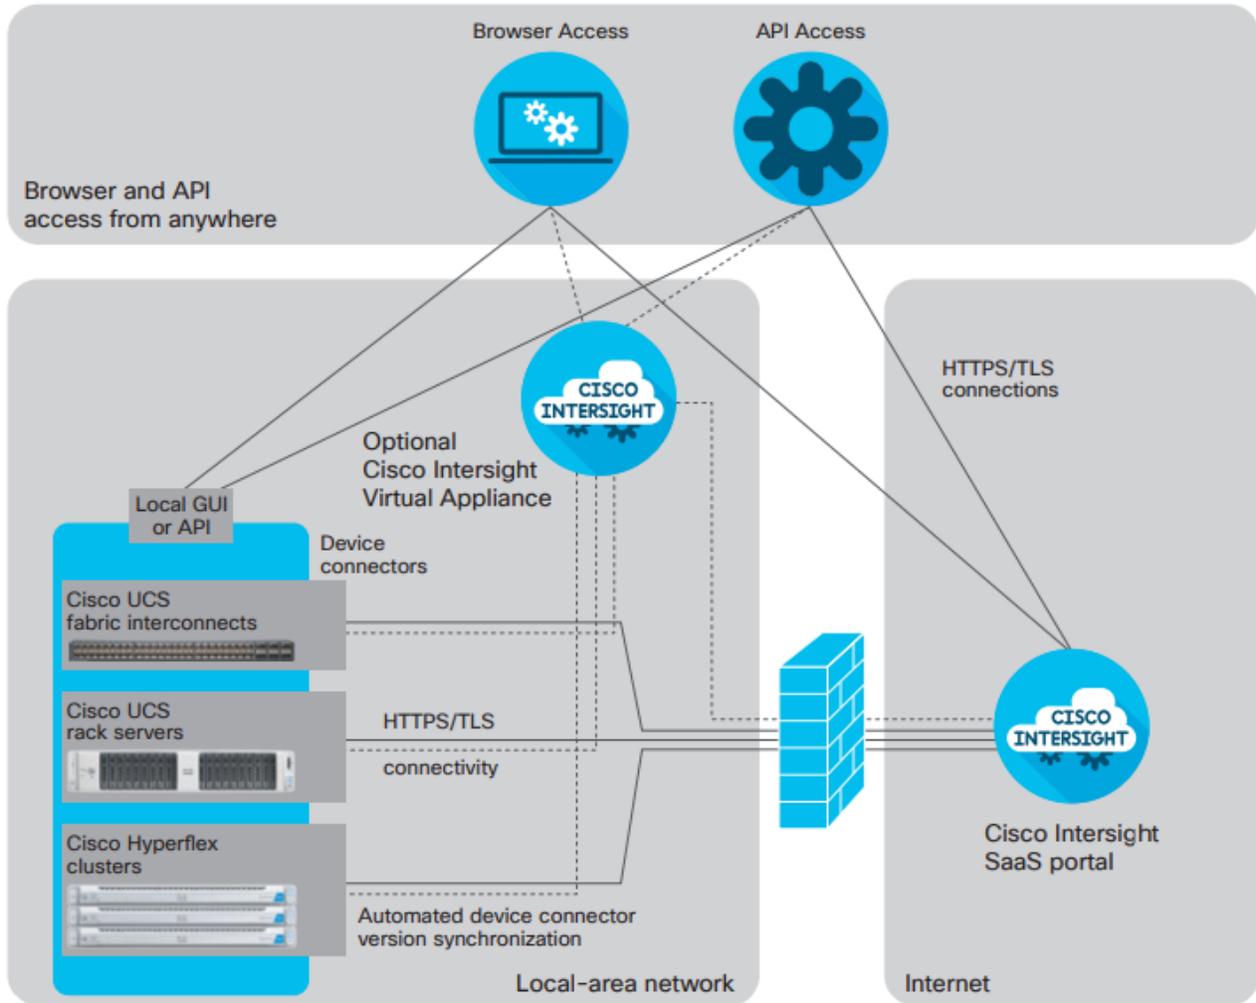
Intersight accounts form the authentication domain for users. The accounts control all resource access, and authenticated users are restricted from seeing any data in accounts where they are not authorized. With the SaaS platform, Cisco login IDs can be used for authentication with the identity provider for Cisco.com, which includes support for multifactor authentication. Both SaaS and on-premises Intersight implementations allow integration with external identity management systems to meet existing customer authentication requirements.



The Cisco Intersight framework uses granular access control with privileges managed per resource. Intersight software allows configuration of users and groups into several roles, and each user or group can be a member of multiple roles. Roles implemented include the following privileges:

- Account administrator: Full control and management capabilities for the Cisco Intersight account and devices under management
- Read-only: Read-only visibility to resources under management
- Device technician: Administrative device actions including device claim to a Cisco Intersight account
- Device administrator: Administrative device actions including device delete from a Cisco Intersight account
- Server administrator: Server lifecycle and policy-based management
- User access administrator: User, group, and identity provider configuration Please see the Intersight help pages for specifics on managing roles and resources.

Cisco UCS and HyperFlex systems are connected to the Intersight SaaS platform or on-premises virtual appliance through a device connector (DC) that is embedded in the management controller of each system.



All data exchanged between devices and the Intersight platform uses industry-standard encryption and security protocols. Connected devices use Transport Layer Security (TLS) with restricted ciphers and HTTPS on the standard HTTPS port 443. All data sent to Intersight is encrypted using the Advanced Encryption Standard (AES) with a 256-bit, randomly generated key that is distributed with a public-key mechanism. In addition, every device connection to the portal is authenticated with a cryptographic token so that only legitimate devices can be managed.

All connections are initiated from the device. Thus, firewalls can block all incoming connection requests; only HTTPS port 443 needs to be enabled for outbound connections. As a result, firewalls do not need any other special configuration to enable Intersight connectivity. Devices can be configured to use HTTPS proxy servers to add an additional layer of security through indirection.

To help ensure connection security and prevent man-in-the-middle attacks, Cisco UCS and Cisco HyperFlex devices connecting directly to the Intersight platform use a single-destination HTTPS URL. The platform presents a certificate signed by a certificate authority (CA). If an unsigned certificate is presented, the devices will not connect to the portal. Intersight software and the device connector create a secure management framework



that provides real-time information related to device security. This approach also allows connected devices and Intersight software to stay synchronized with the latest connection security updates.

To monitor and manage devices with the Intersight platform, they first must be claimed from an Intersight account. Devices can be claimed using a browser by going to the SaaS or virtual appliance portal and clicking on the Claim Devices tab. Device IDs and a claim code, both of which are unique to the device, are retrieved from the device. You can find the device ID and claim code through the device's local management interface. The claim code is refreshed every 10 minutes as an additional safeguard to ensure that the administrator claiming the device has physical access to it.

Two-factor authentication is used to verify the identity and authenticity of each device being claimed. This authentication mechanism adds another layer of security to the device-claiming process. It requires access to the device as well as device identification information that is validated against your Intersight account. If an unauthorized user guesses or learns device information, the user cannot claim a device without access to the device.

The device-claiming process allows the user to set the device as read-only or allowing control from the Intersight platform. Devices configured as read-only cannot be modified by Intersight software regardless of user privileges within the Intersight account. Devices also can be unclaimed or removed from a Cisco Intersight account through the portal.

Compliance with industry security standards

The Intersight platform meets or exceeds InfoSec's requirements applying to numerous industry standards

- Federal Information Processing Standard (FIPS) 140-2: Intersight uses FIPS 140-2 compliant cryptographic modules. Certifications are being planned.
- The platform's out-of-band management architecture makes it out of scope for some standards/audits:
 - Payment Card Industry Data Security Standard (PCI DSS): Customer traffic (including cardholder data) does not flow through the Intersight platform.
 - The Health Insurance Portability and Accountability Act (HIPAA): No individually identifiable health information (IIHI) on the network is ever sent to the Intersight portal.

Data collected and encrypted at rest

- The Intersight platform has complete visibility into and control over managed systems, the same as local API access. Data collected from device connectors on managed systems may include the following:
 - Inventory and configuration data for fabric interconnects and all servers and nodes, including storage controllers, network adapters, I/O modules, and CPUs.
 - Server operational data (such as faults) that can be used by the Intersight platform to provide automated recommendations.
 - Technical support files that can be created when requested by the Cisco Technical Assistance Center (Cisco TAC).

Note that device connectors do not collect sensitive data that may be stored in the connected systems, such as passwords. If you use the Cisco Intersight Virtual Appliance, you have control over whether the above data



is passed on to the cloud-based portal. If you opt out of additional data collection, the above information is kept locally. The Intersight help pages have more information on data collected by the onpremises Cisco Intersight Virtual Appliance.

For all data collected, the following additional security practices are implemented:

- Customer data is kept separate from other customer data through virtual data segregation. Data requests by Cisco Intersight services return data specific to the customer account only, and per-customer encryption keys are used for access.
- Long-term persistent data is encrypted at rest. Block storage or similar volume encryption is enabled for all data and tenant files.
- Third-party access to data is not permitted.

Intersight RBAC specific to HX

Beginning with HXDP 5.5(1a), Role Based Access Controls (RBAC) specific to HX have been added to management with Intersight. These roles are Cluster System Operator, Cluster Data Protection Operator, Cluster Lifecycle Administrator, and Cluster Administrator. See the Intersight administration documentation on setting up users.

https://intersight.com/help/saas/resources/Managing_Roles_and_Privileges

The characteristics of these user roles are listed below.

HX Cluster System Operator

- Can view HX on the dashboard, cluster inventory, cluster details, license status
- Can cross-launch HX Connect
- Can run health check
- Can enable software encryption

HX Cluster Data Protection Operator

- Can view HX on the dashboard, cluster inventory, cluster details, license status
- Can configure N:1 backup and restore

HX Cluster Storage Administrator

- Can view HX on the dashboard, cluster inventory, cluster details, license status
- Can create, update, or delete datastores
- Can configure and manage iSCSI storage

HX Cluster Lifecycle Administrator

- Can view HX on the dashboard, cluster inventory, cluster details, license status
- Can cross-launch HX Connect
- Can run health check
- Can create HX policies, create and deploy cluster profiles, expand and upgrade clusters

HX Cluster Administrator (system-defined role)

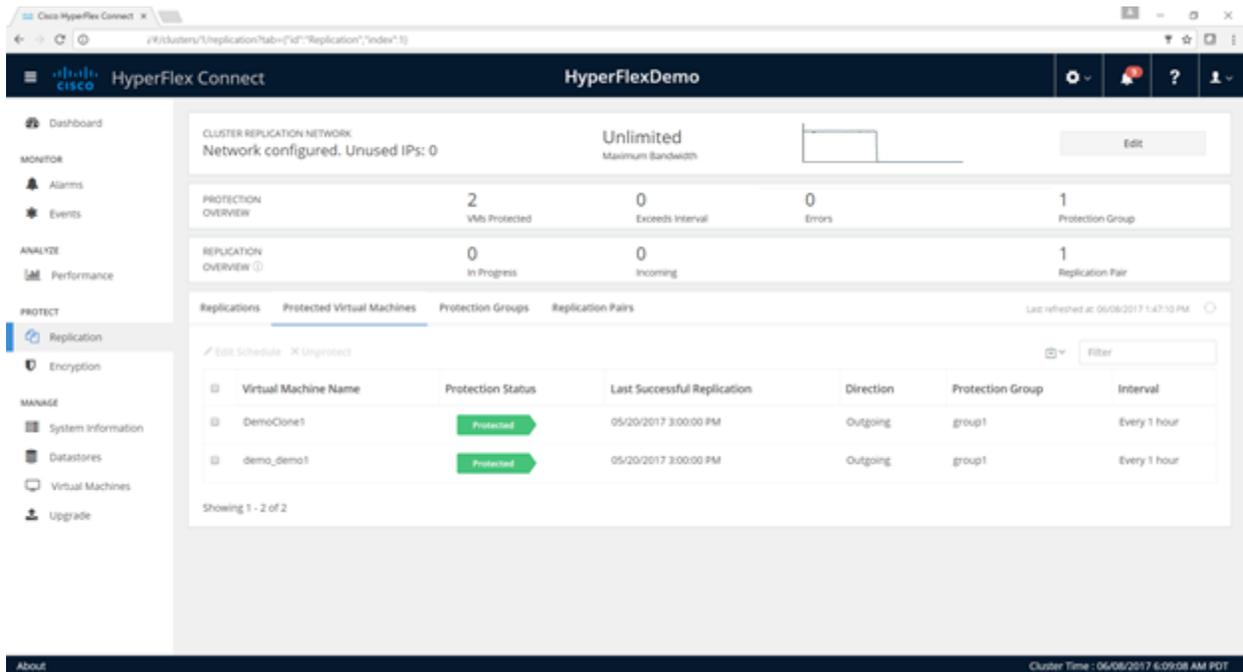
Cisco HX Platform Hardening Guide



- Union of all the above privileges (i.e., super user)

HX Connect

HX Connect is a native UI for managing the HX cluster. This includes configuring replication and encryption along with some VM management functions.



HX Connect has a security warning banner that can be disabled on a global basis by the administrator(s).

The session to the interface is encrypted via FIPS compliant SSL communications. The mechanics of the session are described below, and caution should be taken by the administrator(s) when logging out of sessions to ensure that all tokens are revoked and sessions are terminated.

Session architecture:

1. When a user logs in server provides an access token to the user. This access token is used to validate this user and do all subsequent actions.
2. Idle Timeout is by default set to 30 minutes. You can change/view this idle timeout in UI in user settings (Click on the top right user icon and you will see the user setting.)
3. Idle timeout is global and can be changed or viewed by a user with admin role.
4. If an admin user changes the idle timeout, it will be reflected for all the users.
5. From HX Connect perspective, if a user is not doing any activity on the GUI and is idle then after 30 minutes (default idle timeout), the user is logged out.
6. Once the user logs out, the access token is revoked.

Details of the transaction:



Session Management happens at the HX Connect browser end, and token management happens in AAA [backend]. Once a user logs in, a session starts. The “start of session” implies that HX Connect creates a cookie and installs it in the browser.

This cookie is removed under the following circumstances:

1. When user logs out explicitly.
2. When idle timeout occurs.
3. When user closes the browser completely.

If you log in using HX Connect (session starts):

1. You share the same session if you open another tab in browser window.
2. You share the same session if you open another window from the browser, you logged in.

In addition, this means that if you login using HX Connect and open another window or tab and navigate to the CIP-M, and then logout, you log out from **all** tabs and windows.

Please note that the cookie is not removed when:

1. User closes a tab.
2. User closes a window in the browser [however, the browser process is still running, i.e., another widow of this browser is still alive].

This means that the login session is still active in the above two cases.

Associated with the session is a token. This is managed by AAA. This token will be invalidated when the user logs out.

If you close the browser completely without logging out, you will no longer have a session, but the token will be alive. Therefore, it is recommended that you logout before closing the browser.

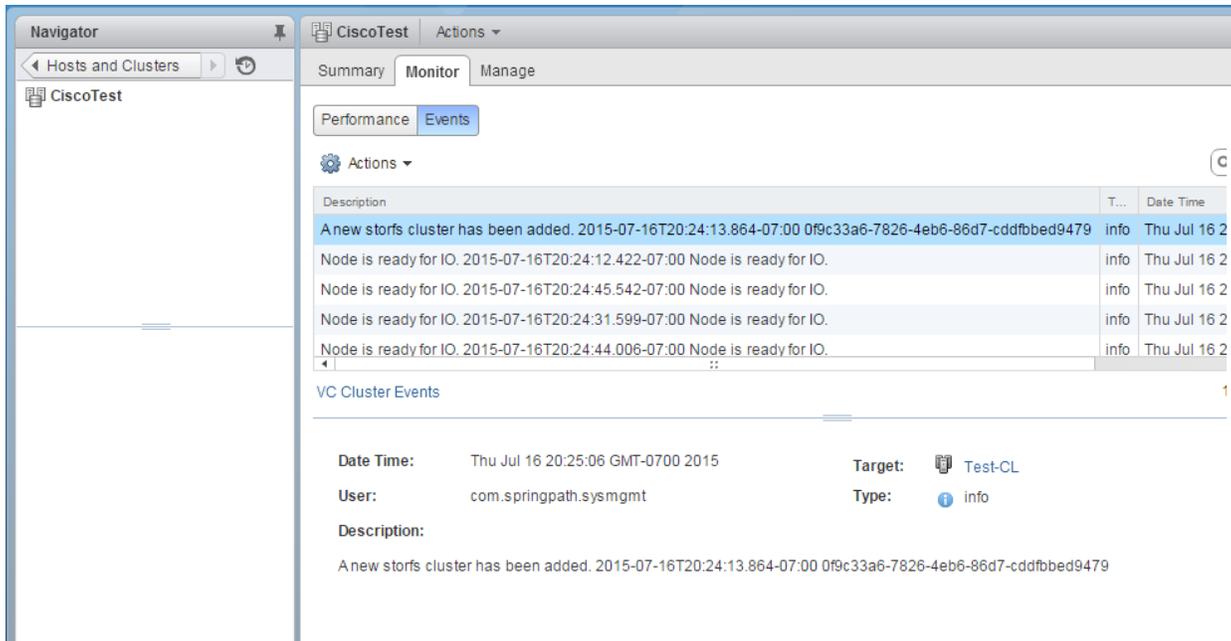
Multiple Sessions for same user are supported if user logs in:

1. From different machines.
2. From different kinds of browser [such as, Chrome and Safari] on the same machine.

HX Connect also provides a support bundle collection interface that allows the user to collect and download all system component logs, including audit files. These can then be examined or uploaded to support.

vCenter Plug-in

The vCenter plug-in is an https accessible UI available after logging in to vCenter. The portlet to access the plug-in is in the summary page for the cluster or accessible in the VC inventory list. The besides providing an admin interface for datastore creation and cluster consumption overviews, the plug-in has a monitor tab that permits event and task browsing along with hardware status.



The plug-in's right click context menu also allows the administrator to create VAAI offloaded snapshots of VMs, perform cloning operations, and generate system wide support bundles for log collection. These can then be examined or uploaded to support.

Plug-in session mechanics operate in the same manner as vCenter sessions and are managed by editing the appropriate vCenter configuration files.

STCLI and HXCLI

A session via FIPS compliant SSH cipher suites is used to access the CLI (STCLI and HXCLI). STCLI is being deprecated for general use and users are encouraged to use HXCLI for their command line administrative needs. All administrative functions along with some extra options are available via the CLI. See the HX STCLI reference for an exhaustive list.

<http://www.cisco.com/c/en/us/support/hyperconverged-systems/hyperflex-hx-data-platform-software/tsd-products-support-series-home.html>

A warning banner can be configured on the control VM (HXDP) for display on access using the MOTD functionality available in the base OS. This can also be done for ESX.

- At the control VM CLI add a file called `/etc/update-motd.d/00-springpath-motd`
- At the ESX CLI use the web or C# client to set `config.etc.issue` for the DCUI and `config.etc.motd` for SSH both under advanced options. Alternatively, use `/etc/issue` for DCUI and `/etc/motd` for SSH.
- There is no customization on the vSphere Web client logging into vCenter.

The STCLI and HXCLI security subset of commands enables the administrator to configure external machines to access the datastore, configure the root password, synchronize SSH keys across the nodes or enable, set, and



disable encryption. Access should not be granted unless the external system is trusted. Access should be revoked when move/copy/migration operations are complete.



STCLI/HXCLI security

usage: stcli/hxcli security [-h] {password,whitelist,ssh,encryption} ...

NOTE: There are significant changes to the CLI interface Starting with HX 4.5.1 (a), described below.

Secure Admin Shell Access (HXDP 4.5.1 (a) and above)

With the release of HXDP 4.5.1 (a) a significant change in security posture has been introduced with respect to the CLI and root users. Users are no longer able to access the root account on a controller VM via the normal routes. Users are also placed into a restricted access shell upon SSH to the system as admin. This restricted admin shell has the following characteristics:

- Removes **root** access over ssh to the command line of Controller VMs via management interfaces
 - Command line access to Controller VM must authenticate as the **admin** user
- Reduces attack surface by restricting **admin** user to execute only allowed commands that cannot manipulate the system
- Restricts changes to the Controller VMs to HXDP upgrades/updates only
 - A full list of commands in the restricted admin shell can be listed by typing "?" or "help" and pressing enter from the admin shell command line.
- Base commands are limited and do not allow executables, downloads, or modification to system files
- Some scripts are allowlisted and available via "priv"
 - E.g., hx_post_install
 - A full list of commands in priv can be listed by typing "priv" and pressing enter from the admin shell.
- root access is available only for troubleshooting by su to root from within secure shell **after customer-initiated Consent Token challenge-response with TAC**

Access the root account is only available once a specific challenge-response workflow has been completed. Only your TAC representative can generate the proper response token. Only this token can grant access and only within the timeframe specified by the user during the initiation of the consent token workflow.

Before contacting TAC be sure you understand the consent token workflow:

- The consent token challenge-response takes place from CLI on the CIP-M node by typing "su" once you are logged in via SSH as admin.

```
kaptain@Kaptain-Linux:~$ ssh admin@10.10.10.10
The authenticity of host '10.10.10.10 (10.10.10.10)' can't be established.
ECDSA key fingerprint is SHA256:0kA9czzcL7I5fybflNtSI+D+Ng5dYp15qk/9C1cQzzk.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.10.10.10' (ECDSA) to the list of known hosts.
HyperFlex StorageController 4.5(1a)
admin@10.10.10.10's password:
This is a Restricted shell.
Type '?' or 'help' to get the list of allowed commands.
admin:~$ su
Password:
WARNING: By accepting this support session, you give your consent and hereby authorize Cisco to have privileged access to
the supported Cisco device for the purpose of providing technical support. At the conclusion of this session you must ex
it root shell from all the open ssh sessions of all the controller vms of the cluster and invalidate the consent token in
order to terminate Cisco's access and close the privileged access portal. You are hereby advised that failure to do so m
ay create a vulnerability in your product.
Accept(Y/n): Y
Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
Enter CLI Option:
█
```

- **Be sure to have access to TAC**
- The successful challenge/response performs a background sync of the token on all nodes.

```
kaptain@Kaptain-Linux:~$ ssh admin@10.10.10.10
The authenticity of host '10.10.10.10 (10.10.10.10)' can't be established.
ECDSA key fingerprint is SHA256:0kA9czzcL7I5fybflNtSI+D+Ng5dYp15qk/9C1cQzzk.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.10.10.10' (ECDSA) to the list of known hosts.
HyperFlex StorageController 4.5(1a)
admin@10.10.10.10's password:
This is a Restricted shell.
Type '?' or 'help' to get the list of allowed commands.
admin:~$ su
Password:
WARNING: By accepting this support session, you give your consent and hereby authorize Cisco to have privileged access to
the supported Cisco device for the purpose of providing technical support. At the conclusion of this session you must ex
it root shell from all the open ssh sessions of all the controller vms of the cluster and invalidate the consent token in
order to terminate Cisco's access and close the privileged access portal. You are hereby advised that failure to do so m
ay create a vulnerability in your product.
Accept(Y/n): Y
Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
Enter CLI Option:
█
```

- **Verify with HX Connect**
 - **system information --> node tab --> columns showing admin/root status.**
- If a node cannot sync, that node will not have token.
- If a node does not have the token, commands will not take effect on that node from the CIP-M root shell
- If a node does not have consent token, you need to reproduce consent token workflow on that node

```

Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
Enter CLI Option:
1
Enter time period in minutes for root shell access(max 4320 mins): 10
Generating Challenge.....
Challenge String (Please copy everything between the asterisk lines exclusively):
*****BEGIN TOKEN*****
rBykQgAAAEBAQAABAgAEAAAAAMACJyfnUou7I6HBAAQtcJXdmscNjht+woFyq76LwUABAAAAoGAALiEXBlcmZsZXgHAAXIeXBlcmZsZXhfQ1QIAAtIW
C1NNVMtSFhEUAKAIDA5NjA2ZjVhMzEwZmE2YWRhYTgyOWJjZGRhODUzOTBl
*****END TOKEN*****
Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
Enter CLI Option:

```

You will have root access on the system for the duration of time entered in time period question that was answered in the workflow. In the screen above, that time is 10 minutes. During this time interval you can log out and log back in to the CLI as admin and su to root without challenge, using the hard root password set at installation. Once this time interval is expired, you will have to re-run the challenge/response with a new set of tokens to regain access.

