White paper

Cisco public

The bridge to possible

Cisco HX Data Platform Edge Encryption Using ESXi VMCrypt

Version 2.1

Contents

Document information

Table 1. Document Summary

Document summary		Prepared for	Prepared by
Cisco HX with VMCrypt	v.2.1	Field	Aaron Kapacinskas
Last Modified:	20 November 2020		
Previous Version:	1.0		
Changes in this version:	Fixed misc. typos. Added Bias Statement		

Intended use and audience

This document contains original configuration and test-result content along with material from various Cisco® and VMware publications. The materials, ideas, and concepts contained herein are to be used exclusively to assist in the configuration of Cisco software solutions. It is intended for Cisco customers and employees who need to understand and implement software-based encryption solutions from VMware on Cisco HyperFlex™ Edge platforms.

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.

Prerequisites

We recommend reviewing the release notes, installation guide, and user guide before proceeding with any configuration. The Cisco HyperFlex™ HX Data Platform (HXDP) should be installed and functioning per the installation guide. Please contact your Cisco representative or the Cisco Technical Assistance Center (TAC) if assistance is needed.

Introduction

A Cisco HyperFlex Data Platform Edge cluster is targeted for smaller, cost-sensitive deployments. It is ideal for branch offices or remote offices. A key difference between Edge clusters and standard Cisco HyperFlex HX Data Platform (HXDP) clusters is that Edge operates without Fabric Interconnects and runs using 1GE or 10GE infrastructure and can be sized from 2 to 4 nodes (starting with Cisco HXDP Release 4.0). Self-Encrypting Drives (SEDs) are not supported on this platform, so, as a result, software-based encryption is required for environments that demand data-at-rest encryption. Beginning with VMware vSphere Hypervisor (ESXi) 6.5, general datastore native encryption is supported by VMware. This document provides guidance for HyperFlex (HX) Edge users in ensuring that their product is deployed in a more secure manner using a key manager and ESXi-native encryption. This document also provides some insight into the performance characteristics of Edge platforms using HX Bench, a free performance benchmarking tool available from Cisco, to compare unencrypted and encrypted VM workloads.

It is necessary to understand the architecture and components of the solution in order to enable data-at-rest encryption properly. This document provides recommended configuration settings and deployment architectures for HXDP Edge solutions using this paradigm. 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 reference Cisco.com.

Cisco HyperFlex (Cisco HX) Edge architecture

The following subsections provide a brief overview of the Cisco HyperFlex Edge system architecture, starting with a general node description. It is followed by a brief discussion of the networking requirements. It is critical to the secure environment that the various parts are hardened as needed.

Typical Cisco 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 Cisco HX Edge configurations with the add-on Intel® PCIe quad port NIC, ensure that any unused Ethernet ports remain disconnected from any virtual switches in ESXi. This will prevent unauthorized access to the virtual switching environment. In general, any unused Ethernet ports should remain disconnected from the virtual networking stack.

SED deployments are currently not supported with Cisco HX Edge. VM encryption by virtue of third-party encryption clients will work to encrypt the VMs deployed on Cisco HX Edge. VMware vSphere 6.5 incorporates a VM encryption capability called VMCrypt. VMCrypt in conjunction with a Key Manager from Thales Vormetric (DSM) is the focus of this paper. Other key managers may be used, subject to VMware’s Compatibility Guide.

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 Cisco HyperFlex HX Data Platform (HXDP) that is deployed on each node as part of cluster installation. The HXDP utilizes PCI pass-through, which removes storage (hardware) operations from the hypervisor; this makes the system highly performant. The HX nodes use special plug-ins for VMware called VIBs that are used to redirect NFS datastore traffic to the correct distributed resource, and for hardware offloading of complex operations such as snapshots and cloning.

The following illustration shows the typical components found within an HX node.