The response token is not a random string. The token hash is encoded with the following information:

- Function: 0x00000001 ROOT_SHELL_ACCESS
- Sub Function: 0x00000000 DEFAULT
- TTL (minutes)
- Product Name: Hyperflex
- Key Name: Hyperflex_CT
- PID: HYPERFLEX
- Serial Number

REST APIs

Cisco HXDP comes with a comprehensive REST based API for use in developing custom software that can access the system. The built-in REST API:

- Contains well-documented syntax and examples with REST API explorer
- Secure token-based access with RBAC and auditing
- Accessed Via: <http://<Cluster-IP>/apiexplorer>

GET /rebalance/status Show rebalance status

Response Class (Status 200)
successful operation

Model | Example Value

```

{
  "rebalanceEnabled": false,
  "rebalanceStatus": {
    "rebalanceState": "CLUSTER_REBALANCE_NOT_RUNNING",
    "percentComplete": 0,
    "setRebalanceState": false,
    "setPercentComplete": false
  },
  "setRebalanceEnabled": false,
  "setRebalanceStatus": false
}

```

Response Content Type

Parameters

Parameter	Value	Description	Parameter Type	Data Type
Authorization	Bearer eyJhbGciOiJIUzI1NiU9.eyJzdWwiOiJ1c2Vycy9		header	string

[Hide Response](#)

REST APIs can be used to authenticate users and grant or validate access tokens. AAA provides role-based access control to REST APIs which allows users to perform various operations on resources in a cluster. These APIs are supported in version v1 of the AAA REST API and are subject to change and deprecation in future versions.

A rate limit is enforced on the /auth API in a 15-minute window. This means that /auth can be invoked (successfully) a maximum of 5 times. A user is allowed to create a maximum of 8 unrevoked tokens. Subsequent call to /auth will automatically revoke the oldest issued token to make room for the new token. A maximum of 16 unrevoked tokens can be present in system. In order to prevent brute-force attacks, after 10 consecutive failed authentication attempts, a user account is locked for a period of 120 seconds. Access Tokens issued are valid for 18 days (1555200 second).

Starting with HXDP 4.0.1a, a subset of the REST API entries are reserved for STIG specific functions. The following is a list of the current set:

- configure_stig_parameters
- configure_stig_parameters_host
- configure_stig_parameters_vm
- configure_stig_parameters_vCenter
- remove_stig_parameters
- check_stig_parameters

Examine the API explorer discussed above for more detailed explanations of the values passed and returned by these STIG API calls.



It is not possible to remove API methods like PUT and DELETE. These methods sometimes show up on OWASP or nmap type scans. Note that the APIs are all behind an authentication barrier. This is sufficient to prevent exploits of these methods by bad actors. More details around API authentication can be found here:

<https://developer.cisco.com/docs/ucs-dev-center-hyperflex/#!authentication/rest-api-explorer--authentication>

AAA Domains

Authentication, authorization, and accounting (AAA) is managed by HX depending on the access method. HX Data Platform supports Role-Based Access Control (RBAC). AAA is implemented with Open Authorization (OAuth), Security Assertion Markup Language (SAML), or Lightweight Directory Access Protocol (LDAP). It is integrated with the ESX cluster authentication mechanism. HX Connect and the STCLI primarily use this database for user authentication. Access to HX Connect or the STCLI is also available using a local admin account in the event that vCenter is unavailable. Beginning in HX 3.5(1a) the local root user is no longer available for HX Connect logins.

vCenter

vCenter maintains a set of user accounts and roles in a database. vCenter itself can be integrated with an external AD or LDAP user management system. HX RBAC integrates directly to this mechanism. See the HX RBAC documentation for configuration steps.

AD Integration

You can join a Platform Services Controller appliance or a vCenter Server Appliance with an embedded Platform Services Controller to an Active Directory domain and attach the users and groups from this Active Directory domain to your vCenter Single Sign-On domain.

<https://docs.vmware.com/en/VMware-vSphere/6.0/com.vmware.vsphere.vcsa.doc/GUID-08EA2F92-78A7-4EFF-880E-2B63ACC962F3.html>

User Management

RBAC settings configure users with one or more roles. Roles are assigned privileges to act on a resource. For example, one role has a privilege to perform virtual machine power on, another role has a privilege to only monitor a virtual machine. Users are created through vCenter. vCenter supports Active Directory (AD) users and groups. Two roles are supported with HX. Privileges associated with these roles cannot be modified.

- Administrator:
 - Most tasks that can be performed on a HX Storage Cluster require administrator privileges.
 - Administrative users grant privileges to the roles.
 - Administrator users have access to the HX Data Platform interfaces: HX Connect, HX Data Platform Plug-in, the Storage Controller VM command line for running STCLI commands, and HX Data Platform REST APIs.
- Read-only:
 - This role allows users to monitor status and summary information through HX Connect and the HX Data Platform Plug-in.
 - Read Only users have access to the HX Data Platform interfaces: HX Connect and the HX Data Platform Plug-in.



Cisco HyperFlex User Overview

The user types allowed to perform actions on or view content in the HX Data Platform, include:

- **admin**—A predefined user included with HX Data Platform. The password is set during HX Cluster creation. Same password is applied to root. This user has read and modify permissions.
- **root**—A predefined user included with HX Data Platform. The password is set during HX Cluster creation. Same password is applied to admin. This user has read and modify permissions.
- **administrator**—A created HX Data Platform user. This user is created through vCenter and assigned the RBAC role, administrator. This user has read and modify permissions. The password is set during user creation.
- **read-only**—A created HX Data Platform user. This user is created through vCenter and assigned the RBAC role, read-only. This user only has read permissions. The password is set during user creation.

You can use REST APIs with the read only user. If you have a read-only user, (s)he can perform only GET operations. They will receive an access error if they perform PUT, POST, or DELETE operations. The non-read only users are called admin users. They are CVM users and users belonging to administrator group in vCenter.

Local Users

The main cluster local user is "admin". The cluster maintains a separate administrative account called root that is created at install time. This root user has full privileges to the system and hard passwords are enforced during creation. In 4.5.x and above, the root user is not accessible by end-users without the use of a challenge-response workflow managed by Cisco TAC.

Creating other users is not supported on the system. Only the admin user can SSH to the system. Vulnerability scanning is not a sanctioned use of the root shell. Creation of new user accounts for the sake of remote access for scans are not supported. Root is a protected account and should never be used. Use of root starting in 4.5.1a is strictly forbidden except if initiated by a TAC support case using a tokenized workflow.

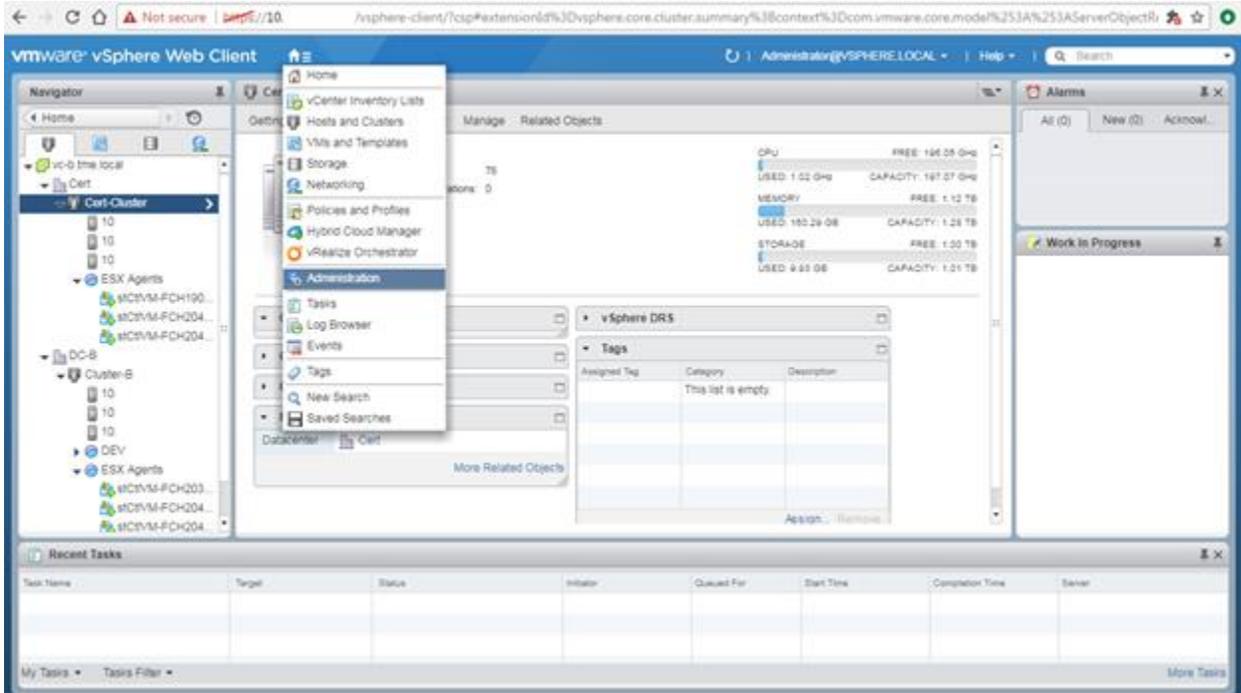
The "diag" user is a user account implemented in 5.0.2a and above. The diag user has a specific subset of permissions to certain files and scripts that can be run under guidance from TAC. This user and any diagnostics performed should only be used when prompted by technical support.

The root user is only available after contacting Cisco TAC. This is essentially a debug user and is generally not available to users. See the Secure Admin Shell Access section above.

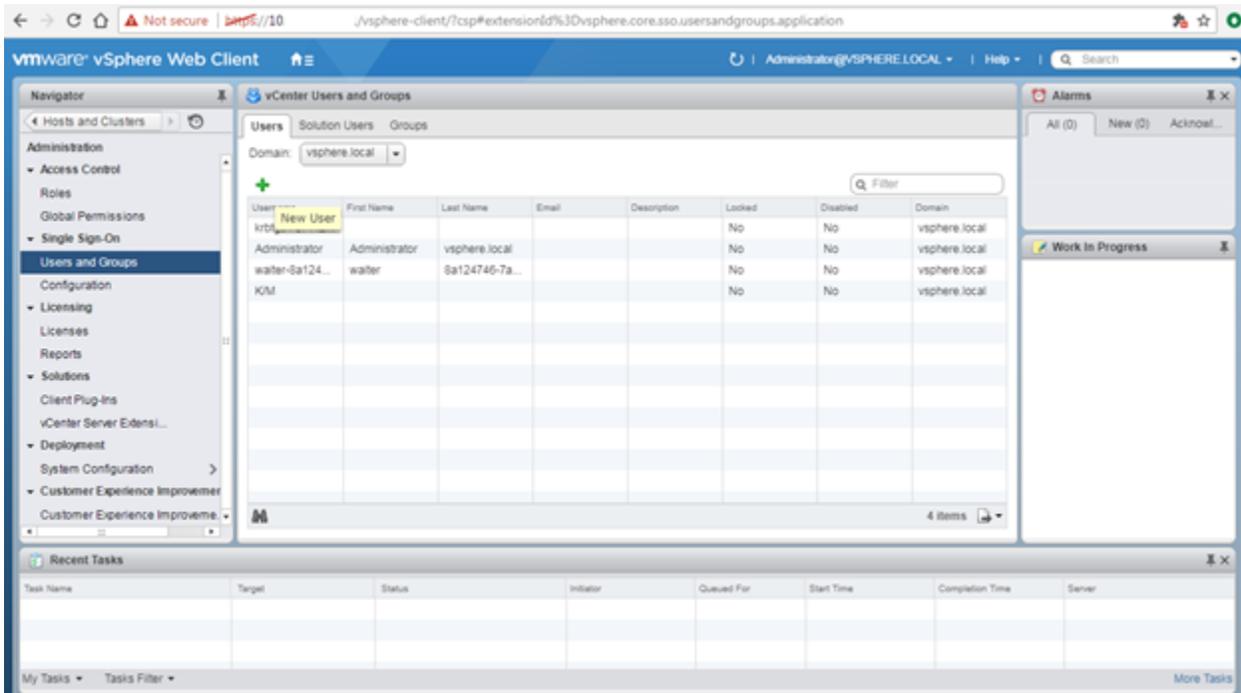
UI Users

Create new users for HX using vCenter with roles. This applies to the HX vCenter Plug-in and HX Connect UIs.

1. Log into the GUI plug-in for the cluster and select **Administration** under the **Home** icon.



2. Under **Single Sign-On**, add the user.





CLI Users

The “admin” user is the only supported CLI user under normal operations.

To set or change the password for the local node admin user:

- Change the admin password in pre-4.5.1a releases using `passwd admin` from root
- Change the admin password in 4.5.1a or greater using `hxcli security password set`
- Change on all nodes for consistency

To set or change the vCenter maintained administrative user password for UI users:

- log in to vCenter
- select Administration → Users and Groups
- Edit the password for the user

Auditing, Logging, Support Bundles

An audit trail, maintained in a set of audit logs, is a security-relevant chronological set of records that provide documentary evidence of the sequence of activities that have affected the system. They contain records of system changes at any time a specific operation, procedure, or event occurs. A full set of logs for the entire system can be gathered with a support bundle. However, STCLI and REST commands are recorded continuously and can be examined by looking at just a few files instead of the generating a comprehensive log dump. STCLI commands use the REST architecture to execute their commands, so they are also captured in the REST log. These audit records are maintained on each node of the system and are contained primarily in the following files on each node in the `/var/log/springpath` directory:

- `stMgr.log`
- `audit-rest.log`

Additional information relevant to an audit may be found here as well:

- `admin.log`
- `hxcli.log`

Auditing is required for compliance purposes and for forensic examination of system activity. A typical `audit-rest.log` entry will look like this:

```
2017-06-29-23:26:38.096 - Audit - 127.0.0.1 -> 127.0.0.1 – create /aaa/v1/auth?grant_type=password; 201; null 3341ms
```

- Timestamp - source IP – dest IP – http api method – URL retrieved – response code (200 success, 4xx error) – user issued – response time

What sources are captured:

- GUI -- REST API auditing – Any calls to REST
 - A method to audit UI usage as well as 3rd party integrated software
 - `/var/log/springpath/Audit-rest.log`
- STCLI (RBAC) auditing



- STCLI calls utilize the API
- Audit trail records will have the keyword "Audit".
- Collect all such Audit trail records and save it

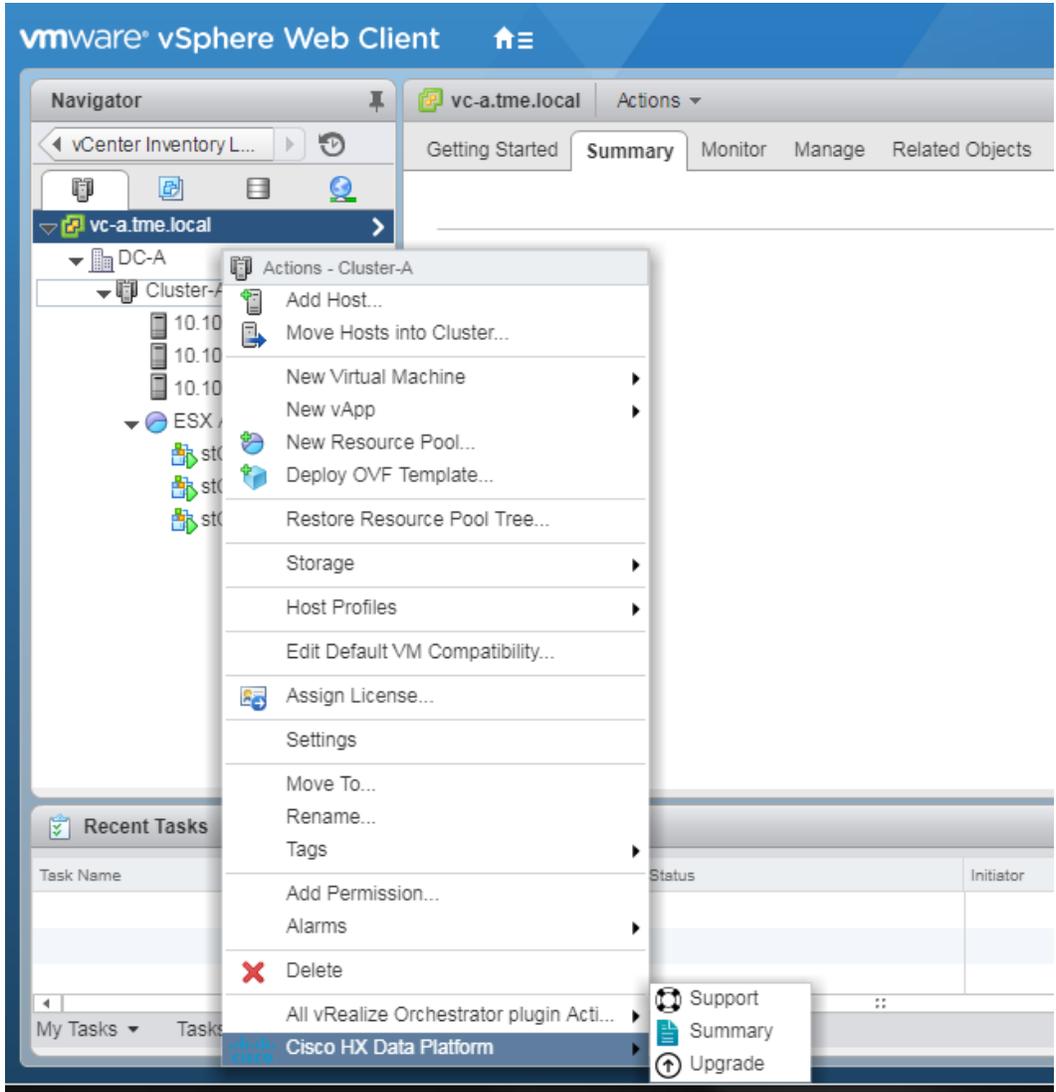
The cluster root user or a node root user can manipulate the audit logs. Read-only users or any other RBAC user account cannot alter the logs files

Replication log files that can be used for auditing traffic or general troubleshooting are listed below:

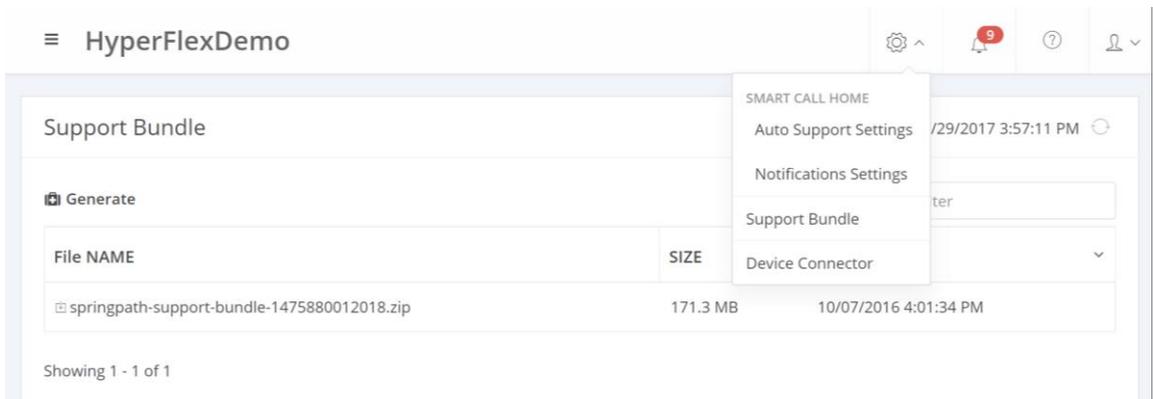
```
/var/log/springpath/nrcli.log  
/var/log/springpath/debug-repl-cipmonitor.log  
/var/log/springpath/nr-stat-history.log  
/var/log/springpath/user-replsvc.log  
/var/log/springpath/error-replsvc.log  
/var/log/springpath/replicationNetworkConfig.log
```

Support Bundles for the HX system can be generated in two ways. There are menu interfaces to generate them in both HX Connect and the vCenter Plug-in UI.

1. Generating the Support Bundle using vCenter:



2. Similarly, the Support Bundle can be generated in HX Connect:





If Auto Support has not been configured during install, be sure to configure it now. This can be done via HX Connect (see the Support menu in the above illustration) or via STCLI/HXCLI using STCLI services ASUP. Auto Support enables:

- HTTP based auto-support data collection for proactive case creation
 - Continuous monitoring thru auto-support for 30+ events to detect early problems
 - Critical events integrated with Smart Call Home
 - Auto-generate SRs (tickets)
- Email notifications for critical events

The log bundle in vCenter includes the plug-in log file and is located in the regular vCenter log location. For Windows, the default is C:\ProgramData\VMware\vCenterServer\logs. For VCSA (vCenter Server Appliance) the default log location is /var/log/vmware.

The retention policy for the log files created by the various HX services are defined in their corresponding logback.xml files. For example, the stMgr log files retention policy is defined in /opt/springpath/storfs-mgmt/stMgr-1.0/conf/logback.xml. For log files created by scripts or other workflows, 'logrotate' is used for log rotation. It runs as a cron job every 10 minutes. The retention policy for these files is defined in /usr/share/springpath/storfs-misc/syslog.logrotate.

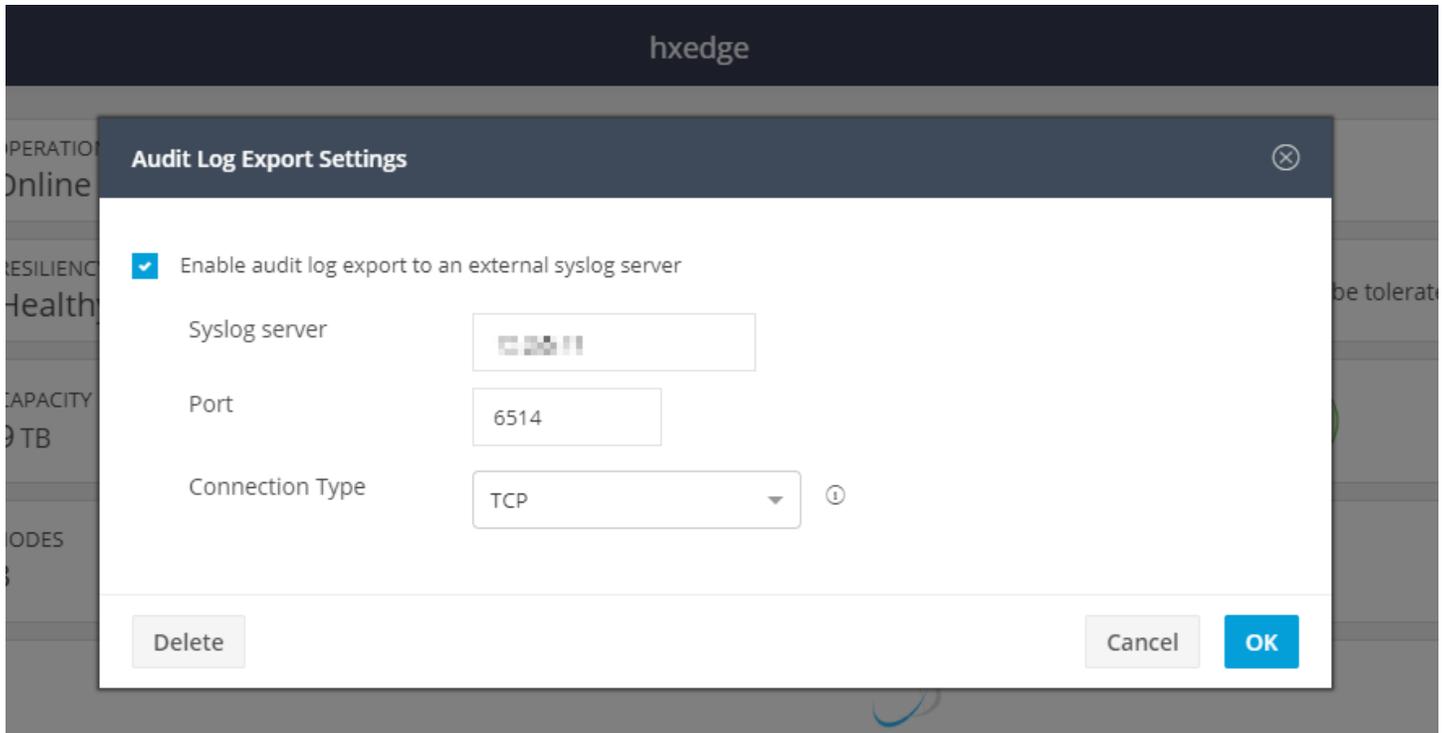
Here is a sample logback.xml file for the stMgr service:

```
<configuration scan="true" scanPeriod="60 minutes">
  <appender name="stMgrRoller" class="ch.qos.logback.core.rolling.RollingFileAppender">
    <file>/var/log/springpath/stMgr.log</file>
    <rollingPolicy class="ch.qos.logback.core.rolling.FixedWindowRollingPolicy">
      <fileNamePattern>/var/log/springpath/stMgr.%i.log.gz</fileNamePattern>
      <minIndex>1</minIndex>
      <maxIndex>10</maxIndex>
    </rollingPolicy>
    <triggeringPolicy class="ch.qos.logback.core.rolling.SizeBasedTriggeringPolicy">
      <maxFileSize>100MB</maxFileSize>
    </triggeringPolicy>
    <encoder>
      <pattern>%d{YYYY-MM-dd-HH:mm:ss.SSS,UTC} [%marker] [%X] [%thread] %-5level %logger{35}
- %msg%n</pattern>
    </encoder>
  </appender>
```

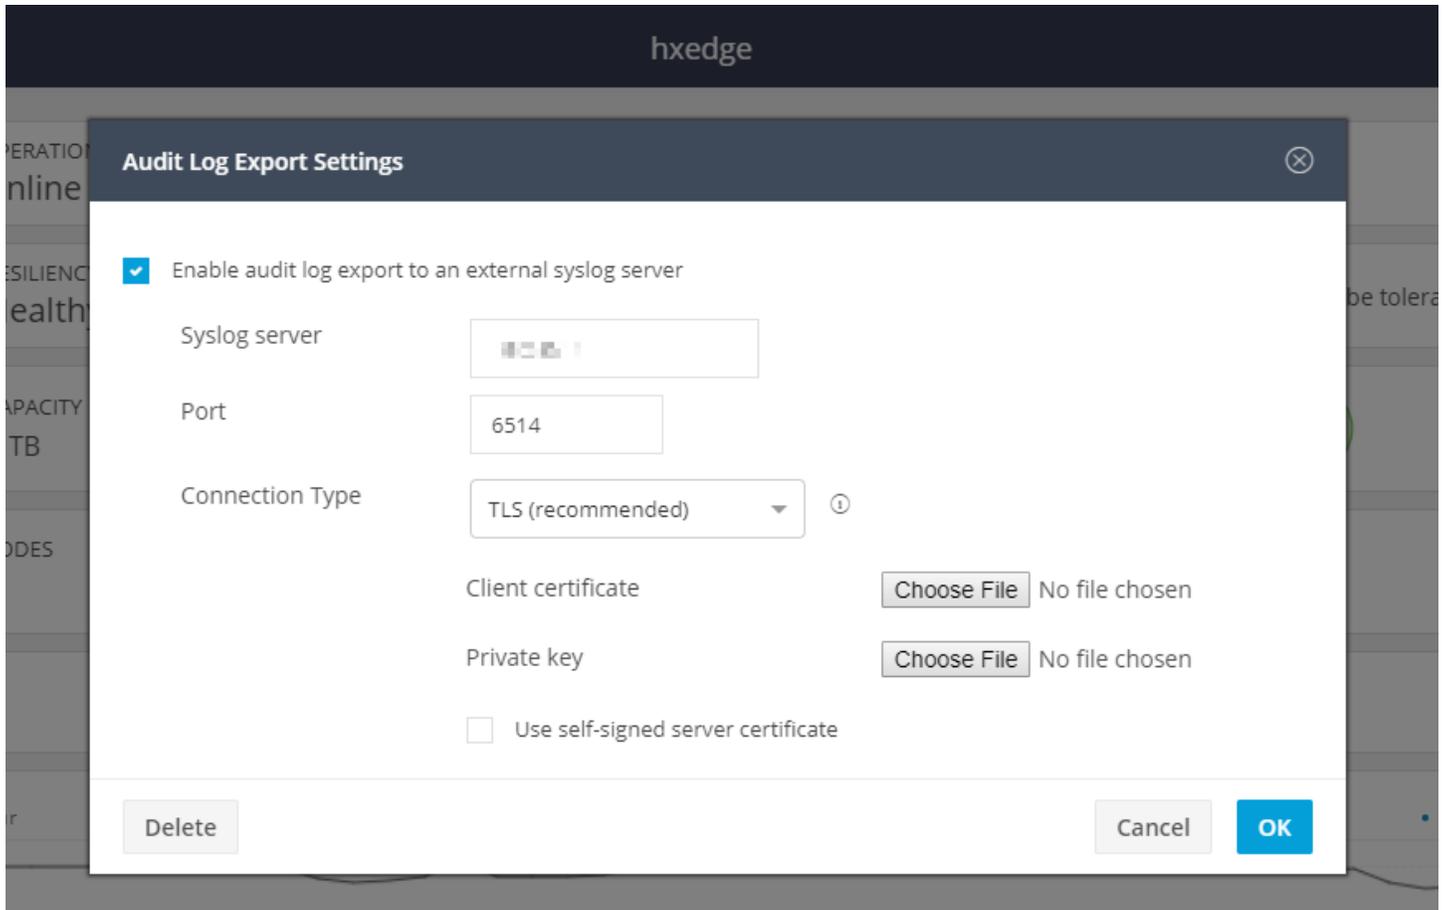
Setting Up Remote Logging for HXDP 4.0.1.a and Later

Beginning with HXDP 4.0.1a, HyperFlex includes a built-in syslog mechanism. Using the HX Connect Remote Logging wizard, you are able to enter the information required for each HX Cluster node to send its audit log records to a centralized remote server. You are required to have a remote log collection server built and accessible by the management interfaces on each cluster node. See Appendix E for a sample syslog-ng configuration file used for an Ubuntu-based log collector.

By clicking the gear symbol in the top right of the HX Connect UI, you can select remote logging. You will be presented with the following:



This is the default configuration, and it is for unencrypted transport over TCP using port 6514. The port is configurable; however, you must use either plain text or encrypted transport. If you select the drop-down you are given the option to change the connection type to TLS (encrypted).



The wizard then prompts you to upload a client certificate and key pair. This certificate and key will be used for each node to securely communicate with the remote log collection server. If the certificate is CA signed, it does not need to be uploaded.

If using self-signed certificates on the remote server, place them under `/etc/syslog-ng/CA`. Regardless of connection type selected, the system will attempt to connect to the remote server immediately. The server certificate will automatically be placed into the trusted certificate store on the syslog-ng client nodes.

The following syntax can be used with `openssl` to generate a self-signed certificate and key for both the client system and the remote server in the absence of a CA certificate:

```
openssl req -x509 -newkey rsa:4096 -keyout key.pem -out cert.pem -days 365
```

Password Requirements

Hard passwords are required for the cluster root user during installation. This password can be updated using the CLI from any node. SSH to the node and issue `stcli security password`.

Passwords for users maintained in the vCenter authentication database can have password difficulty set based on vCenter configuration. See your vCenter documentation for this.



Password Guidelines

The storage controller VM password for the predefined users admin and root are specified during HX Installer deployment. After installation, you can change passwords through the stcli or hxcli command line.

Component	Permission Level	Username	Password	Notes
HX Data Platform OVA	root	root	Cisco123	
HX Data Platform Installer VM	root	root	Cisco123	
HX Connect	administrator or read-only	User defined through vCenter.	User defined through vCenter.	
		Predefined admin or root users.	As specified during HX installation.	Requires leading local/ for login: local/admin or local/root
HX Storage Controller VM	admin	User defined during HX installation. User defined through vCenter.	As specified during HX installation. Strong password required.	Should match across all nodes in storage cluster. In 4.5.1a+ use the hxcli security password set command
HX Storage Controller VM	root	User defined during HX installation. User defined through vCenter.	As specified during HX installation. Strong password required.	Must match across all nodes in storage cluster. Use the stcli command when changing the password after installation.
vCenter	admin	administrator@vsphere.local default. SSO enabled.	SSO enabled. As configured.	Ensure the vCenter credentials meet the vSphere 5.5 requirements if the

Component	Permission Level	Username	Password	Notes
		As configured, MYDOMAIN\name or name@mydomain. com		ESX servers are at version 5.5. Read only users do not have access to HX Data Platform Plug-in.
ESX Server	root	SSO enabled. As configured.	SSO enabled. As configured.	Must match across all ESX servers in storage cluster.
Hypervisor	root	root	As specified during HX installation.	Use vCenter or esxcli command when changing the password after HX installation.
UCS Manager	admin	As configured.	As configured.	
FI	admin	As configured.	As configured.	

Session Timeouts

vCenter session timeouts are managed by vCenter configuration settings. Idle timeouts for TLS connections when using the HyperFlex plug-in are set in the following file making the noted changes:

- vSphere Web Client sessions terminate after 120 minutes by default. You can change this default in the `webclient.properties` file, as discussed in the *vCenter Server and Host Management* documentation.
- Login to the vCenter host system and navigate to this properties file. The location of this file depends on the base operating system on which vCenter is installed.
- Edit the file to include the line `session.timeout = value` where value is the timeout value in minutes. For example, to set the timeout value to 60 minutes, include the line `session.timeout = 60`.
- Restart the service

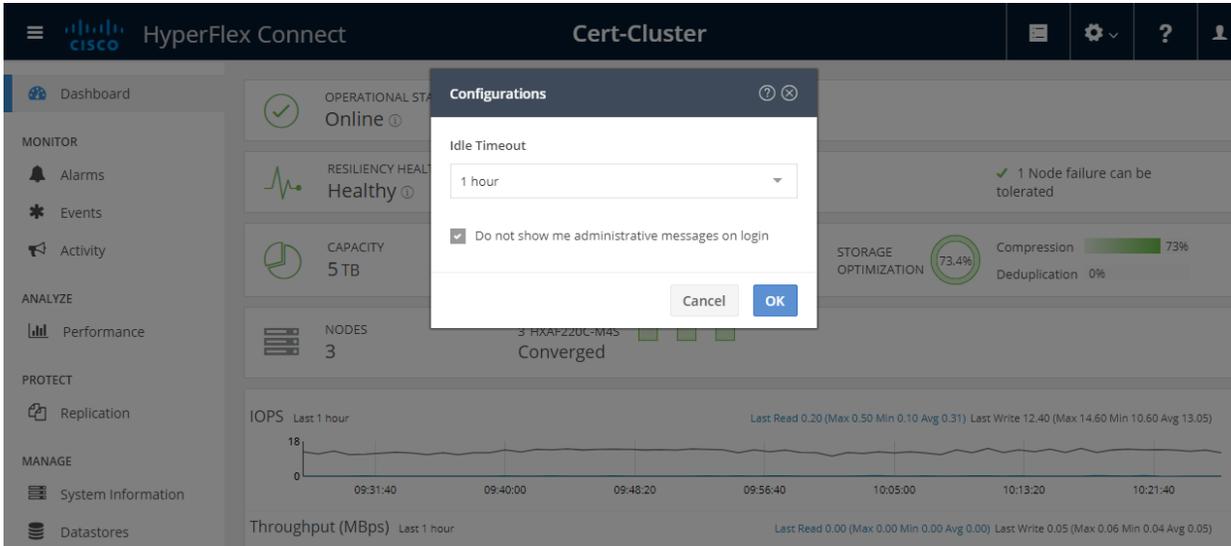
Alternatively:

- In the vSphere Web Client, navigate to the vCenter Server instance.
- Select the Manage tab.
- Under Settings, select General.
- Click Edit.
- Select Timeout settings.
- In Normal operations, type the timeout interval in seconds for normal operations.



HX Connect

Idle session timeouts for HX Connect sessions can be set in the dashboard view under the administrative icon.



CLI

Prior to 4.5.1a, the idle timeout for an STCLI session can be set on each HXDP CVM. SSH to the STCLI of each node and navigate to the `/etc/ssh/sshd_config`. Uncomment and change the `ClientAliveInterval` by setting a time. Once editing is complete, restart the `sshd` services. `ClientAliveInterval 60` would drop the connection after 60 seconds of inactivity.

For deployments of HX that are 4.5.1a or later, setting the SSH timeout for the admin shell requires editing the following file:

`/etc/lshell.conf`

This file is accessible via root shell only. The admin user will need to complete the token challenge-response workflow in order to edit the file. Once you can edit the file, simply change the last line here by specifying the number of seconds before session termination:

```
# 6 hours of inactivity will kill the admin session
timer      : 21600
```

Manually Clearing Sessions

To clear session data associated with a given user immediately you can run this command from any CVM while logged in as root:

```
> python /opt/springpath/clearsession.py root
```



Users can run into this situation if they close the browser session without logging out.

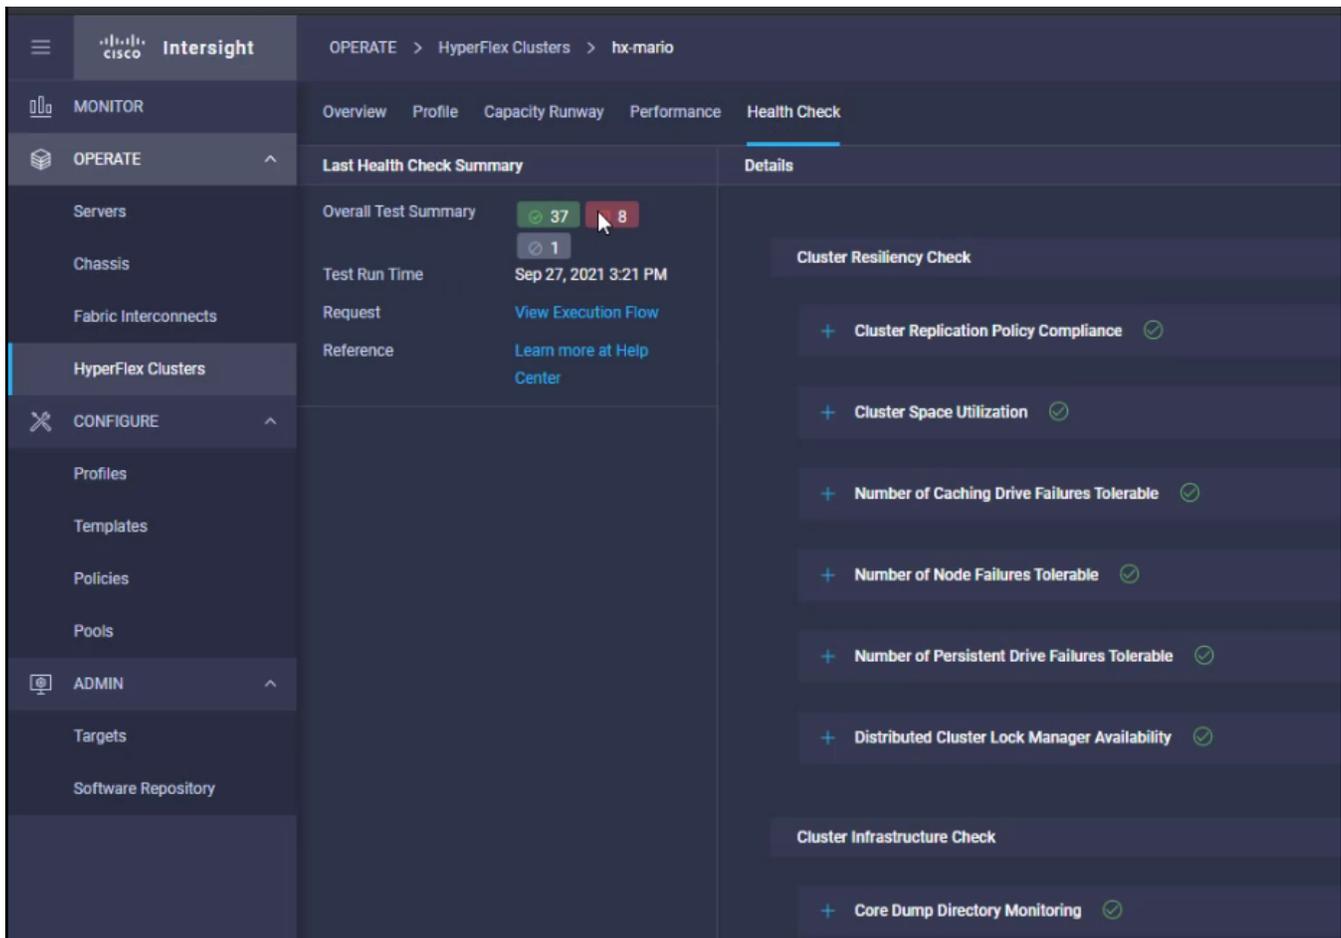
HX Platform Hardening

This section provides information on setting specific configurations for HX, ESX and UCS to further enhance system security.

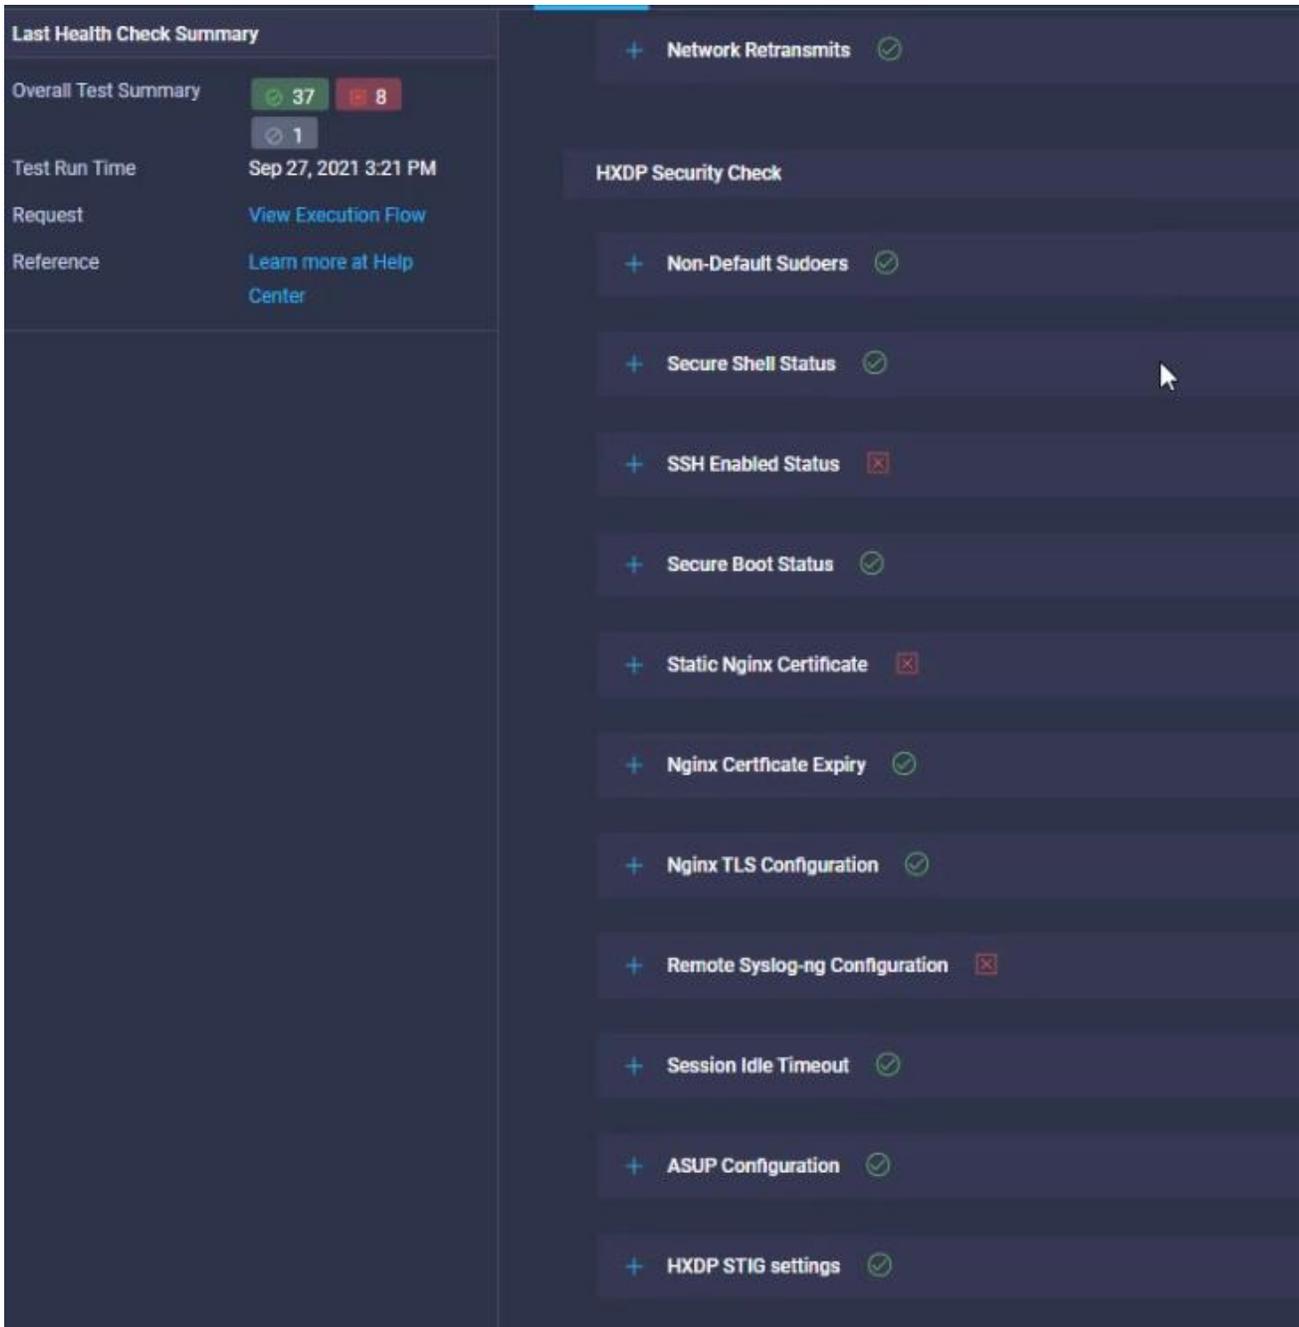
HyperFlex Security Advisor (SA)

Security Advisor (SA) is a subset of the Health Check function in Intersight. It is available in the HyperFlex ClustersàOperateàHealth Check tab. It is a best practice and recommended that general health checks along with security checks be performed on a regular basis and prior to upgrades.

Cisco Intersight has an integrated health check that is available for claimed systems.



The health check has a subsection for security related items:



These Security Advisor checks give a quick overview of the cluster with the top-level security posture items called out. For example, certificate expiry is displayed along with STIG setting compliance, among others. Note also that the security check is "dynamic", that is it is decoupled from cluster upgrades. The check is:

- Built into the Intersight framework
- Removed from HXDP updates
- Adaptable, with components added or removed as needed with product evolution

US Federal STIG (Secure Technical Implementation Guide) Settings



The controller VM is not a generic Ubuntu server and users should not generally apply Ubuntu specific STIGs to the CVM on their own. The controller VM is a virtual appliance and is highly customized for the purpose of running HXDP. The CVM has many built in security mitigations e.g., stripping out software packages that are not in use, iptables rules, secure admin shell, consent token root access, runtime defenses, secure boot etc.

Cisco implements relevant Secure Technical Implementation Guide (STIG) settings as defined by the Defense Information Security Agency (DISA) for several aspects of the HXDP ecosystem. The STIG adherence is accomplished through implementation of settings explicitly called out in the following DISA STIGS:

- U_General_Purpose_Operating_System_V1R4_SRG
- U_VMware_vSphere_6-0_ESXi_V1R4_STIG
- U_VMware_vSphere_6-0_vCenter_Server_for_Windows_V1R4_STIG
- U_VMware_vSphere_6-0_Virtual_Machine_V1R1_STIG

The corresponding STIG settings for vSphere 6.5 and 6.7 are supported starting in HX 5.0.2a. Beginning in HXDP 6.0(1a) support for the ESXi STIG 7.0 will be available.

These can be found and downloaded from the Federal DISA site [here](#). These STIG settings are automated via script. Please note that DISA STIGs are dynamic, and as such, will be updated frequently, so you can expect this list to change. The corresponding automation in HXDP will follow suit. Some of these settings, while desirable for secure daily operation, have potential repercussions for cluster upgrade and expansion. As such, some settings may need to be temporarily disabled to accommodate changes of this nature. See the administration guide for your version of HXDP for instructions on running STIG automation and caveats around them for certain cluster operations.

Some settings derived from the DISA STIG set have become the default. For example, it is now the default to set promiscuous mode, forged transmits, and MAC change to REJECT in ESXi. Verify your version with respect to the STIG settings that are applied. Settings that ship with a specific version have been thoroughly tested with that version.

STIG settings can be found in the following file on any CVM for the ESXi, vCenter, and CVM settings.
`/usr/share/springpath/storfs-misc/hx-scripts/stig_config.ini`

For vCenter, only one parameter is currently set:
'config.nfc.useSSL' set to 'true'

A technote on the STIG settings can be found here:

https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/TechNotes/b_How_to_Configure_vCenter_Security_Hardening_Settings.html

Note that the STIG scripts must be run from each CVM node. SSH to each CVM management IP and change to the following path:

`/usr/share/springpath/storfs-misc/hx-scripts`



From this location run the `stig_security_settings.py` script. Note that this script has default configuration values in the `stig_config.ini` file located in the same directory. These may be edited as needed but will no longer match the vetted settings. Every setting set by the STIG script is idempotent, so multiple executions of the script will not adversely affect the system and you can reset your compliance baseline at any time by running it if things have changed in the interim.

A subset of the REST API contains STIG related entries. See the section on REST APIs above.

Starting with HXDP version 4.5.1a, there is no access to the root user without the tokenized workflow initiated by TAC, so you are unable to simply run the STIG script. You can, however, still activate and reset the STIG settings from the command line (Secure Admin Shell), using a CURL call to the HX API.

Here is the general format:

<https://developer.cisco.com/docs/ucs-dev-center-hyperflex/#!api-documentation-format/try-it-out>

Location of STIG APIs:

<https://developer.cisco.com/docs/ucs-dev-center-hyperflex/#!support-service>

To apply STIG run this API using CURL: configure stig parameters

SSL Certificate Replacement

During Cisco HyperFlex deployment, a set of local certificates is generated between the components to allow for trusted communication. Many organizations have their own certificate authority already in place. It is recommended that you replace the default SSL Certificates with your certificates. This process has been automated using scripts. Please see Appendix F for a full treatment of certificate management.

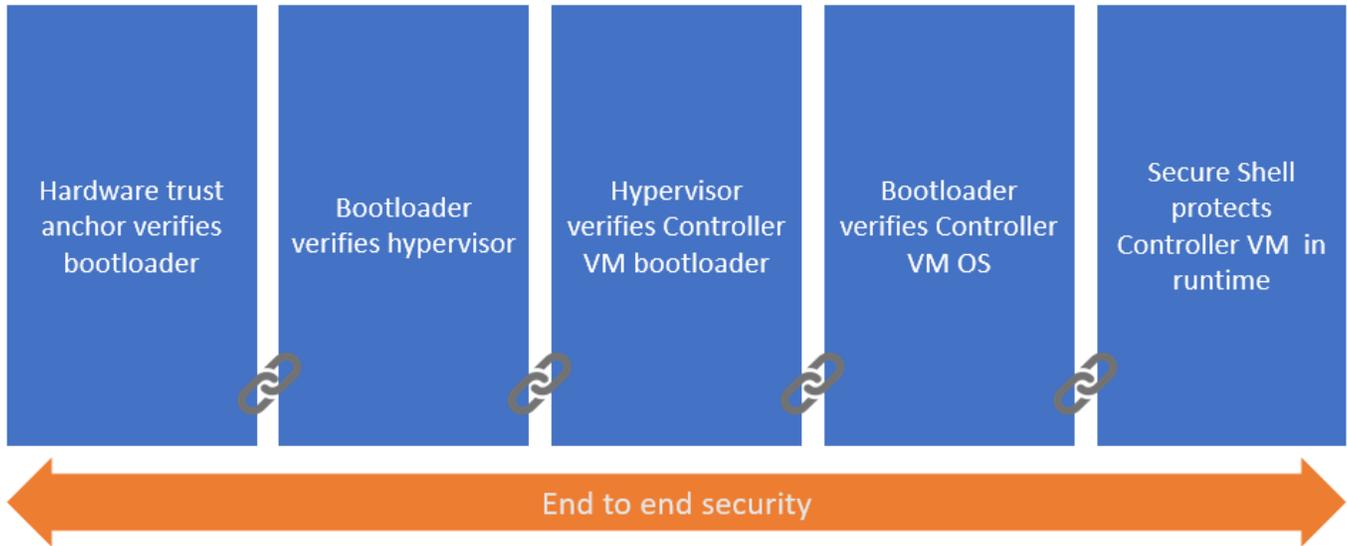
Secure Boot

Beginning with HXDP 4.5.1 (a) HX supports an option for secure boot through a setting on the UCS UEFI BIOS that is enabled in the HX Connect UI under the Upgrade tab. It is recommended that you use this operation to automate enabling UEFI secure boot on all nodes in the Hyperflex cluster. Once UEFI secure boot is enabled, this operation cannot be reversed from within HX Connect. Changing boot settings manually on UCS servers managed by HyperFlex is not recommended. SSH needs to be enabled and ESXi Lockdown mode disabled in order to enable secure boot.

Currently the secure boot process, when enabled, is in effect during boot including the hypervisor but not the control VM. CVM secure boot will be released in a future version. The end-to-end security model that this enables, when combined with the secure admin shell, encompasses hardware trust anchor, to secure hypervisor boot, to (eventual) CVM secure boot with access protected by a secure appliance shell. This is all externally verifiable using attestation with vCenter.

HX Secure Boot with Hardware Trust Anchor

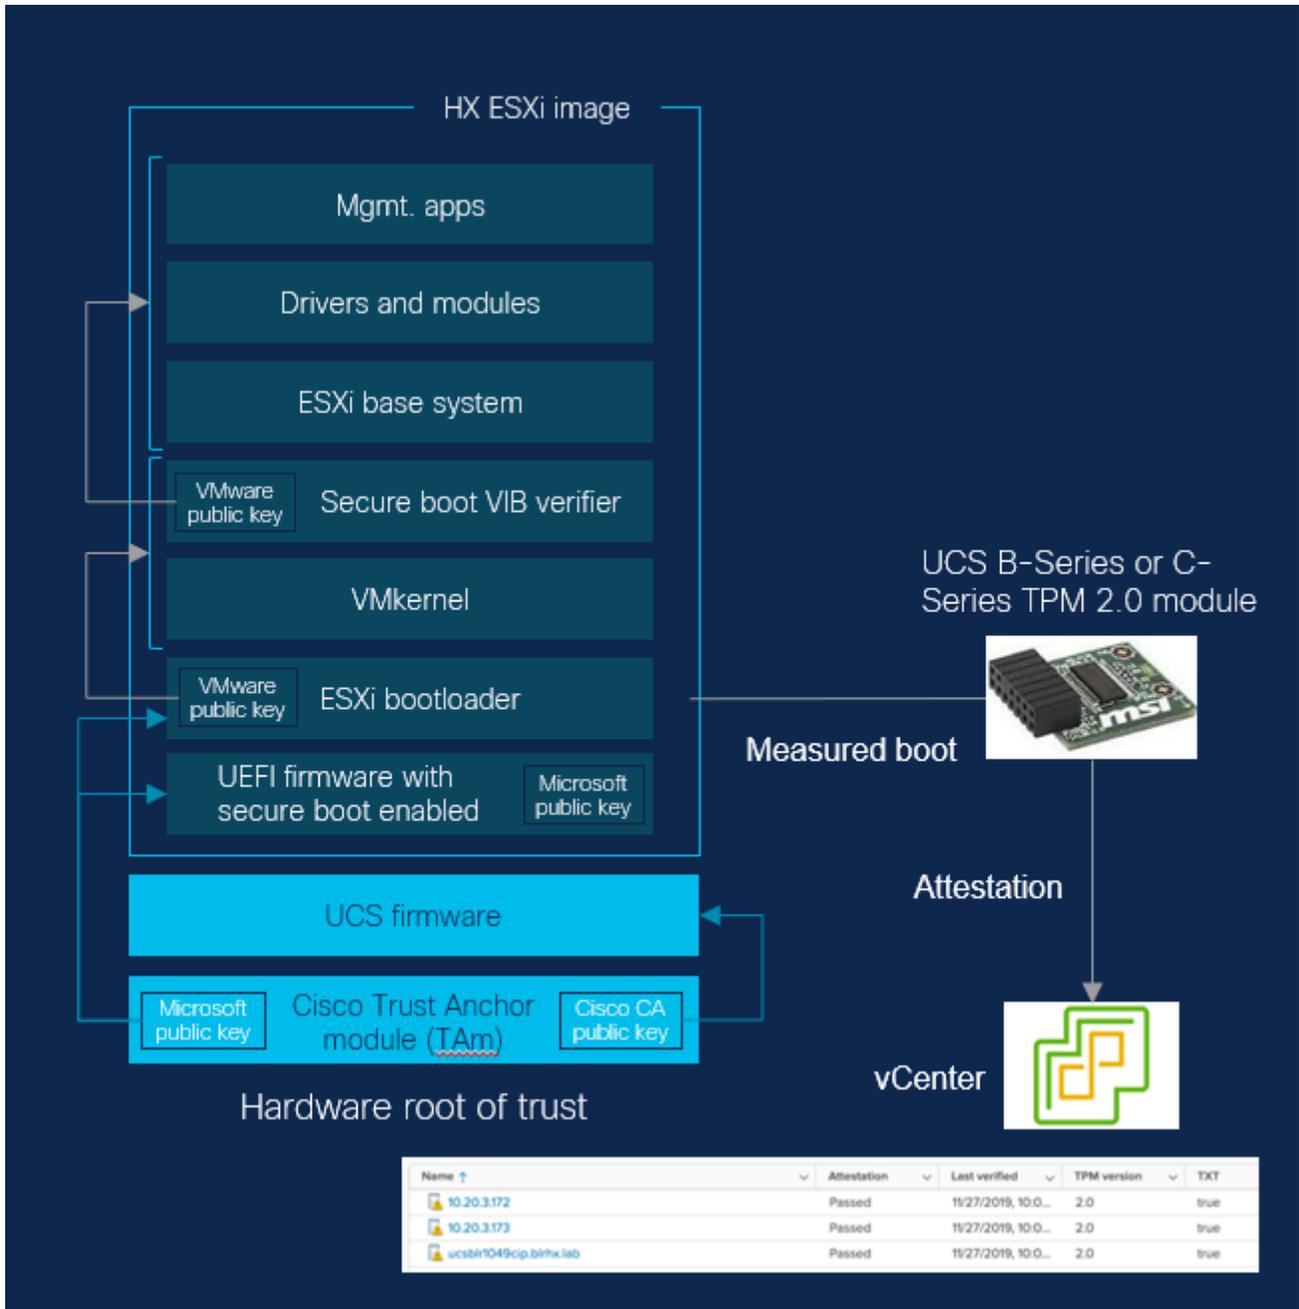
Secure Shell for HX Controller VM



The current implementation covers the following:

- Hypervisor secure boot, secured by public keys stored in the write-protected hardware root of trust
- Ensures only a trusted HX ESXi image, including drivers, is booted by verifying signatures
- Supports attestation of secure boot of ESXi by vCenter (requires min. ESXi 6.7 and TPM 2.0)

The detailed process flow for secure boot of the hypervisor and attestation capability is shown below. Note that the certificate-based hardware root of trust validates the UCS firmware which ensures a clean BIOS that is set for key validation of the hypervisor bootloader and so on. This guarantees that the hardware and hypervisor in the HX system have not been tampered with. External validation of this can be made through attestation with vCenter using the TPM 2.0 module in UCS.



The Cisco “HW Root of Trust” ensures secure boot by enabling a trusted hardware module. The Cisco IMC secure boot is handled via (Hardware) HW Root of Trust. Immutable keys are embedded in write-protected devices on every UCS server. Additionally, system BIOS secure boot is also encoded at manufacturing, and Cisco resolves both Firmware and BIOS via HW Root of Trust measures. Cisco also employs anti-counterfeit measures to ensure the physical hardware is authentic and signed by Cisco.

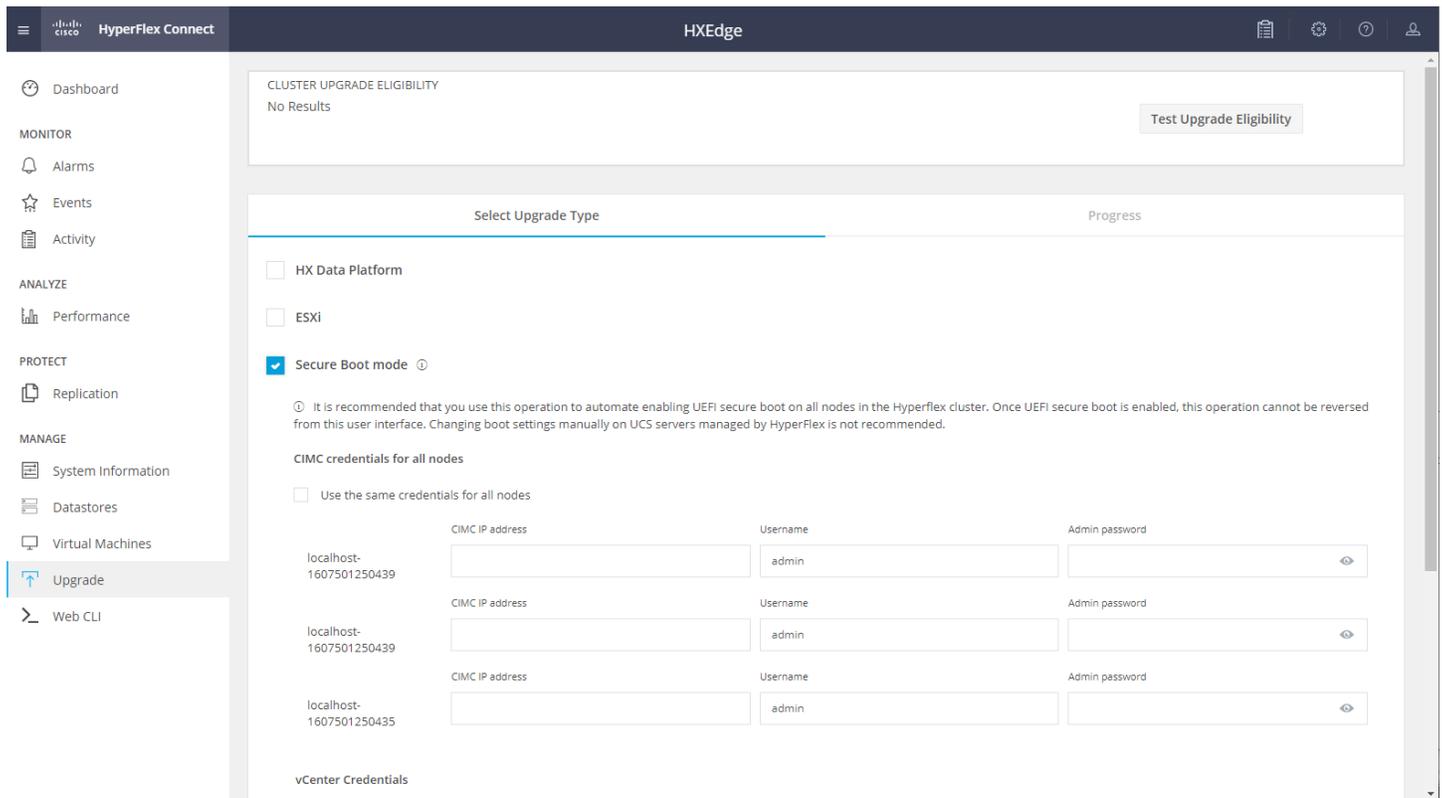
The HW Root of Trust is a Cisco ACT2 Trust Anchor Module (TAM). This module has the following characteristics:

- Immutable Identity with IEEE 802.1AR (Secure UDI- X.509 cert)
- Anti-Theft & Anti-Counterfeiting
- Built-In Cryptographic Functions
- Secure Storage for Certificates and Objects
- Certifiable NIST SP800-92 Random Number Generation

Once a system is securely booted, it is often important to get external verification that this is indeed the case. This is done through attestation. "Attestation" is evidence of a result, i.e., "The host was booted with secure boot enabled and signed code". Host Secure Boot Assurance via Attestation:

- Requires minimum ESXi 6.7 and Trusted Platform Module (TPM) 2.0
- TPM stores platform measurements of a known good boot; vCenter compares current boot against values stored in TPM
- For more information, see the link in the references at the end of this paper.

Secure boot is enabled in the HX Connect UI under the Upgrade tab.





Select the Secure Boot Mode checkbox to change the boot mode of the ESXi hosts from Legacy to UEFI Secure Boot anchored to the Cisco hardware root of trust on the Cisco Integrated Management Controller (CIMC). After Secure Boot is enabled, it cannot be disabled. Changing boot mode will cause a rolling reboot of each ESXi host in the HX cluster. A maintenance window is recommended for this operation. Enabling Secure Boot is a one-time operation and cannot be combined with other upgrade workflows.

To check the secure boot status in HX 5.x, you can select the "Check Secure Boot Status" in the drop-down Action Menu in the top right of HX Connect.

NOTE: The custom Cisco ESXi image has some first boot scripts that need to be run during first install or clean up. Secure boot with Edge and DC-no-FI prevents this first boot ESXi script dependence from running. Secure boot must be disabled for the portion of an install or re-install that requires a fresh ESXi image.

SSL Certificate Thumbprint (Hash) and Signatures

Many vendors have SSL thumbprints used on all HX certificates used the SHA1 hash. SHA1 has an astronomically rare chance (only achieved several years ago) of generating a hash collision that could potentially offer an avenue whereby two certificates have the same thumbprint. The thumbprint is used for reference and does not represent a security threat when used with SHA1 (see Windows certificates for example). However, HX uses a SHA256 hash to generate the thumbprint for enhanced uniqueness. All signatures are SHA256 and have been for a while.

Dynamic Self-Signed Certificates in HX

From 4.0.2a onwards, all certificates are dynamically generated and unique per cluster, but the same per CVM within that cluster. The requirement to keep the certificate the same per CVM within a cluster is to ensure that whichever node gets the management CIP (which can change during a failure event) will not break secure access to HX Connect.

During a fresh install of HX, a set of self-signed certificates is generated for secure intra-cluster communication. This certificate is created when you run the post-install script. The post-install script must be run on fresh installations.

The script creates a new self-signed certificate that is unique to the cluster. This certificate is pushed to each node. Automatic re-registration of the cluster with vCenter using the new certificate is performed. The post-install script can be re-run later to generate a new certificate if desired. This script is located in each CVM at the following path:

```
/usr/share/springpath/storfs-misc/hx-scripts/post_install.sh
```

Dynamic self-signed certificates created using the post-install script use the "admin" account and password when prompted. Upon upgrade of HXDP from pre-4.0.2a to 4.0.2a, and if self-signed certificates are used, the upgrade regenerates a unique certificate per cluster dynamically, and installs it on each CVM.

The post install script is a direct command in HX 4.5.1a and above. You invoke it with `hx_post_install` from the admin shell command line.



It is recommended to replace your self-signed certificates with CA signed certificates to improve your security posture.

UCSM Certificate Management

Setting UCSM certificates is covered in the UCSM management guide here:

<https://www.cisco.com/c/en/us/support/docs/servers-unified-computing/ucs-infrastructure-ucs-manager-software/213523-creating-and-using-3rd-party-certificate.html>

This can be automated using the UCSM API references for larger deployments.

Setting the CIMC/KVM certificate is a separate procedure. This can be conducted through the UCSM UI as described here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Admin-Management/4-0/b_Cisco_UCS_Admin_Mgmt_Guide_4-0/b_Cisco_UCS_Admin_Mgmt_Guide_4-0_chapter_0110.html#task_gxr_3qg_ncb

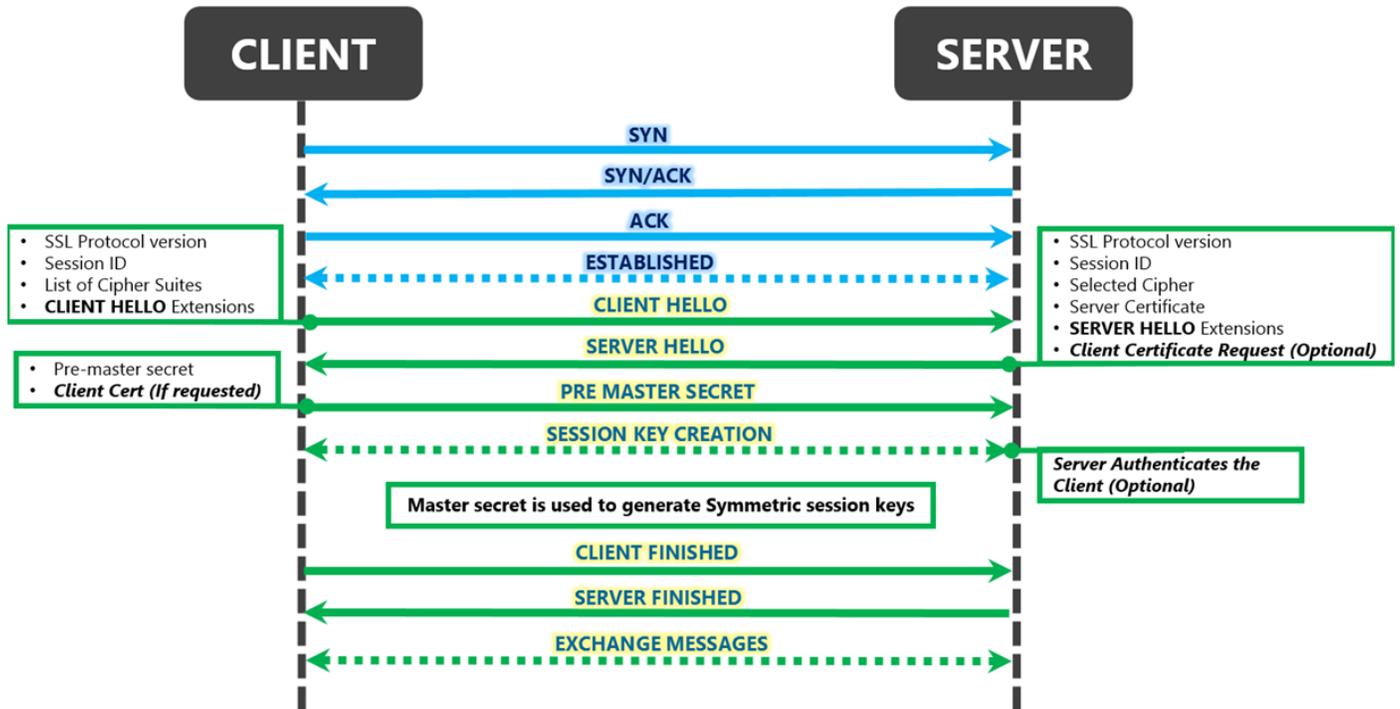
It is also possible to automate this using the XML API for IMC. See the reference here:

<https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-integrated-management-controller/products-programming-reference-guides-list.html>

When using these guides for HX CIMC/KVM be sure to use the C-series references.

HX and Perfect Forward Secrecy (PFS)

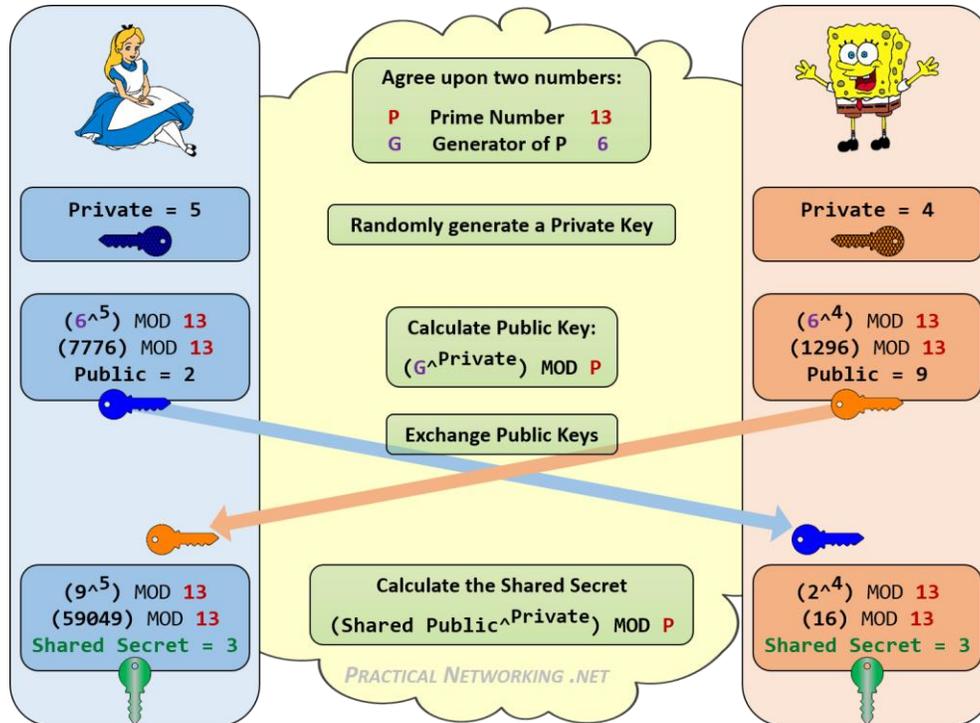
Perfect Forward Secrecy (PFS) is a mechanism to prevent private key compromise from affecting a larger set of historical or future communications. To understand the problem, let us first review how TLS handshake works.



Source: <https://blogs.msdn.microsoft.com/kaushal/2013/08/02/ssl-handshake-and-https-bindings-on-iis/>

As the diagram shows, the server's private key is used to encrypt the pre-master secret. The pre-master secret is used to derive the symmetric key which encrypts and decrypts messages in the TLS session. This dependency on the server private key means if an attack is somehow able to get a hold of the private key on the server, they will be able to guess at the session key. If the attacker has also been eavesdropping and recording the HTTP traffic, they can then decrypt and read any past message, and even future messages if the private key is not changed.

The solution is to disable RSA based key exchange (see the next two sections) in favor of newer Diffie-Hellman Ephemeral (Elliptic) Key Exchange methods. These methods are known as Perfect Forward Secrecy.



Source: <https://security.stackexchange.com/questions/45963/diffie-hellman-key-exchange-in-plain-english>

In this architecture, new parameters are used for each session. Private keys are not used to deduce the shared (session) key. Some environments require a subset of the default TLS ciphers to be disabled in order to accomplish this. The procedure to do this must be performed on each controller VM.

TLS Weak Protocol Disable

Some environments require a subset of the default TLS ciphers to be disabled. This procedure to do this must be performed on each controller VM.

On each controller VM, edit `/etc/nginx/conf.d/springpath.conf` and change the line starting with `ssl_protocols`:

Comment out the existing one, and replace with the new one as shown below:

```
#ssl_protocols    TLSv1 TLSv1.1 TLSv1.2;
ssl_protocols    TLSv1.2;
```

Save the file exit your editor (vi). Restart nginx using **service nginx restart** from the CVM CLI.

All components are strict TLS 1.2 implementations, including new provisioning and upgrades. HX does not support anything less than TLS 1.2.



TLS Weak Cipher Disable

To remove weak ciphers, you must edit the `/etc/nginx/conf.d/springpath.conf` file. Change the following section in bold:

```
server
{
    ### server port and name ###
    listen    443;
    ...
    ...
    ...

    ### Add SSL specific settings here ###
    ### Disable SSLv3 & RC4 cipher, to suppress POODLE & BEAST attacks
    ssl_protocols    TLSv1.2;
    ssl_ciphers    !AES256-SHA:!ECDHE-RSA-AES256-SHA:!ECDHE-RSA-DES-CBC3-SHA:!DES-CBC3-SHA:!ECDHE-RSA-AES128-SHA:!AES128-SHA:!aNULL:!eNULL:FIPS@STRENGTH:!RSA;
}
```

Note: append `!RSA` at the end.

Save the file exit your editor (`vi`). Restart nginx using **`service nginx restart`** from the CVM CLI.

SSH (ESX) Lockdown Mode and Root Logins

ESX SSH lockdown mode can be enabled on each ESX node of the HX cluster. This applies only to a post-install system. SSH traffic must not be blocked during install. SSH needs to be enabled before cluster expansion can take place. It can be disabled again afterwards.

HX snapshots and native replication do not use SSH to interface with ESXi, i.e., neither “root” nor “hxuser” based SSH logins are performed. With respect to logins to `hostd` (ESXi), for vSphere API access, only “hxuser” is used. The hxuser password is randomly generated on cluster creation during setup.

All nodes within the same cluster have the same hxuser password. This password is used in several workflows and upgrades. Changing this password is not supported without cluster re-creation. Root login is only used during cluster creation, node expansion, and initial installation.

Lockdown Mode is either Disabled, Normal or Strict. When Lockdown is enabled, the ESXi host can only be accessed through the vCenter server or the Direct Console User Interface (DCUI). Enabling Lockdown mode affects which users are authorized to access host services. Once Lockdown mode is enabled, and if root or `administrator@vsphere.local` or any other user is not part of the Exception user list, SSH to that ESX is not allowed. Similarly, if the host has been removed from the vCenter for some reason, adding the host back to vCenter is not allowed. Here is an overview of the features:

- Lockdown Mode exists in three states:
 - Disabled à Can SSH to host
 - Normal à Can connect through DCUI or VC
 - Strict à Can connect only using VC
- Upgrade checks whether Lockdown Mode is enabled
 - If enabled, prompts the user to disable for upgrade to proceed
- Upgrade will not proceed even in normal Lockdown mode

Normal vs. Strict mode have additional different behaviors and exceptions. For a comprehensive examination of system behavior in each mode and for troubleshooting guidelines for Lockdown, see the HyperFlex Installation Guide.

SSH, Lockdown, and HX Upgrades

Many systems will have various lockdowns and access methods disabled as part of an overall security posture. This is the recommended configuration. During upgrades of the HyperFlex software, however, some access methods and settings need to be enabled:

- Ensure SSH is ENABLED to the HX system
- Ensure that ESXi Lockdown mode is DISABLED
- Ensure that DRS is ENABLED in vCenter so that rolling upgrades can automatically migrate VMs
 - If this is not available or enabled, the upgrade will not proceed until VMs are manually migrated off of the node that is currently being upgraded.

Once the upgrade is complete, you can change the access methods back to the secure posture that was set prior to upgrade.

Tech Support Mode

Tech Support Mode, also called "Controller Access Over SSH", is specifically designed to allow for CVM troubleshooting.

- Tech Support Mode is enabled by default
 - Allows SSH access to the CVM management interface
- Tech Support Mode can be disabled
 - SSH to CVM management IP is disallowed
- Status of Tech Support Mode is listed in the status banner at the top of System Information in HX Connect
- If Tech Support Mode is disabled, the user will be prompted to enable it for upgrades to proceed

Third Party Software Execution on FIs and HXDP

Cisco does not support the installation of 3rd party software on either Fabric Interconnects (FIs) or on HXDP nodes (ESXi or HX CVM). For FI's, external software is not supported by virtue of the UCSM kernel-space type management shell. It is not possible to load or run any applications. For HXDP, Cisco does not recommend or support the installation and/or execution of 3rd party applications. In 4.0.X it is recommended that you use HXDP's tech support mode along with ESXi's lockdown mode at the same time in order to safeguard against



accidental or malicious attempts to run external applications on the HX CVM or the node hypervisor. In HX 4.5.1a and above, HXDP uses the admin shell, making this precaution redundant.

Whitelisting and other STCLI Security Commands

The HX datastores are a protected resource only mountable by HX nodes participating in the cluster (either by installation or by expansion). These protected datastore(s) cannot be mounted by other systems unless they are whitelisted. To whitelist a system for the cluster, ssh to a node and use the stcli security whitelist commands:

Remove systems from the list when not in immediate use.

```
root@SpringpathControllerEWA35H09RF:~# stcli security
usage: stcli security [-h] {password,whitelist,ssh,encryption} ...
```

```
root@SpringpathControllerEWA35H09RF:~# stcli security password
usage: stcli security password [-h] {set} ...
```

```
root@SpringpathControllerEWA35H09RF:~# stcli security whitelist
usage: stcli security whitelist [-h] {list,add,remove,clear} ...
```

```
root@SpringpathControllerEWA35H09RF:~# stcli security ssh
usage: stcli security ssh [-h] {resync} ...
```

```
root@SpringpathControllerEWA35H09RF:~# stcli security encryption
usage: stcli security encryption [-h] {ucsm-ro-user} ...
root@SpringpathControllerEWA35H09RF:~#
```

HX Data Platform Firewalling: IP Tables

Each HXDP node maintains a set of IP Tables firewall entries. This serves to explicitly set traffic that is allowed to communicate in and out with the HXDP node. The table is maintained automatically and should not have to be edited. These entries are listed for reference below. They are also automatically updated when HX Native Replication is enabled so that cluster-cluster traffic is permitted.

```
root@ucs-stctlv-137-1:~# iptables -S
-P INPUT ACCEPT
-P FORWARD ACCEPT
-P OUTPUT ACCEPT
-A INPUT -i lo -j ACCEPT
-A INPUT -m conntrack --ctstate RELATED,ESTABLISHED -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p tcp -m tcp --dport 443 -j ACCEPT
-A INPUT -d 10.a.b.d/32 -i eth0 -p tcp -m tcp --dport 443 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p tcp -m tcp --dport 443 -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p tcp -m tcp --dport 8888 -j ACCEPT
-A INPUT -d 10.a.b.d/32 -i eth0 -p tcp -m tcp --dport 8888 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p tcp -m tcp --dport 8888 -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p tcp -m tcp --dport 22 -j ACCEPT
```

```

-A INPUT -d 10.a.b.d/32 -i eth0 -p tcp -m tcp --dport 22 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p tcp -m tcp --dport 22 -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p tcp -m tcp --dport 80 -j ACCEPT
-A INPUT -d 10.a.b.d/32 -i eth0 -p tcp -m tcp --dport 80 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p tcp -m tcp --dport 80 -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p tcp -m tcp --dport 123 -j ACCEPT
-A INPUT -d 10.a.b.d/32 -i eth0 -p tcp -m tcp --dport 123 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p tcp -m tcp --dport 123 -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p udp -m udp --dport 427 -j ACCEPT
-A INPUT -d 10.a.b.d/32 -i eth0 -p udp -m udp --dport 427 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p udp -m udp --dport 427 -j ACCEPT
-A INPUT -d 10.a.b.c/32 -i eth0 -p udp -m udp --dport 8125 -j ACCEPT
-A INPUT -d 10.a.b.d/32 -i eth0 -p udp -m udp --dport 8125 -j ACCEPT
-A INPUT -d 10.a.b.e/32 -i eth1 -p udp -m udp --dport 8125 -j ACCEPT
-A INPUT -s 10.a.b.g/32 -d 10.a.b.l/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.g/32 -d 10.a.b.e/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.f/32 -d 10.a.b.l/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.f/32 -d 10.a.b.e/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.h/32 -d 10.a.b.l/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.h/32 -d 10.a.b.e/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.i/32 -d 10.a.b.l/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.i/32 -d 10.a.b.e/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.j/32 -d 10.a.b.l/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.j/32 -d 10.a.b.e/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.k/32 -d 10.a.b.l/32 -i eth1 -j ACCEPT
-A INPUT -s 10.a.b.k/32 -d 10.a.b.e/32 -i eth1 -j ACCEPT
-A INPUT -p udp -m udp --dport 32768:65535 -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -j DROP
root@ucs-stctlv-137-1:~#
root@ucs-stctlv-137-1:~# iptables -L
Chain INPUT (policy ACCEPT)
target prot opt source destination
ACCEPT all -- anywhere
ACCEPT all -- anywhere ctstate RELATED,ESTABLISHED
ACCEPT tcp -- anywhere ucs139-cip-m.eng.test-domain.com tcp dpt:https
ACCEPT tcp -- anywhere ucs-stctlv-137-1.eng.test-domain.com tcp dpt:https
ACCEPT tcp -- anywhere ucs-stctlv-137.eng.test-domain.com tcp dpt:https
ACCEPT tcp -- anywhere ucs139-cip-m.eng.test-domain.com tcp dpt:8888
ACCEPT tcp -- anywhere ucs-stctlv-137-1.eng.test-domain.com tcp dpt:8888
ACCEPT tcp -- anywhere ucs-stctlv-137.eng.test-domain.com tcp dpt:8888
ACCEPT tcp -- anywhere ucs139-cip-m.eng.test-domain.com tcp dpt:ssh
ACCEPT tcp -- anywhere ucs-stctlv-137-1.eng.test-domain.com tcp dpt:ssh
ACCEPT tcp -- anywhere ucs-stctlv-137.eng.test-domain.com tcp dpt:ssh
ACCEPT tcp -- anywhere ucs139-cip-m.eng.test-domain.com tcp dpt:http
ACCEPT tcp -- anywhere ucs-stctlv-137-1.eng.test-domain.com tcp dpt:http
ACCEPT tcp -- anywhere ucs-stctlv-137.eng.test-domain.com tcp dpt:http
ACCEPT tcp -- anywhere ucs139-cip-m.eng.test-domain.com tcp dpt:ntp
ACCEPT tcp -- anywhere ucs-stctlv-137-1.eng.test-domain.com tcp dpt:ntp
ACCEPT tcp -- anywhere ucs-stctlv-137.eng.test-domain.com tcp dpt:ntp
ACCEPT udp -- anywhere ucs139-cip-m.eng.test-domain.com udp dpt:svrloc
ACCEPT udp -- anywhere ucs-stctlv-137-1.eng.test-domain.com udp dpt:svrloc

```



```
ACCEPT udp -- anywhere ucs-stctlv-137.eng.test-domain.com udp dpt:svrloc
ACCEPT udp -- anywhere ucs139-cip-m.eng.test-domain.com udp dpt:8125
ACCEPT udp -- anywhere ucs-stctlv-137-1.eng.test-domain.com udp dpt:8125
ACCEPT udp -- anywhere ucs-stctlv-137.eng.test-domain.com udp dpt:8125
ACCEPT all -- ucs-stctlv-139.eng.test-domain.com ucs139-cip.eng.test-domain.com
ACCEPT all -- ucs-stctlv-139.eng.test-domain.com ucs-stctlv-137.eng.test-domain.com \\ ACCEPT all --
ucs139-v.eng.test-domain.com ucs139-cip.eng.test-domain.com
ACCEPT all -- ucs139-v.eng.test-domain.com ucs-stctlv-137.eng.test-domain.com \\ ACCEPT all -- ucs136-
v.eng.test-domain.com ucs139-cip.eng.test-domain.com
ACCEPT all -- ucs136-v.eng.test-domain.com ucs-stctlv-137.eng.test-domain.com \\ ACCEPT all -- ucs137-
v.eng.test-domain.com ucs139-cip.eng.test-domain.com
ACCEPT all -- ucs137-v.eng.test-domain.com ucs-stctlv-137.eng.test-domain.com \\ ACCEPT all -- ucs-stctlv-
138.eng.test-domain.com ucs139-cip.eng.test-domain.com
ACCEPT all -- ucs-stctlv-138.eng.test-domain.com ucs-stctlv-137.eng.test-domain.com \\ ACCEPT all --
ucs138-v.eng.test-domain.com ucs139-cip.eng.test-domain.com
ACCEPT all -- ucs138-v.eng.test-domain.com ucs-stctlv-137.eng.test-domain.com \\ ACCEPT udp --
anywhere udp dpts:32768:65535
ACCEPT icmp -- anywhere
DROP all -- anywhere
Chain FORWARD (policy ACCEPT)
target prot opt source destination \\

Chain OUTPUT (policy ACCEPT)
target prot opt source destination \\
```

Replication

Replication setting changes are maintained globally once replication is enabled on the cluster. Firewall entries are updated for ports needed for replication (see Networking Requirements above).

When replication is enabled, a new NIC is non-disruptively added to HXDP. This NIC is assigned an IP address in a new replication VLAN. The HX Service Profile on the FI (via UCSM) is automatically updated.

Replication traffic is not encrypted on the wire from the cluster. Secure replication requires an IPSEC capable WAN connection or relies on a trusted network. Data on the wire is always compressed so its general appearance is not plain text.

Specific ESX Environment Hardening Settings Relevant to HXDP

See [Appendix B](#) for a set of ESX hardening configuration settings. These items are general recommendations from the UCS verified ESX hardening guide.

Specific UCS Environment Hardening Settings Relevant to HXDP

The UCSM build used for the system must match the supported UCSM version in the preinstall checklist.

http://www.cisco.com/c/dam/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/HyperFlex_preinstall_checklist/Cisco_HX_Data_Platform_Preinstallation_Checklist_form.pdf



Refer to the UCS Hardening guide specifically for settings relevant to the build you are running.

Control VM (SCVM) Customization

Control VM customization is not supported and can be problematic. You should not modify the CVM hardware settings in vCenter. For example, the USB0 NIC interface is present in the CVM and some environments will try to remove any hardware devices from the CVM that do not seem to be in use. This should not be done.

USB0 is used for SED communications but should still be left alone in non-SED deployments. The IP address assigned to usb0 is a 169.254/16 IPv4 APIPA private address and is not routable. Iptables rules are also pre-configured so that all inbound packets to USB0 from any network will be dropped. In other words, an attacker would be required to first break into the Controller VM itself and gain local access in order to exploit the usb0 interface.

It is important to reiterate that changes to the CVM configuration like disabling this interface should not be made. It is configured on all clusters starting with the M5 platform and, although this communication channel may not be used today except for SED, in the future things like software encryption and other capabilities may start to use it and expect it to be present. Any other devices present in the factory CVM should, similarly, not be altered.

References

Intersight Security Guides and Certifications

- Star Registry Listing <https://cloudsecurityalliance.org/star/registry/cisco-systems/services/cisco-intersight/>
- FIPS 140-2 Compliance https://trustportal.cisco.com/c/r/ctp/trust-portal.html?search_keyword=intersight#/1631765240165158
- ISO/IEC 27001:2013 https://trustportal.cisco.com/c/r/ctp/trust-portal.html?search_keyword=intersight#/1608332940718694
- ISO/IEC 27017:2015 https://trustportal.cisco.com/c/r/ctp/trust-portal.html?search_keyword=intersight#/19695138127125035
- Cisco Intersight Platform Privacy Data Map https://trustportal.cisco.com/c/r/ctp/trust-portal.html?search_keyword=intersight#/1583721933988107
- Cisco Intersight Platform Security Brief https://trustportal.cisco.com/c/r/ctp/trust-portal.html?search_keyword=intersight#/1600118331971304
- Cisco Intersight SOC 3 https://trustportal.cisco.com/c/r/ctp/trust-portal.html?search_keyword=intersight#/1632370623703433

ESX Hardening Guide

- ESX <https://www.vmware.com/security/hardening-guides.html>



UCS Hardening Guide

- UCS <http://www.cisco.com/c/en/us/about/security-center/ucs-hardening.html>

Cisco CSDL

- CSDL <http://www.cisco.com/c/en/us/about/security-center/security-programs/secure-development-lifecycle/sdl-process.html>

Syslog-ng Configuration

- Syslog-ng configuration: <https://www.techrepublic.com/article/how-to-use-syslog-ng-to-collect-logs-from-remote-linux-machines/>

Secure Boot Assurance Via Attestation

- VMWare Blog: <https://www.yelof.com/2018/04/30/vsphere-6-7-esxi-and-tpm-2-0/>

Appendix A: Networking Ports

The following table lists the ports required for component communication for the HyperFlex solution.

Component	Service	Port	Protocol	Source	Destination	Notes
Time Server						
	NTP	123	UDP	Each ESX Node	Time Server	
				Each SCVM Node	Time Server	
				UCSM	Time Server	
HX Installer						
	SSH	22	TCP	HX Installer	Each ESX Node	Mgmt addresses
				HX Installer	Each SCVM Node	Mgmt addresses
				HX Installer	CIP-M	Cluster Mgmt
				HX Installer	UCSM	UCSM mgmt addresses
	HTTP	80	TCP	HX Installer	Each ESX Node	Mgmt addresses
				HX Installer	Each SCVM Node	Mgmt addresses
				HX Installer	CIP-M	Cluster Mgmt
				HX Installer	UCSM	UCSM mgmt addresses. (port not required, can be disabled)
	HTTPS	443	TCP	HX Installer	Each ESX Node	Mgmt addresses
				HX Installer	Each SCVM Node	Mgmt addresses
				HX Installer	CIP-M	Cluster Mgmt
				HX Installer	UCSM	UCSM mgmt addresses
	vSphere SDK	8089	TCP	HX Installer	Each ESX Node	Mgmt addresses
		9333	TCP	HX Installer	Each ESX Node	Cluster Data Network
	Heartbeat	902	TCP/UDP	HX Installer	vCenter, Each ESX Node	
	CIMC SoL	2400	TCP	CIMC OOB	Mgmt Network	bidirectional

	ICMP			HX Installer, CVM IPs	ESX/CVM IPs, vCenter	Mgmt addresses. Also required for cluster re- registration with vCenter.
Mail Server						
	SMTP	25	TCP	Each SCVM Node	Mail Server	
				CIP-M	Mail Server	
				UCSM	Mail Server	
Monitoring						
	SNMP Poll	161	UDP	Monitoring Server	UCSM	
	SNMP Trap	162	UDP	UCSM	Monitoring Server	
Name Server						
	DNS	53	TCP/UDP	Each ESX Node	Name Server	Mgmt addresses
				Each SCVM Node	Name Server	Mgmt addresses
				CIP-M	Name Server	Cluster Mgmt
				UCSM	Name Server	
vCenter						
	HTTP	80	TCP	vCenter	Each SCVM Node	Bidirectional
				vCenter	CIP-M	Bidirectional
	HTTPS (plugin)	443	TCP	vCenter	Each ESX Node	Bidirectional
				vCenter	Each SCVM Node	Bidirectional
				vCenter	CIP-M	Bidirectional
	HTTPS (VC SSO)	7444	TCP	vCenter	Each ESX Node	Bidirectional
				vCenter	Each SCVM Node	Bidirectional
				vCenter	CIP-M	Bidirectional
	HTTPS (plugin)	9443	TCP	vCenter	Each ESX Node	Bidirectional
				vCenter	Each SCVM Node	Bidirectional
				vCenter	CIP-M	Bidirectional
	CIM Server	5989	TCP	vCenter	Each ESX Node	
		9080	TCP	vCenter	Each ESX Node	Introduced in ESXi 6.5
	Heartbeat	902	TCP/UDP	vCenter	Each ESX Node	

User						
	SSH	22	TCP	User	Each ESX Node	Mgmt addresses
				User	Each SCVM Node	Mgmt addresses
				User	CIP-M	Cluster Mgmt
				User	HX Installer	
				User	UCSM	UCSM mgmt addresses
				User	vCenter	
				User	SSO Server	
	HTTP	80	TCP	User	Each SCVM Node	Mgmt addresses
				User	CIP-M	Cluster Mgmt
				User	UCSM	
				User	HX Installer	
				User	vCenter	
	HTTPS	443	TCP	User	Each SCVM Node	
				User	CIP-M	
				User	UCSM	UCSM mgmt addresses
				User	HX Installer	
				User	vCenter	
	HTTPS (SSO)	7444	TCP	User	vCenter	
				User	SSO Server	
	HTTPS (plugin)	9443	TCP	User	vCenter	
	KVM	2068	TCP	User	UCSM	UCSM mgmt addresses
SSO Server						
	HTTPS (SSO)	7444	TCP	SSO Server	Each ESX Node	Bidirectional
				SSO Server	Each SCVM Node	Bidirectional
				SSO Server	CIP-M	Bidirectional
Stretch Witness						
	Zookeeper	2181	TCP	Witness	Each CVM Node	Bidirectional, Mgmt Addresses
		2888	TCP	Witness	Each CVM Node	Bidirectional, Mgmt Addresses

		3888	TCP	Witness	Each CVM Node	Bidirectional, Mgmt Addresses
	Exhibitor (Zookeeper Lifecycle)	8180	TCP	Witness	Each CVM Node	Bidirectional, Mgmt Addresses
	HTTP	80	TCP	Witness	Each CVM Node	Potential Future Req.
	HTTPS	443	TCP	Witness	Each CVM Node	Potential Future Req.
Replication						
	ICMP			Each CVM Node	Each CVM Node	Include Cluster Mgmt IP
	Data Services Manager Peer	9338	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	Data Services Manager Peer	9339	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	Replication for CVM	3049	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	NRDR	9350	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	Cluster Map	4049	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	NR NFS	4059	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	Replication Service	9098	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	NR Master for Coordination	8889	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
UCSM						

	Encryption etc.	443	TCP	Each CVM Node	CIMC OOB	Bidirectional for each UCS node
	KVM	81	HTTP	User	UCSM	OOB KVM
	KVM	743	HTTPS	User	UCSM	OOB KVM Encrypted
Misc						
	Hypervisor Service	9350	TCP	Each CVM Node	Each CVM Node	Bidirectional, Include Cluster Mgmt IP as well
	CIP-M Failover	9097	TCP	Each CVM Node	Each CVM Node	Bidirectional for each CVM to other CVMs
	RPC Bind	111	TCP	Each CVM Node	Each CVM Node	
	Installer	8002	TCP	Each CVM Node	Installer	CVM outbound to Installer
	Apache Tomcat	8080	TCP	Each CVM Node	Each CVM Node	stDeploy makes connection, any request with uri /stdeploy
	Auth Service	8082	TCP	Each CVM Node	Each CVM Node	any request with uri /auth/
	hxRoboControl	9335	TCP	Each CVM Node	Each CVM Node	Robo deployments
	syslog-ng	6514	TCP	Each CVM Node	Remote syslog-ng collector	log aggregation
	TLS	5696	TCP	CIMC from each node	KMS Server	Bidirectional, Key Exchange
	Thrift RPC	10207	TCP	Management IP from each node	Management IP from each node	Platform thrift server port (closed to outside world) responds to requests from various internal management agents
SED Cluster						
	HTTPS	443	TCP	Each CVM Mgmt IP	UCSM A/B and VIP	Policy Configuration

				including CIP-M		
	TLS	5696	TCP	CIMC from each node	KMS Server	Key Exchange
	HXManager	9002	TCP	Node	Node	This is an internal service between nodes
	License Management Service	9346	TCP	Node	Node	This is an internal service between nodes
iSCSI						
	iSCSI	10152	TCP	Initiators	iSCSI CIP	SAN protocol

The following links are relevant specifically to ESXi and vCenter:

ESX port requirements:

<https://kb.vmware.com/s/article/2039095>

Comprehensive vSphere Port requirements

<https://ports.esp.vmware.com/home/vSphere+vSphere-7>

Please note that the following ports are shown as open but not needed for installation or general operation:

TCP port 81: HTTP KVM direct to CIMC (UCSM credentials required)

TCP port 743: HTTPS KVM direct to CIMC (UCSM credentials required)

TCP port 8888: Storage data network port for file system rebuilds

TCP port 843: UCS Central port on the FI for application integration

Note: UDP ports 427 (Service Location Protocol) and 8125 (Graphite) are open on the SCVM. Ports 32k-65k are also open for SCVM outbound communication. These UDP ports can be seen in the IP Tables ACCEPT syntax above.

Note: Ports 80 and 443 are required, bidirectionally, for tools.cisco.com to function properly with the smart call home functionality.

Note: Cisco HyperFlex does not use Cisco Discovery Protocol (CDP).

For Container Storage Infrastructure (CSI), the following ports must be available:

- Control plane:
 - From the Kubernetes nodes to the HX Storage Control VIP port 443 for API calls to provision the LUNS
- Data Plane:



- From the Kubernetes nodes to the HX nodes, on the iSCSI network, port 3260 and 860 for mounting the iSCSI LUNS
- The apps run inside a containers that have access to the iSCSI LUNS which are mounted on the Kubernetes hosts. You only need connectivity from the Kubernetes hosts to the HX iSCSI IPs.
- Port 860 for iSCSI is no longer the default and HX is not using it. See <https://www.rfc-editor.org/rfc/rfc7143.html>

Please note that the following ports are required for successful operation of the HX Profiler in order to gather input for the HX Sizer:

- vCenter to profiler: 443(https)
- hyperV/LBM/WBM to profiler: 5986(HTTPS)/5985(HTTP) for remote powershell & WMIC query execution

Connectivity between a SED-enabled cluster and the KMS have a few requirements in environments that are firewall segmented. Access between FI-A/FI-B/VIP and the KMS is not required. Only the CIMC IPs of each node need to access the KMS. Allowing CIMC IPs along with KMS IP and port 5696 is sufficient.

NRDR Port Requirement summary: ICMP, 3049, 4049, 4059, 8889, 9098, 9338, 9339, and 9350 for replication. See the table above for specifics.

Intersight

Network Communication Requirements for CIMC:

- Communication between CIMC and vCenter via ports 80, 443 and 8089 during installation phase.
- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the Hyperflex Hardening Guide.
- This communication needs to be persistent. It is required for any and all upgrades (including firmware), monitoring, and UI cross-launch.
- CIMC to Intersight should only require 443. Per the preinstall guide:
- *All device connectors must properly resolve svc.intersight.com and allow outbound-initiated HTTPS connections on port 443. The current HX Installer supports the use of an HTTP proxy. The IP addresses of ESXi management must be reachable from Cisco UCS Manager over all the ports that are listed as being needed from installer to ESXi management, to ensure deployment of ESXi management from Cisco Intersight.*
- Allow port 22 between the UCSM (or CIMC) VLAN and the ESXi/SCVM management VLAN

Intersight Connectivity Consider the following prerequisites pertaining to Intersight connectivity:

- Before installing the HX cluster on a set of HX servers, make sure that the device connector on the corresponding Cisco IMC instance is properly configured to connect to Cisco Intersight and claimed.
- Communication between CIMC and vCenter via ports 80, 443 and 8089 during installation phase.
- All device connectors must properly resolve svc.intersight.com and allow outbound initiated HTTPS connections on port 443. The current version of the HX Installer supports the use of an HTTP proxy.
- All controller VM management interfaces must properly resolve svc.intersight.com and allow outbound initiated HTTPS connections on port 443. The current version of HX Installer supports the use of an HTTP proxy if direct Internet connectivity is unavailable.
- *IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter*

server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the *Hyperflex Hardening Guide*.

- Starting with HXDP release 3.5(2a), the Intersight installer does not require a factory installed controller VM to be present on the HyperFlex servers. When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.
- Post-cluster deployment, the new HX cluster is automatically claimed in Intersight for ongoing management.

Appendix B: URLs Needed for Smart Call Home, Post Install Scripts, Intersight

Quick List

- `diag.hyperflex.io`
- `svc.intersight.com`
- `svc.ucs-connect.com`
- `tools.cisco.com` ports 80 and 443
- `upload.hyperflex.io`
- `ftp.springpath.com`

These are for autosupport, Smart Call Home, and the device connector to Intersight for management, trending (diags), install, upgrade. It is recommended to add the corresponding IP addresses for these FQDNs as well in the event DNS is unavailable. Access to `cisco.com` URLs is only needed in HXDP 4.x.

For example:

the following external URLs exist in some scripts:

```
/usr/share/springpath/storfs-misc/hx-scripts/support.py:
```

```
def verifyConnectivity():
```

```
    try:
```

```
        try:
```

```
            url = "https://upload.hyperflex.io/admin/api2/ping"
```

```
            r = requests.get(url, verify=False)
```

```
            return True
```

```
        except requests.exceptions.ConnectionError:
```

```
            url = "https://38.140.50.205/admin/api2/ping"
```

```
            r = requests.get(url, verify=False)
```

```
        outputfile = open('/etc/hosts', 'a')
```



```
entry = "\n38.140.50.205" + "\t upload.hyperflex.io \n"
```

and:

```
/usr/share/springpath/storfs-misc/hx-scripts/check_vswitch.py:
```

```
def uploadSupportBundle(fileName, folderTag, bundle):
```

```
try:
```

```
    folder = '/upload/' + folderTag + '/'
```

```
    data = {'command': 'mkdir', 'path': folder}
```

```
    auth = ('upload', 'upload')
```

```
    r = requests.post("https://54.88.201.239",  
                    verify=False, auth=auth, data=data)
```

```
    fileobj = open(bundle, 'rb')
```

```
    files = {'uploadPath': (None, folder), 'the_action': (  
        None, 'STOR'), 'file': (fileName, fileobj, 'application/x-gzip ')}
```

```
    requests.post("https://54.88.201.239",  
                verify=False, auth=auth, files=files)
```

54.88.201.239 reverse resolves to ftp.springpathinc.com

Smart Call Home (SCH):

```
root@hx-6-scvms-01:~# stcli services sch show
```

```
proxyPort: 8080
```

```
enableProxy: True
```

```
enabled: True
```

```
proxyPassword:
```

```
proxyUser:
```

```
cloudEnvironment: production
```

```
proxyUrl: proxy.esl.cisco.com
```

```
emailAddress: dummy\_address@cisco.com
```

```
portalUrl:
```

```
cloudAsupEndpoint: https://diag.hyperflex.io/
```

```
root@hx-6-scvms-01:~#
```

Post Install:

```
root@Cisco-HX-Installer-Appliance:~# vi /usr/share/springpath/storfs-misc/hx-scripts/update.sh
```

```
#!/bin/sh
```

```
FILENAME="hx-tools.zip"
```

```
URL="http://cs.co/hx-scripts"
```

```
Cisco HX Platform Hardening Guide
```



```
cd /usr/share/springpath/storfs-misc/hx-scripts
wget --no-check-certificate -q -T1 -t1 ${URL} -O ${FILENAME} > /dev/null 2>&1

if [ $? -gt 0 ]; then
    echo "Could not download latest tools. Please verify internet connection"
    rm -f ${FILENAME} > /dev/null 2>&1
    exit 1
fi

unzip -oj ${FILENAME} > /dev/null 2>&1
rm -f ${FILENAME} > /dev/null 2>&1
echo "Scripts successfully updated"
```

Intersight Device connector:

- svc.intersight.com **(Preferred)**
- svc.ucs-connect.com **(Will be deprecated in the future)**

Appendix C: ESX Hardening Settings

ESX hardening settings:

ESXi.apply-patches	Keep ESXi system properly patched	By staying up to date on ESXi patches, vulnerabilities in the hypervisor can be mitigated. An educated attacker can exploit known vulnerabilities when attempting to attain access or elevate privileges on an ESXi host.
ESXi.audit-exception-users	Audit the list of users who are on the Exception Users List and whether they have administrator privileges	In vSphere 6.0 and later, you can add users to the Exception Users list from the vSphere Web Client. These users do not lose their permissions when the host enters lockdown mode. Usually, you may want to add service accounts such as a backup agent to the Exception Users list. Verify that the list of users who are exempted from losing permissions is legitimate and as needed per your environment. Users who do not require special permissions should not be exempted from lockdown mode.

ESXi.config-ntp	Configure NTP time synchronization	By ensuring that all systems use the same relative time source (including the relevant localization offset), and that the relative time source can be correlated to an agreed-upon time standard (such as Coordinated Universal Time—UTC), you can make it simpler to track and correlate an intruder's actions when reviewing the relevant log files. Incorrect time settings can make it difficult to inspect and correlate log files to detect attacks and can make auditing inaccurate.
ESXi.config-persistent-logs	Configure persistent logging for all ESXi host	ESXi can be configured to store log files on an in-memory file system. This occurs when the host's "/scratch" directory is linked to "/tmp/scratch". When this is done only a single day's worth of logs are stored at any time. In addition, log files will be reinitialized upon each reboot. This presents a security risk as user activity logged on the host is only stored temporarily and will not persistent across reboots. This can also complicate auditing and make it harder to monitor events and diagnose issues. ESXi host logging should always be configured to a persistent datastore.
ESXi.config-snmp	Ensure proper SNMP configuration	If SNMP is not being used, it should remain disabled. If it is being used, the proper trap destination should be configured. If SNMP is not properly configured, monitoring information can be sent to a malicious host that can then use this information to plan an attack. Note: ESXi 5.1 and later supports SNMPv3 which provides stronger security than SNMPv1 or SNMPv2, including key authentication and encryption.
ESXi.disable-mob	Disable Managed Object Browser (MOB)	The managed object browser (MOB) provides a way to explore the object model used by the VMkernel to manage the host; it enables configurations to be changed as well. This interface is meant to be used primarily for debugging the vSphere SDK. In Sphere 6.0 this is disabled by default

ESXi.firewall-enabled	Configure the ESXi host firewall to restrict access to services running on the host	Unrestricted access to services running on an ESXi host can expose a host to outside attacks and unauthorized access. Reduce the risk by configuring the ESXi firewall to only allow access from authorized networks.
ESXi.set-account-auto-unlock-time	Set the time after which a locked account is automatically unlocked	<p>Multiple account login failures for the same account could possibly be a threat vector trying to brute force the system or cause denial of service. Such attempts to brute force the system should be limited by locking out the account after reaching a threshold.</p> <p>In case, you would want to auto unlock the account, i.e., unlock the account without administrative action, set the time for which the account remains locked. Setting a high duration for which account remains locked would deter and severely slow down the brute force method of logging in.</p>
ESXi.set-account-lockout	Set the count of maximum failed login attempts before the account is locked out	Multiple account login failures for the same account could possibly be a threat vector trying to brute force the system or cause denial of service. Such attempts to brute force the system should be limited by locking out the account after reaching a threshold.
ESXi.set-dcui-access	Set DCUI.Access to allow trusted users to override lockdown mode	Lockdown mode disables direct host access requiring that admins manage hosts from vCenter Server. However, if a host becomes isolated from vCenter Server, the admin is locked out and can no longer manage the host. If you are using normal lockdown mode, you can avoid becoming locked out of an ESXi host that is running in lockdown mode, by setting DCUI.Access to a list of highly trusted users who can override lockdown mode and access the DCUI. The DCUI is not running in strict lockdown mode.
ESXi.set-dcui-timeout	Audit DCUI timeout value	DCUI is used for directly logging into ESXi host and carrying out host management tasks. The idle connections to DCUI must be terminated to avoid any unintended usage of the DCUI originating from a left-over login session.

ESXi.set-password-policies	Establish a password policy for password complexity	ESXi uses the pam_passwdqc.so plug-in to set password strength and complexity. It is important to use passwords that are not easily guessed and that are difficult for password generators to determine. Password strength and complexity rules apply to all ESXi users, including root. They do not apply to Active Directory users when the ESX host is joined to a domain. Those password policies are enforced by AD.
ESXi.set-shell-interactive-timeout	Set a timeout to automatically terminate idle ESXi Shell and SSH sessions	If a user forgets to log out of their SSH session, the idle connection will remain open indefinitely, increasing the potential for someone to gain privileged access to the host. The ESXiShellInteractiveTimeout allows you to automatically terminate idle shell sessions.
ESXi.set-shell-timeout	Set a timeout to limit how long the ESXi Shell and SSH services are allowed to run	When the ESXi Shell or SSH services are enabled on a host they will run indefinitely. To avoid having these services left running set the ESXiShellTimeOut. The ESXiShellTimeOut defines a window of time after which the ESXi Shell and SSH services will automatically be terminated.
ESXi.TransparentPageSharing-intra-enabled	Ensure default setting for intra-VM TPS is correct	<p>Acknowledgement of the recent academic research that leverages Transparent Page Sharing (TPS) to gain unauthorized access to data under certain highly controlled conditions and documents VMware's precautionary measure of restricting TPS to individual virtual machines by default in upcoming ESXi releases. At this time, VMware believes that the published information disclosure due to TPS between virtual machines is impractical in a real-world deployment.</p> <p>VMs that do not have the sched.mem.pshare.salt option set cannot share memory with any other VMs.</p>

vCenter.verify-nfc-ssl	Enable SSL for Network File copy (NFC)	<p>NFC (Network File Copy) is the name of the mechanism used to migrate or clone a VM between two ESXi hosts over the network.</p> <p>***By default, NFC over SSL is enabled (i.e.: "True") within a vSphere cluster but the value of the setting is null.***</p> <p>Clients check the value of the setting and default to not using SSL for performance reasons if the value is null. This behavior can be changed by ensuring the setting has been explicitly created and set to "True". This will force clients to use SSL.</p>
VM.disable-console-copy	Explicitly disable copy/paste operations	Copy and paste operations are disabled by default. However, if you explicitly disable this feature audit controls can check that this setting is correct.
VM.disable-console-drag-n-drop	Explicitly disable copy/paste operations	<p>Copy and paste operations are disabled by default however by explicitly disabling this feature it will enable audit controls to check that this setting is correct.</p> <p>The default value is null. Setting this to true is just for audit.</p>
VM.disable-console-gui-options	Explicitly disable copy/paste operations	Copy and paste operations are disabled by default however by explicitly disabling this feature it will enable audit controls to check that this setting is correct.
VM.disable-console-paste	Explicitly disable copy/paste operations	Copy and paste operations are disabled by default, however, if you explicitly disable this feature, audit controls can check that this setting is correct.

<p>VM.disable-disk-shrinking-shrink</p>	<p>Disable virtual disk shrinking</p>	<p>Shrinking a virtual disk reclaims unused space in it. The shrinking process itself, which takes place on the host, reduces the size of the disk's files by the amount of disk space reclaimed in the wipe process. If there is empty space in the disk, this process reduces the amount of space the virtual disk occupies on the host drive. Normal users and processes—that is, users and processes without root or administrator privileges—within virtual machines have the capability to invoke this procedure. A non-root user cannot wipe the parts of the virtual disk that require root-level permissions. However, if this is done repeatedly, the virtual disk can become unavailable while this shrinking is being performed, effectively causing a denial of service. In most datacenter environments, disk shrinking is not done, so you should disable this feature. Repeated disk shrinking can make a virtual disk unavailable. Limited capability is available to non-administrative users in the guest.</p>
<p>VM.disable-disk-shrinking-wiper</p>	<p>Disable virtual disk shrinking</p>	<p>Shrinking a virtual disk reclaims unused space in it. VMware Tools reclaims all unused portions of disk partitions (such as deleted files) and prepares them for shrinking. Wiping takes place in the guest operating system. If there is empty space in the disk, this process reduces the amount of space the virtual disk occupies on the host drive. Normal users and processes—that is, users and processes without root or administrator privileges—within virtual machines have the capability to invoke this procedure. A non-root user cannot wipe the parts of the virtual disk that require root-level permissions. However, if this is done repeatedly, the virtual disk can become unavailable while this shrinking is being performed, effectively causing a denial of service. In most datacenter environments, disk shrinking is not done, so you should disable this feature. Repeated disk shrinking can make a virtual disk</p>

		<p>unavailable. Limited capability is available to non-administrative users in the guest.</p>
VM.disable-hgfs	Disable HGFS file transfers	<p>Certain automated operations such as automated tools upgrades use a component in the hypervisor called "Host Guest File System" and an attacker could potentially use this to transfer files inside the guest OS</p>
VM.disconnect-devices-floppy	Disconnect unauthorized devices	<p>Ensure that no device is connected to a virtual machine if it is not required. For example, serial and parallel ports are rarely used for virtual machines in a datacenter environment, and CD/DVD drives are usually connected only temporarily during software installation. For less commonly used devices that are not required, either the parameter should not be present, or its value must be FALSE. NOTE: The parameters listed are not sufficient to ensure that a device is usable; other required parameters specify how each device is instantiated. Any enabled or connected device represents a potential attack channel.</p> <p>When setting is set to FALSE, functionality is disabled, however the device may still show up within the guest operation system.</p>

<p>VM.disconnect-devices-parallel</p>	<p>Disconnect unauthorized devices</p>	<p>Ensure that no device is connected to a virtual machine if it is not required. For example, serial and parallel ports are rarely used for virtual machines in a datacenter environment, and CD/DVD drives are usually connected only temporarily during software installation. For less commonly used devices that are not required, either the parameter should not be present, or its value must be FALSE. NOTE: The parameters listed are not sufficient to ensure that a device is usable; other required parameters specify how each device is instantiated. Any enabled or connected device represents a potential attack channel.</p> <p>When setting is set to FALSE, functionality is disabled, however the device may still show up within the guest operation system.</p>
<p>VM.disconnect-devices-serial</p>	<p>Disconnect unauthorized devices</p>	<p>Ensure that no device is connected to a virtual machine if it is not required. For example, serial and parallel ports are rarely used for virtual machines in a datacenter environment, and CD/DVD drives are usually connected only temporarily during software installation. For less commonly used devices that are not required, either the parameter should not be present, or its value must be FALSE. NOTE: The parameters listed are not sufficient to ensure that a device is usable; other required parameters specify how each device is instantiated. Any enabled or connected device represents a potential attack channel.</p> <p>When setting is set to FALSE, functionality is disabled, however the device may still show up within the guest operation system.</p>
<p>VM.limit-setinfo-size</p>	<p>Limit informational messages from the VM to the VMX file</p>	<p>The configuration file containing these name-value pairs is limited to a size of 1MB. This 1MB capacity should be sufficient for most cases, but you can change this value if necessary. You might increase this value if large amounts of custom information are being</p>

		<p>stored in the configuration file. The default limit is 1MB;this limit is applied even when the sizeLimit parameter is not listed in the .vmx file. Uncontrolled size for the VMX file can lead to denial of service if the datastore is filled.</p>
<p>VM.prevent-device-interaction-connect</p>	<p>Prevent unauthorized removal, connection, and modification of devices</p>	<p>In a virtual machine, users and processes without root or administrator privileges can connect or disconnect devices, such as network adaptors and CD-ROM drives, and can modify device settings. Use the virtual machine settings editor or configuration editor to remove unneeded or unused hardware devices. If you want to use the device again, you can prevent a user or running process in the virtual machine from connecting, disconnecting, or modifying a device from within the guest operating system. By default, a rogue user with non-administrator privileges in a virtual machine can:</p> <ol style="list-style-type: none"> 1. Connect a disconnected CD-ROM drive and access sensitive information on the media left in the drive 2. Disconnect a network adaptor to isolate the virtual machine from its network, which is a denial of service 3. Modify settings on a device
<p>VM.prevent-device-interaction-edit</p>	<p>Prevent unauthorized removal, connection, and modification of devices</p>	<p>In a virtual machine, users and processes without root or administrator privileges can connect or disconnect devices, such as network adaptors and CD-ROM drives, and can modify device settings. Use the virtual machine settings editor or configuration editor to remove unneeded or unused hardware devices. If you want to use the device again, you can prevent a user or running process in the virtual machine from connecting, disconnecting, or modifying a device from within the guest operating system. By default, a rogue user with non-administrator privileges in a virtual machine can:</p> <ol style="list-style-type: none"> 1. Connect a disconnected CD-ROM drive

		<p>and access sensitive information on the media left in the drive</p> <ol style="list-style-type: none"> 2. Disconnect a network adaptor to isolate the virtual machine from its network, which is a denial of service 3. Modify settings on a device
VM.restrict-host-info	Do not send host information to guests	<p>By enabling a VM to get detailed information about the physical host, an adversary could potentially use this information to inform further attacks on the host.</p> <p>If set to "True" a VM can obtain detailed information about the physical host. *The default value for the parameter is False but is displayed as Null. Setting to False is purely for audit purposes.*</p> <p>This setting should not be TRUE unless a particular VM requires this information for performance monitoring.</p>
VM.verify-network-filter	Control access to VMs through the dvfilter network APIs	<p>An attacker might compromise a VM by making use the dvFilter API. Configure only those VMs to use the API that need this access.</p> <p>This setting is considered an "Audit Only" guideline. If there is a value present, the admin should check it to ensure it is correct.</p>

<p>VM.verify-PCI-Passthrough</p>	<p>Audit all uses of PCI or PCIe passthrough functionality</p>	<p>Using the VMware DirectPath I/O feature to pass through a PCI or PCIe device to a virtual machine results in a potential security vulnerability. The vulnerability can be triggered by buggy or malicious code running in privileged mode in the guest OS, such as a device driver. Industry-standard hardware and firmware does not currently have sufficient error containment support to make it possible for ESXi to close the vulnerability fully.</p> <p>There can be a valid business reason for a VM to have this configured. This is an audit-only guideline. You should be aware of what virtual machines are configured with direct passthrough of PCI and PCIe devices and ensure that their guest OS is monitored carefully for malicious or buggy drivers that could crash the host.</p>
<p>vNetwork.limit-network-healthcheck</p>	<p>Enable VDS network healthcheck only if you need it</p>	<p>Network Healthcheck is disabled by default. Once enabled, the healthcheck packets contain information on host#, vds# port#, which an attacker would find useful. It is recommended that network healthcheck be used for troubleshooting and turned off when troubleshooting is finished.</p>
<p>vNetwork.restrict-netflow-usage</p>	<p>Ensure that VDS Netflow traffic is only being sent to authorized collector IPs</p>	<p>The vSphere VDS can export Netflow information about traffic crossing the VDS. Netflow exports are not encrypted and can contain information about the virtual network making it easier for a MITM attack to be executed successfully. If Netflow export is required, verify that all VDS Netflow target IPs are correct.</p>
<p>vNetwork.restrict-port-level-overrides</p>	<p>Restrict port-level configuration overrides on VDS</p>	<p>Port-level configuration overrides are disabled by default. Once enabled, this allows for different security settings to be set from what is established at the Port-Group level. There are cases where particular VM's require unique configurations, but this should be monitored so it is only used when authorized. If overrides are not monitored,</p>

		anyone who gains access to a VM with a less secure VDS configuration could surreptitiously exploit that broader access.
--	--	---



Appendix D: Acronym Glossary

AAA - authentication, authorization, and accounting

AD – Active Directory

API – Application Programming Interface

CC – Common Criteria

CERT – Computer Emergency Response Team, or Certificate

CIMC – Cisco Integrated Management Console

CIP – Cluster IP (data)

CIP-M – Cluster IP Management

CIS – Center for Internet Security

CLI – Command Line Interface

CSDL – Cisco Secure Development Lifecycle

CVM – Control Virtual Machine

CMVP – Cryptographic Module Validation Program

CSR – Certificate Signing Request

DISA – Defense Information Systems Agency

DNS – Domain Name Service

DSM – Vormetric Data Security Manager

DRS – Distributed Resource Scheduler

EAL – Evaluation Assurance Level

ESX - ESXi replaces Service Console (a rudimentary operating system) with a more closely integrated OS.

ESX/ESXi is the primary component in the VMware Infrastructure software suite. The name ESX originated as an abbreviation of **Elastic Sky X**

FedRAMP -- Federal Risk and Authorization Management Program

FI – Fabric Interconnect

FIPS – Federal Information Processing Standard

FISMA -- Federal Information Security Management Act

FQDN – Fully Qualified Domain Name

HX/HXDP – HyperFlex/HyperFlex Data Platform

IMC – Integrated Management Console

ISO – International Standards Organization

KMIP – Key Management Interoperability Protocol

KMS – Key Management Server

KVM – Keyboard Video Mouse

LAN – Local Area Network

MAC – Media Access Control (unique identifier)

MM – Maintenance Mode

NERC -- North American Electric Reliability Corporation Critical Infrastructure Protection

NTP – Network Time Protocol

OOB – Out of Band

POC – Proof of Concept

QoS – Quality of Service



REST – Representational State Transfer
RHEL – Red Hat Enterprise Linux
SCH – Smart Call Home
SCVM – Storage Control Virtual Machine
SED – Self Encrypting Drive
SL – Smart Licensing
SLP – Service Location Protocol
SNMP – Simple Network Monitoring Protocol
SSH – Secure Shell
SSL – Secure Sockets Layer
SSO – Single Sign On
STCLI -- Storage Command Line Interface
TLS – Transport Layer Security
TPM – Trusted Platform Module
UCARP – Userland Common Address Redundancy Protocol
UCS – Unified Computing System
UCSM – UCS Manager
UI – User Interface
VLAN – Virtual Local Area Network
VM – Virtual Machine
vNIC – Virtual Network Interface Card
vWAAS -- Virtual Wide Area Application Services (WAN acceleration device)
WAN – Wide area Network

Appendix E: Sample Syslog-ng Configuration File

Sample syslog-ng collection server configuration file. Note that HX defaults to port 6514 for syslog-ng traffic. The config file(s) below uses port 6515 for encrypted TLS transport. The default location for this file in Ubuntu is: `/etc/syslog-ng/syslog-ng.conf`

It is recommended to back up the original configuration file using the following:

```
#> sudo cp /etc/syslog-ng/syslog-ng.conf /etc/syslog-ng/syslog-ng.confg.BAK
```

Here is a sample `syslog-ng.conf` that works for TLS secure shipping. It imports configuration files from `/etc/syslog-ng/conf.d`

```
kaptain@kaptain-syslog:/etc/syslog-ng$ cat syslog-ng.conf
@version: 3.5
@include "scl.conf"
@include "'scl-root`/system/tty10.conf"
options {
    time-reap(30);
    mark-freq(10);
    keep-hostname(yes);
};
source s_local { system(); internal(); };
# source s_network {
#     syslog(transport(tcp) port(6514));
# };
# source tls_source {
#     network(ip(0.0.0.0) port(6515)
#     transport("tls")
#     tls( key-file("/etc/syslog-ng/cert.d/serverkey.pem")
#         cert-file("/etc/syslog-ng/cert.d/servercert.pem")
#         ca-dir("/etc/syslog-ng/ca.d"))
# );};
destination d_local {
file("/var/log/syslog-ng/messages_${HOST}"); };
destination d_logs {
file(
    "/var/log/syslog-ng/logs-enc.txt"
    owner("root")
    group("root")
    perm(0777)
); };
log { source(s_local); destination(d_logs); };
```



```
###
# Include all config files in /etc/syslog-ng/conf.d/
###
@include "/etc/syslog-ng/conf.d/"

kaptain@kaptain-syslog:/etc/syslog-ng/conf.d$ cat audit.conf
## Audit Logging Configuration ###
source demo_tls_src {
    tcp(ip(0.0.0.0) port(6515)
        tls(
            key-file("/etc/syslog-ng/cert.d/serverkey.pem")
            cert-file("/etc/syslog-ng/cert.d/servercert.pem")
            peer-verify(optional-untrusted)
        )
    );
};

filter f_audit_rest { match("hx-audit-rest" value("MSGHDR")); };
filter f_device_conn { match("hx-device-connector" value("MSGHDR")); };
filter f_stssomgr { match("hx-stSSOMgr" value("MSGHDR")); };
filter f_ssl_access { match("hx-ssl-access" value("MSGHDR")); };
filter f_hxmanager { match("hx-manager" value("MSGHDR")); };
filter f_hx_shell { match("hx-shell" value("MSGHDR")); };
filter f_stcli { match("hx-stcli" value("MSGHDR")); };
filter f_hxcli { match("hx-cli" value("MSGHDR")); };

destination d_audit_rest { file("/var/log/syslog-ng/audit_rest.log"); };
destination d_device_conn { file("/var/log/syslog-ng/hx_device_connector.log"); };
destination d_stssomgr { file("/var/log/syslog-ng/stSSOMgr.log"); };
destination d_ssl_access { file("/var/log/syslog-ng/ssl_access.log"); };
destination d_hxmanager { file("/var/log/syslog-ng/hxmanager.log"); };
destination d_hx_shell { file("/var/log/syslog-ng/shell.log"); };
destination d_stcli { file("/var/log/syslog-ng/stcli.log"); };
destination d_hxcli { file("/var/log/syslog-ng/hxcli.log"); };

log { source(demo_tls_src); filter(f_audit_rest); destination(d_audit_rest); flags(final); };
log { source(demo_tls_src); filter(f_device_conn); destination(d_device_conn); flags(final); };
log { source(demo_tls_src); filter(f_stssomgr); destination(d_stssomgr); flags(final); };
log { source(demo_tls_src); filter(f_ssl_access); destination(d_ssl_access); flags(final); };
log { source(demo_tls_src); filter(f_hxmanager); destination(d_hxmanager); flags(final); };
log { source(demo_tls_src); filter(f_hx_shell); destination(d_hx_shell); flags(final); };
log { source(demo_tls_src); filter(f_stcli); destination(d_stcli); flags(final); };
log { source(demo_tls_src); filter(f_hxcli); destination(d_hxcli); flags(final); };

#####
```

It would be possible to use the same system for both TCP and TLS log transport (for example, from 2 different systems). The files above have the TCP part commented out, but if you wanted to configure it you would just create one more file like `audit.conf` in `/etc/syslog-ng/conf.d` and name it something like `audit_tcp.conf` with the configuration as mentioned in the documentation. However, `syslog-ng` will not allow the same property / identifier name as `'demo_tls_src'` (which would be the same in both TCP and TLS configurations above if the file were simply copied over) so it would need to be renamed (e.g., `'demo_tcp_src'`).

Appendix F: Certificate Management and Use Cases

SCVM: How to generate and replace External CA Certificate in HX 4.0.2a+

In 4.0.2a and above, Import CA certificate is automated through a shell script. Generate the CSR from any SCVM preferably from the CIP node. Once you get the CA certificate, import the certificate using the automated script. The script will update the certificate in ZK and push it to other SCVMs.

- Script Location in SCVM: `/usr/share/springpath/storfs-misc/hx-scripts/`
 1. **certificate_import_input.sh**
 2. **updateNginxCertificate.py (not exposed to user)**
- In the Controller VM (Pointing to CIP), execute these commands to generate the CSR request.
 1. **`openssl req -nodes -newkey rsa:2048 -keyout /etc/ssl/private/<Host Name of the CVM>.key -out /etc/ssl/certs/<Host Name of the CVM>.csr`**
 2. **`cat /etc/ssl/certs/<Host Name of the CVM>.csr` - Copy the request to any notepad.**
- Send the request to CA to generate the certificate
- Once you receive the certificate from CA (.crt files), copy the certificate to respective CVM.
- Then use this script to import the certificate: `./certificate_import_input.sh`
 1. `root@SpringpathControllerVUFSTDS58L:/usr/share/springpath/storfs-misc/hx-scripts# ./certificate_import_input.sh`
- Enter the path for the key: `/etc/ssl/private/<Host Name of the CVM>.key`
- Enter the path for the certificate in crt format: `<Path to the CA .crt file>`
- After providing all the inputs, it takes some time to finish the import process
- ***The script prompts to reregister the cluster with vCenter. Its mandatory to reregister the cluster once the certificate is imported.***

The CSR process is updated for 4.5.1a and greater due to the restricted shell.

In the command used above:

```
openssl req -nodes -newkey rsa:2048 -keyout /etc/ssl/private/<Host Name of the CVM>.key -out /etc/ssl/certs/<Host Name of the CVM>.csr
```

While the `openssl` command is part of the allowed admin commands, this command will fail because of lack of write privilege on the `/etc/` folder for the admin user.



Modify the command to use the /tmp/ folder, for example, so that the write proceeds. The command used for 4.5.1a (i.e., any HX version with the secure admin shell):

```
openssl req -nodes -newkey rsa:2048 -keyout /tmp/<Host Name of the CVM>.key -out /tmp/<Host Name of the CVM>.csr
```

This can be completed by changing to the diag user as well. The "diag" user is a user account implemented in 5.0.2a and above. The diag user has a specific subset of permissions to certain files and scripts that can be run under guidance from TAC. This user and any diagnostics performed should only be used when prompted by technical support.

```
admin@hx-shell>su diag
password:
diag@hx-shell>
```

NOTE: The content of the X.509 CSR is entered by the user. There are no backend checks on the contents of the entry. If the user specifies the [multiple] hostnames or IPs of the nodes as subject alternative names, or if they had used the wildcard character to specify the hostname for the Common Name, a single certificate can be used for all nodes.

NOTE: Chrome no longer supports the usage of Common Name and now requires Subject Alternative Names (SAN) to be present in the certificate.

1. Copy your default openssl.cnf file to openssl-san.cnf file
 1. Command: `cp /etc/ssl/openssl.cnf /etc/ssl/openssl-san.cnf`
2. Edit the openssl-san.cnf file to add Subject Alternative Name (SAN) list and required parameters
 1. Uncomment the req-extensions line in the [req] section
Example:
[req]
req_extensions = v3_req
3. Add the SAN lines in the [v3_req] section
Example:
[v3_req]
4. ~lines omitted

```
subjectAltName = @alt_names
```

```
[alt_names]
```

```
DNS.1 = SCVM1
```

```
DNS.2 = SCVM1.example.com
```

```
DNS.3 = SCVM2
```

```
DNS.4 = SCVM2.example.com
```

```
DNS.5 = SCVM3
```

```
DNS.6 = SCVM3.example.com
```

```
DNS.7 = cluster-name
```

```
DNS.8 = cluster-name.example.com
```

Create CSR

1. `openssl req -nodes -newkey rsa:2048 -keyout /etc/ssl/private/<Host Name of the CVM>.key -out /etc/ssl/certs/<Host Name of the CVM>.csr -config openssl-san.cnf`

1. Check CSR SAN entries

1. `openssl req -text -noout -in <yourcsrfile>.csr`

2. `cat /etc/ssl/certs/<Host Name of the CVM>.csr` - Copy the request to any notepad.

3. Send the request to CA to generate the certificate

4. Once you receive the certificate from CA (.crt files), copy the certificate to respective CVM.

1. You may need to convert a CER from Windows CA to crt format

1. `openssl x509 -inform PEM -in <filepath>/certificate.cer -out certificate.crt`

5. Then use this script to import the certificate: **./certificate_import_input.sh**

[/usr/share/springpath/storfs-misc/hx-scripts#](#) ./certificate_import_input.sh

Enter the path for the key: /etc/ssl/private/<Host Name of the CVM>.key

Enter the path for the certificate in crt format: <Path to the CA .crt file>

Have any bundles(y/n)? <In case of bundle, write y, else n>

6. After providing all the inputs, it takes some time to finish the import process

7. Also the script asks to reregister to vCenter. Its mandatory to reregister the cluster once certificate is imported.

HX Certificate Management

vCenter: How to generate and replace External CA Certificate

If ESXi is using a 3rd party CA certificate, certMgmt Mode in vCenter should be set to Custom. The default mode is VMSA. Once the mode is set to Custom, then the hosts with a 3rd party CA can be added.

Follow these steps to update the Mode

- Select the vCenter Server that manages the hosts and click Settings.
- Click Advanced Settings and click Edit.
- In the Filter box, enter certmgmt to display only certificate management keys.
- Change the value of vpxd.certmgmt.mode to custom and click OK.
- Restart the vCenter Server service. To restart services
 1. `https://<VC URL>:5480/ui/services`

Reference:

- <https://docs.vmware.com/en/VMware-vSphere/6.0/com.vmware.vsphere.security.doc/GUID-122A4236-9696-4E1F-B9E8-738855946A93.html#GUID-122A4236-9696-4E1F-B9E8-738855946A93>
- http://engineering.pivotal.io/post/vcenter_6.7_tls/

ESX: How to generate and replace External CA Certificate

Do not replace the CA certificates in all the hosts at the same time. Replace one host and then wait for the cluster to be healthy and then replace the certs for the other nodes.

Follow these steps:

- Generate csr for each ESX Node. Provide proper hostname/FQDN of the ESX host while generating .key and .csr file.
- You should have proper .csr and .key file as part of the key generate procedure i.e. rui.csr and rui.key
- Send rui.csr to 3rd part CA to sign and send back the certificate
- Once you receive the rui.crt from CA
- Put the Node into MM mode
 1. Note: Do not put MM for all the ESX node at a time. You should enter MM in a rolling fashion.
- ssh to the node. Take backup of current rui.key and rui.crt in /etc/vmware/ssl
- Upload the new rui.key and rui.crt to the same directory
- Restart hostd and vpxa service and check status if its running
 1. /etc/init.d/hostd restart
 2. /etc/init.d/vpxa restart
 3. /etc/init.d/hostd status
 4. /etc/init.d/vpxa status
- Reconnect the Host in vCenter
- Exit MM in the host

Follow the same procedure for all other nodes. You can verify the certificate of each node by accessing through web.

Reference:

- <https://kb.vmware.com/s/article/2113926>
- <http://buildvirtual.net/replacing-the-esxi-host-default-certificate-with-a-ca-signed-certificate/>
- <http://buildvirtual.net/how-to-generate-new-esxi-host-certificates/>

ESX: To re-generate self-signed certificate

- SSH to ESX and delete rui.key and rui.crt
- Then /sbin/generate-certificates
- Restart hostd and vpxa

Reference:

https://pubs.vmware.com/vsphere-50/index.jsp?topic=%2Fcom.vmware.vsphere.security.doc_50%2FGUID-EA0587C7-5151-40B4-88F0-C341E6B1F8D0.html

HX Cluster: Re-registration

Once all the hosts are added to the vCenter, reregister the HX Cluster to the vCenter using *stcli cluster reregister* Command

Hyperflex Use Cases

Case: A - Create and Expand HX Cluster with vCenter having CA certificate and certMgmt mode = vmsa (Default Mode)

- Both Create and Expand cluster will work as usual and ESX nodes will be deployed with self-signed certificate
- If user would like to replace the ESX self-signed with CA signed - Then follow these steps
 - Update the certMgmt mode to custom in vCenter
 - Replace self-signed cert in ESX to CA certs
 - Reconnect/add host to cluster
 - Reregister the HX cluster to vCenter

Case: B - Create and Expand HX Cluster with vCenter having CA certificate and certMgmt mode = custom

- Both Create and Expand cluster will fail in deploy stage with error: Not able to add host to vCenter
- Recommendation is:
 - Perform Custom UCSM+Hypervisor configuration as part of Installer workflow
 - After hypervisor configuration is completed, replace the self-signed certificate of ESX Nodes with the CA signed certificates
 - Then perform Deploy+Create Cluster or Deploy+Expand Cluster
 - With this approach add hosts to vCenter will be successful and other steps will work as usual.

Case: C - Replace CA certificate in ESX Nodes in running HX Cluster

- Replace self-signed certificate in vCenter with CA certificate and update the certMgmt mode to custom
- In a rolling fashion, replace the self-signed certificate in ESX Nodes with the CA certificate
- Once replaced for all the nodes, re-register the cluster to vCenter
- In future when trying to expand the cluster
 - Perform UCSM+Hypervisor for the expanded node
 - Replace the self-signed cert in the new node with CA cert
 - Perform Deploy+Expand Cluster

Case: D - Support Workflow - Mgmt IP Change in case ESX Nodes has CA cert

- When changing the hostname/IP addr of ESX as part of support workflow, then the CA certs are getting replaced with self-signed certs by VMware when changing the host name of ESX.
- VMware regenerates self-signed certs with the new hostname.
- Recommendation would be: Again get the CA certs with the new host name and replace in each ESX.
- Steps to be followed to change the host name/IP address of ESX
 - Run the support workflow script to change Mgmt IP
 - It will fail in the step "add/reconnect host to vCenter"
 - Now generate CA cert with the new host name and replace the self-signed certificate
 - Now run the checkpoint of the support workflow to finish updating the remaining tasks

Observations and Notes

vCenter: self-signed cert with certMgmt mode = vmsa (Default Mode)

- ESX with self-signed certificates can be added
- Does not allow adding ESX with 3rd party CA certs
- If already a cluster with self-signed ESX is there in vCenter - After replacing CA certs in ESX,



1. when connecting to vCenter [if certMgmt mode = vmsa], then it asks to replace the CA with self-signed and then allows to add it. Else it will not allow to add.
 2. If certMgmt mode = custom in vCenter, it does not say to replace but it gives error as ssl thumbprint mismatch and add host fails
- Note:
 1. Put the node into MM - replace rui.crt and rui.key in ESX
 2. Restart vpxa and hostd service - CA cert comes in the new node
 3. Right click and connect to vCenter - the CA certs and gets replaced with self-signed by vmware

vCenter: CA signed cert with certMgmt mode = vmsa (Default Mode)

- ESX with self-signed certificates can be added
- Does not allow adding ESX with 3rd party CA certs

vCenter: CA signed cert with certMgmt mode = custom

- ESX with self-signed cannot be added to the vCenter
- ESX needs be replaced with the same CA certificate (CA for the vCenter and CA for the hosts should be same)

Hyperflex Use Cases: HX 4.5.1a+ and NGINX Self-Signed Without CA Signed

- Generate the self-signed certificate if you do not have CA signed certificate
- Use the below REST API to set the custom certificate. The REST API can be invoked from any REST client. Below we are using "Advanced REST Client (ARC)" (<https://install.advancedrestclient.com/install>)
- REST call Construct
 - **URL:** https://<cip_ip>/securityservice/v1/certificate?option=custom
 - **Method Type:** "PUT" request
 - **Authorization:** Basic (Add your username and password)

Method: PUT | Request URL: https://x.x.x.x/securityservice/v1/certificate?option=custom | SEND

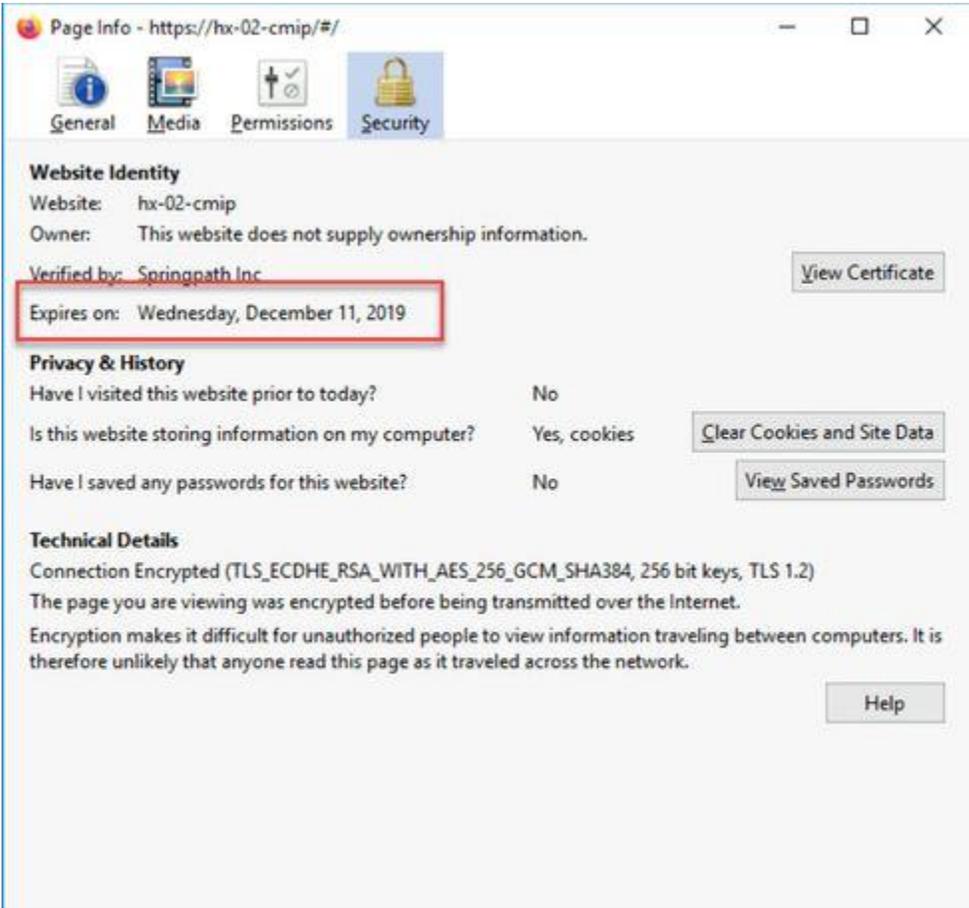
Request parameters ^

HEADERS | BODY | AUTHORIZATION | ACTIONS | CONFIG | CODE

Select authorization: Basic

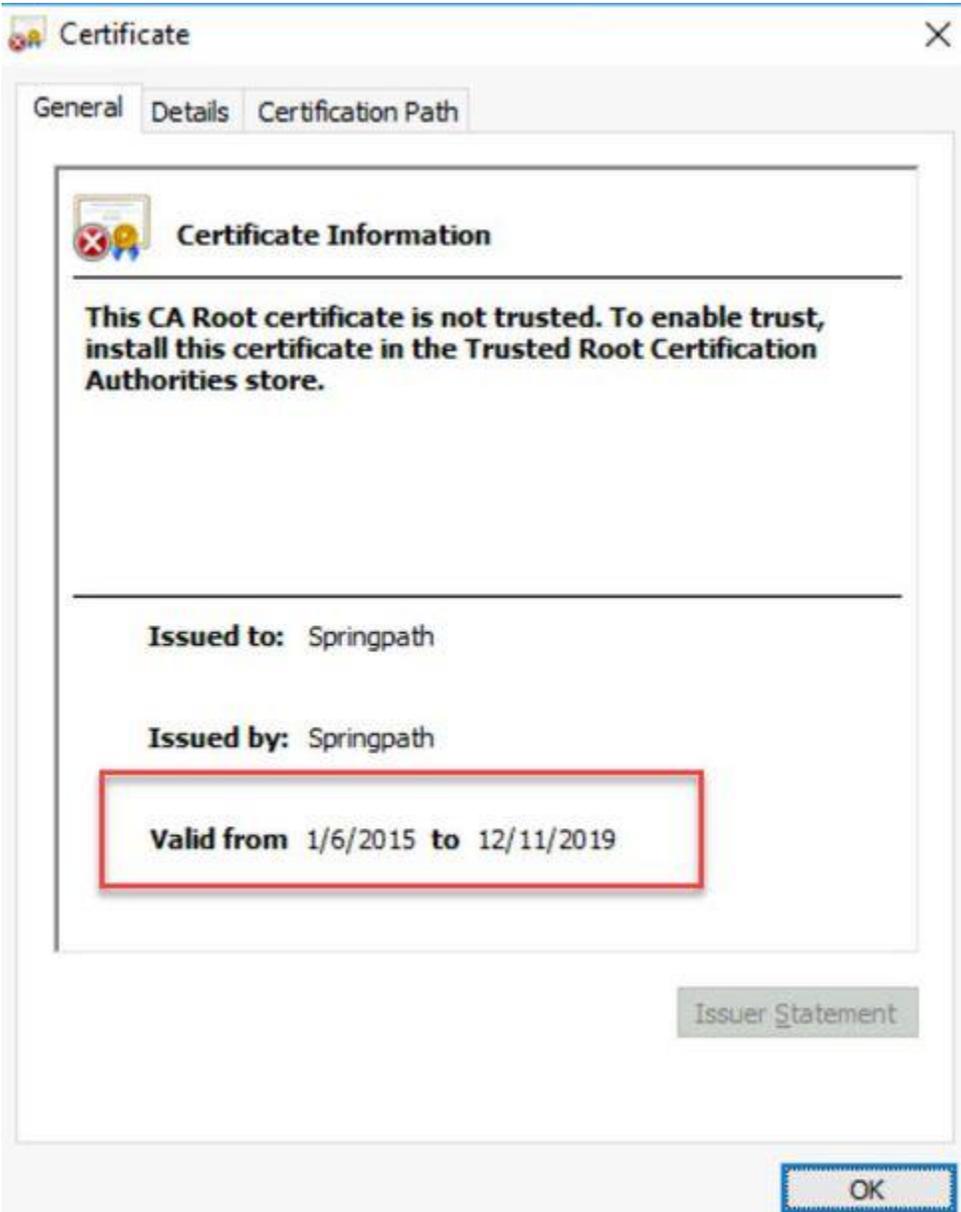
User name: admin

Password: [masked]



The screenshot shows a browser window titled "Page Info - https://hx-02-cmip/#/". The "Security" tab is selected, displaying the following information:

- Website Identity**
 - Website: hx-02-cmip
 - Owner: This website does not supply ownership information.
 - Verified by: Springpath Inc
 - Expires on: Wednesday, December 11, 2019
 - View Certificate
- Privacy & History**
 - Have I visited this website prior to today? No
 - Is this website storing information on my computer? Yes, cookies (Clear Cookies and Site Data)
 - Have I saved any passwords for this website? No (View Saved Passwords)
- Technical Details**
 - Connection Encrypted (TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, 256 bit keys, TLS 1.2)
 - The page you are viewing was encrypted before being transmitted over the Internet.
 - Encryption makes it difficult for unauthorized people to view information traveling between computers. It is therefore unlikely that anyone read this page as it traveled across the network.
 - Help



HX versions with self-signed certificates that are set to expire December 11, 2019 include:

- HX 4.0(1b) and earlier
- HX 3.5(2g) and earlier
- HX 3.5(1a)
- HX 3.0(X)
- HX 2.X
- HX 1.8(X)
- HX 1.7.1

Regenerate and Replace the Certificate



In order to regenerate and replace the certificate, complete these steps:

SSH into the node with the cluster management IP address and enter the required information:

```
root@hx-02-scvm-03:~# openssl req -newkey rsa:2048 -nodes -keyout /etc/nginx/server.key -x509
-days 3650 -out /etc/nginx/server.crt
Generating a 2048 bit RSA private key
.....
.....
writing new private key to '/etc/nginx/server.key'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]: US
State or Province Name (full name) [Some-State]: CA
Locality Name (eg, city) []: San Jose
Organization Name (eg, company) [Internet Widgits Pty Ltd]: Cisco Systems, Inc.
Organizational Unit Name (eg, section) []: HyperFlex
Common Name (e.g. server FQDN or YOUR name) []:<Server_FQDN_Or_IP_Address>
Email Address []: support@cisco.com
root@hx-02-scvm-03:~#
```

After the new certificate and key is generated, copy to all of the storage platform controller virtual machines (SCVMs) in the cluster. This example utilizes Secure Copy (SCP) to copy the certificate files to the other SCVMs:

```
root@hx-02-scvm-03:~# scp /etc/nginx/server.key hx-02-scvm-01.rchs.local:/etc/nginx/server.key
Operating in CiscoSSL FIPS mode
FIPS mode initialized
The authenticity of host 'hx-02-scvm-01.rchs.local (192.168.200.30)' can't be established.
ECDSA key fingerprint is SHA256:OkA9czzcL715fYbfLNtSI+D+Ng5dYp15qk/9C1cQzzk.
Are you sure you want to continue connecting (yes/no)? yes
```



```
Warning: Permanently added 'hx-02-scvm-01.rchs.local,192.168.200.30' (ECDSA) to the list of known hosts.
```

```
HyperFlex StorageController 3.5(2g)
root@hx-02-scvm-01.rchs.local's password:
server.key 100% 1708 1.7KB/s 00:00
```

```
root@hx-02-scvm-03:~# scp /etc/nginx/server.key hx-02-scvm-02.rchs.local:/etc/nginx/server.key
Operating in CiscoSSL FIPS mode
```

```
FIPS mode initialized
```

```
The authenticity of host 'hx-02-scvm-02.rchs.local (192.168.200.31)' can't be established.
ECDSA key fingerprint is SHA256:OkA9czzcL715fYbfLNtSI+D+Ng5dYp15qk/9C1cQzzk.
```

```
Are you sure you want to continue connecting (yes/no)? yes
```

```
Warning: Permanently added 'hx-02-scvm-02.rchs.local,192.168.200.31' (ECDSA) to the list of known hosts.
```

```
HyperFlex StorageController 3.5(2g)
root@hx-02-scvm-02.rchs.local's password:
server.key 100% 1708 1.7KB/s 00:00
```

```
root@hx-02-scvm-03:~# scp /etc/nginx/server.crt hx-02-scvm-01.rchs.local:/etc/nginx/server.crt
Operating in CiscoSSL FIPS mode
```

```
FIPS mode initialized
```

```
HyperFlex StorageController 3.5(2g)
root@hx-02-scvm-01.rchs.local's password:
server.crt 100% 1371 1.3KB/s 00:00
```

```
root@hx-02-scvm-03:~# scp /etc/nginx/server.crt hx-02-scvm-02.rchs.local:/etc/nginx/server.crt
Operating in CiscoSSL FIPS mode
```

```
FIPS mode initialized
```

```
HyperFlex StorageController 3.5(2g)
root@hx-02-scvm-02.rchs.local's password:
server.crt 100% 1371 1.3KB/s 00:00
```

Once the .crt and .key files are in place, restart Nginx on all SCVMs:

```
root@hx-02-scvm-03:~# service nginx restart
* Restarting nginx nginx [ OK ]
```

```
root@hx-02-scvm-02:~# service nginx restart  
* Restarting nginx nginx
```

```
root@hx-02-scvm-01:~# service nginx restart  
* Restarting nginx nginx
```

The certificate will now be valid for 10 more years:



If your HX cluster runs on VMware, you will need to reregister the cluster to vCenter. This command is used to reregister the cluster:

```
root@hx-02-scvm-03:~# stcli cluster reregister --vcenter-datacenter <vc_datacenter_name> --vcenter-cluster <vc_cluster_name> --vcenter-url <VCSA_FQDN_OR_IP> --vcenter-user administrator@vsphere.local
```

Appendix G: SCH Configuration and Proxy

Configuring Smart Call Home for Data Collection

Data collection is enabled by default but, during installation, you can opt-out (disable). You can also enable data collection post cluster creation. During an upgrade, Smart Call Home is set up based on your legacy configuration. For example, if `stcli services asup show` is enabled, Smart Call Home is enabled on upgrade.

Data collection about your HX storage cluster is forwarded to Cisco TAC through https. If you have a firewall installed, configuring a proxy server for Smart Call Home is completed post cluster creation. Using Smart Call Home requires the following:

- A Cisco.com ID associated with a corresponding Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service contract for your company.
- Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service for the device to be registered.

Procedure

Step 1 Log in to a storage controller VM in your HX storage cluster.

Step 2 Register your HX storage cluster with Support.

Registering your HX storage cluster adds identification to the collected data and automatically enables Smart Call Home. To register your HX storage cluster, you need to specify an email address. After registration, this email address receives support notifications whenever there is an issue and a TAC service request is generated.

Note Upon configuring Smart Call Home in Hyperflex, an email will be sent to the configured address containing a link to complete registration. If this step is not completed, the device will remain in an inactive state and an automatic Service Request will not be opened.

Syntax:

```
stcli services sch set [-h] --email EMAILADDRESS
```

Example:

```
# stcli services sch set --email name@company.com
```

Step 3 Verify data flow from your HX storage cluster to Support is operational.

Operational data flow ensures that pertinent information is readily available to help Support troubleshoot any issues that might arise.

--all option runs the commands on all the nodes in the HX cluster.

asupcli [--all] ping

If you upgraded your HX storage cluster from HyperFlex 1.7.1 to 2.1.1b, also run the following command:

asupcli [--all] post --type alert

Contact Support if you receive the following error:

```
root@ucs-stctlvm-554-1:/tmp# asupcli post --type alert
```

```
/bin/sh: 1: ansible: not found
```

```
Failed to post - not enough arguments for format string
```

```
root@ucs-stctlvm-554-1:/tmp#
```

Step 4 (Optional) Configure a proxy server to enable Smart Call Home access through port 443.

If your HX storage cluster is behind a firewall, after cluster creation, you must configure the Smart Call Home proxy server. Support collects data at the url: <https://diag.hyperflex.io:443> endpoint.

- a. Clear any existing registration email and proxy settings.

stcli services sch clear

- b. Set the proxy and registration email.

Syntax:

stcli services sch set [-h] --email EMAILADDRESS [--proxy-url PROXYURL] [--proxy-port PROXYPORT] [--proxy-user PROXYUSER] [--portal-url PORTALURL] [--enable-proxy ENABLEPROXY]

Syntax Description

Option	Required or Optional	Description
--email EMAILADDRESS	Required.	Add an email address for someone to receive email from Cisco support. Recommendation is to use a distribution list or alias.
--enable-proxy ENABLEPROXY	Optional.	Explicitly enable or disable use of proxy.
--portal-url PORTALURL	Optional.	Specify an alternative Smart Call Home portal URL, if applicable.
--proxy-url PROXYURL	Optional.	Specify the HTTP proxy URL, if applicable.

--proxy-port PROXYPORT	Optional.	Specify the HTTP proxy port, if applicable.
--proxy-user PROXYUSER	Optional.	Specify the HTTP proxy user, if applicable. Specify the HTTP proxy password, when prompted.

Example:

```
# stcli services sch set

--email name@company.com

--proxy-url www.company.com

--proxy-port 443

--proxy-user admin

--proxy-password adminpassword
```

- c. Ping to verify the proxy server is working and data can flow from your HX storage cluster to the Support location.

```
# asupcli [--all] ping

--all option runs the command on all the nodes in the HX cluster.
```

Step 5 Verify Smart Call Home is enabled.
When Smart Call Home configuration is set, it is automatically enabled.

```
# stcli services sch show

If Smart Call Home is disabled, enable it manually.
```

```
# stcli services sch enable
```

Step 6 Enable Auto Support (ASUP) notifications.

Typically, Auto Support (ASUP) is configured during HX storage cluster creation. If it was not, you can enable it post cluster creation using HX Connect or CLI. For more information, see [Auto Support and Smart Call Home for HyperFlex](#).

Appendix H: Changing the SSO Timeout for vCenter

Some HyperFlex customers are using Two Factor Authentication (2FA) on their vCenter(s). End users need sufficient time to approve the 2FA request before the create session timeout occurs. The current create session timeout is set to 10 seconds. This is too short for most people to accept the 2FA request before the timeout expires.

From /var/log/springpath/hxmanager.log, we see the **Completed 401 Unauthorized in 13.569296805s**.

```
2022-12-01-17:20:32.207 Started POST /hx/api/auth
2022-12-01-17:20:32.207 [opID=14138cca7b9a35] Request URL POST
https://localhost/aaa/v1/auth?grant_type=password
2022-12-01-17:20:45.776 [opID=14138cca7b9a35] Got response 401 in 13.56910124s for POST
https://localhost/aaa/v1/auth?grant_type=password
2022-12-01-17:20:45.776 [opID=14138cca7b9a35] Error code 401 | <nil> |
/aaa/v1/auth?grant_type=password
2022-12-01-17:20:45.776 [error] 401 | {"code":100005,"message":"Authentication failed or account
locked out. Please check your user name and password and verify that the system configuration
adheres to prerequisites"} | /aaa/v1/auth?grant_type=password
2022-12-01-17:20:45.776 Completed 401 Unauthorized in 13.569296805s
```

To adjust the authentication timeout setting, log into each SCVM as root and update the following file and parameter:
/var/lib/tomcat8/webapps/aaa/WEB-INF/classes/application.conf :

```
timeouts {
    CreateSessionTimeoutInSeconds = 10 << updated to 60
},
```



Appendix I: Password Recovery

Starting with HyperFlex 4.5(1a), the Secure Shell feature was added. To gain root access TAC support is required for root token generation.

Password recovery for 4.5(1a), 4.5(2a)

- SSH as root to any ESXi host.
- Use the following command to SSH to the SCVM that resides on this ESXi host using the certificate over the Storage Controller Data Network. Hit Y to accept.

```
[root@hx-03-esxi-01:/opt/cisco/support] ssh root@`/opt/cisco/support/getstctlvmpip.sh "Storage Controller Data Network" -i /etc/ssh/ssh_host_rsa_key
```

HyperFlex StorageController 4.5(1a)

!!! ALERT !!!
This service is restricted to authorized users only.
All activities on this system are logged. Unauthorized access will be reported.

HyperFlex StorageController 4.5(1a)
Last login: Mon Mar 22 17:10:19 2021 from 192.168.150.67
WARNING: By accepting this support session, you give your consent and hereby authorize Cisco to have privileged access to the supported Cisco device for the purpose of providing technical support. At the conclusion of this session you must exit root shell from all the open ssh sessions of all the controller vms of the cluster and invalidate the consent token in order to terminate Cisco's access and close the privileged access portal. You are hereby advised that failure to do so may create a vulnerability in your product.
Accept(Y/n): y

3. Enter 1 to Generate Challenge For root Shell Access. In the example above, the time period for root access is set to 4320.

Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
Enter CLI Option:
1
Enter time period in minutes for root shell access(max 4320 mins): 4320
Generating Challenge.....
Challenge String (Please copy everything between the asterisk lines exclusively):
*****BEGIN TOKEN*****



```
D9J8cAAAAQEBAQAQAAAABAgAEAAAAAAMACEoGYXnObY4CBAAQASmygYy6+pRygOAjgA/5R
AUABAAAEEOAGAAIleXBICmZsZXgHAAXleXBICmZsZXhfQ1QIAAIIWVBFUKZMRVgJACBkNjYzOTM2ZG
Q5MDIxMjAwODNhZWNmNGQ4OGEyNjIhZA==
*****END TOKEN*****
```

4. Use the TAC generated response key.

Consent token is needed to access root shell !!

1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit

Enter CLI Option:

2

Starting background timer of 30 mins

Please input the response when you are ready

```
M8affwAAAQEBAQAQAAAABAgAEAAAAAAMBYnJVNWpLNmU1bUZ4aFRnQStBTXBKeTUwQWhINm
1FNFmzcFgwb05CM29ZNXN5SWt2L25oOXdfbldmekpKRysxZHgNCm8xMUNHWWJZNmw5cEI1U01
GUFpBYTivSDhKdHdZV2hnd2VwbGc4VfItWkIMQi9iOFBZL3ZjajBqZTczSzZUSDUNCnIJSzA4UkdDSzcw
RDJ6R2tKRTM3SURZeVV1UjROMUIXUEY0OTNrOThyeXhQOEhPUDhBdXc3cGVSNEZQS3FsTWkNckI2d
mQ1ckVLNIRHaW85ZExGU0pNK05RUVJQUzFBK1M0U1VxRU1rdmlBaVc1ejIPNnZUeFJKMDY0OXZrVjJ
YdDINcmVOamZLQTJ5NEowQzB4L2h4WDFxemhvZXRzUFBRWVYvc2xMR0d2VWp3Z05YL3grak5KW
GI2SDdOUEk3djVTek0NCmt0RkZabWJ6WWVub2ZiOGwxN1ZZYWc9PQ==
```

Response Signature Verified successfully !

Response processed successfully.

5. Make sure to enter "n" when asked "Do you want to sync the consent token to other controller VMs in the cluster(y/n)?" since we don't know the root password.

Response Signature Verified successfully !

Response processed successfully.

Do you want to sync the consent token to other controller VMs in the cluster(y/n)? : n <<< Must hit No since we don't know the root password

Other controller VMs will need new token generation to access root shell.

Run 'ct_engine --sync' to sync the token.

Providing root shell for this controller VM for 4320 minutes!!

bash-4.2#

6. Now, change the root password locally using passwd.

```
bash-4.2# passwd <<< change the password locally
```

New password:

Retype new password:

passwd: password updated successfully

7. Next, set the root password globally using stcli security password set --user root.



```
bash-4.2# stcli security password set --user root <<< Change the password globally
Please enter the current password for user root:
Enter new password for user root:
Re-enter new password for user root:
```

8. Lastly, now that we have the root password, exit the secure shell session and invalidate the token.

```
bash-4.2# exit
exit
Do you also want to invalidate consent token(y/n)? : y
To invalidate token across the cluster, please enter root password:
Token invalidated !!
Connection to 192.168.150.73 closed.
```

Now you can use the normal process of gaining secure shell access for this version.

Password recovery for 4.5(2b) and 5.0(1x)

The steps above do not work starting with 4.5(2b). Please contact TAC for this procedure.

Admin password recovery for 4.5(2c) and 5.0(2a) onward

For ESXi version >= 7.0, use the following command to access hxshell:

```
root@bsv-hx-7-A-1:~] ssh -o PubkeyAcceptedKeyTypes=+ssh-rsa admin@`/opt/cisco/support/get
stctlvmpip.sh` -i /etc/ssh/ssh_host_rsa_key
HyperFlex StorageController 5.0(2b)
hostfile_replace_entries: link /.ssh/known_hosts to /.ssh/known_hosts.old: Function not implement
ed
update_known_hosts: hostfile_replace_entries failed for /.ssh/known_hosts: Function not impleme
nted
This is a Restricted shell.
Type '?' or 'help' to get the list of allowed commands.
hxshell:~$
```

After getting access, proceed with the recover-password command as explained below

For ESXi version < 7.0, use the following command

```
[root@hx-08-esxi-01:~] ssh admin@`/opt/cisco/support/getstctlvmpip.sh "Storage Controller Data
Network" -i /etc/ssh/ssh_host_rsa_key
The authenticity of host '192.168.150.232 (192.168.150.232)' can't be established.
ECDSA key fingerprint is SHA256:OkA9czzcL7l5fYbfLNtSI+D+Ng5dYp15qk/9C1cQzzk.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '192.168.150.232' (ECDSA) to the list of known hosts.
```



HyperFlex StorageController 5.0(2a)

!!! ALERT !!!

This service is restricted to authorized users only.
All activities on this system are logged. Unauthorized access will be reported.

HyperFlex StorageController 5.0(2a)

Last login: Tue Aug 30 11:54:10 2022 from 10.201.254.84

This is a Restricted shell.

Type '?' or 'help' to get the list of allowed commands.

hxshell:~\$ recover-password

Consent token is needed to reset password. Do you want to continue?(y/[n]):

y

1. Generate Challenge
2. Accept Response
3. Exit

Enter Option:

1

Generate the response key and then:

Response Signature Verified successfully !

Response processed successfully.

Consent token workflow is successful, allowing password reset.

Enter the new password for admin:

Re-enter the new password for admin:

Changing password for admin...

Password changed successfully for user admin.

Sync the password globally with :

admin:~\$ stcli security password set

Enter new password for user admin:

Re-enter new password for user admin:

admin:~\$



Appendix J: Persistent Root Air-Gapped Revocation

Air-gapped systems have a SLR/PLR type license. There is a specific process to revoke persistent root on these platforms because of their limited or non-existent ability to get online.

With air-gapped clusters, the method is to run an stcli command to "return" the license reservation, which would trigger the controller VM to automatically revoke the long-lived TAC support token that enables persistent root. Then the cluster can be re-licensed by running another stcli command to re-apply the license reservation.

Perform the following procedure once per cluster from the node hosting the cluster management IP.

The license show status command will show the system as 'Registered':

```
hxshell:~$ stcli license show status
```

```
Smart Licensing is ENABLED
```

```
License Reservation is ENABLED
```

```
Registration:
```

```
Status: REGISTERED - UNIVERSAL LICENSE RESERVATION
```

Return and reinstall the license by following these steps:

Returning Specific License Reservation (SLR) Licenses

https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/Installation_VMware_ESXi/4_0/b_HyperFlexSystems_Installation_Guide_for_VMware_ESXi_4_0/b_HyperFlexSystems_Installation_Guide_for_VMware_ESXi_4_0_chapter_0110.html#id_125881

Run 'stcli license reservation return' this will give a return code. This code is used to remove the license instance from CSSM as mentioned in above link. After this the license will need to be reinstalled per the 'Installing Specific License Reservation (SLR) Licenses' section in above link.

NOTE: The guide is for SLR, but the same steps apply for PLR licenses. The only difference is during generation of the authorization code from CSSM during the license install, the selection of license has to be type PLR.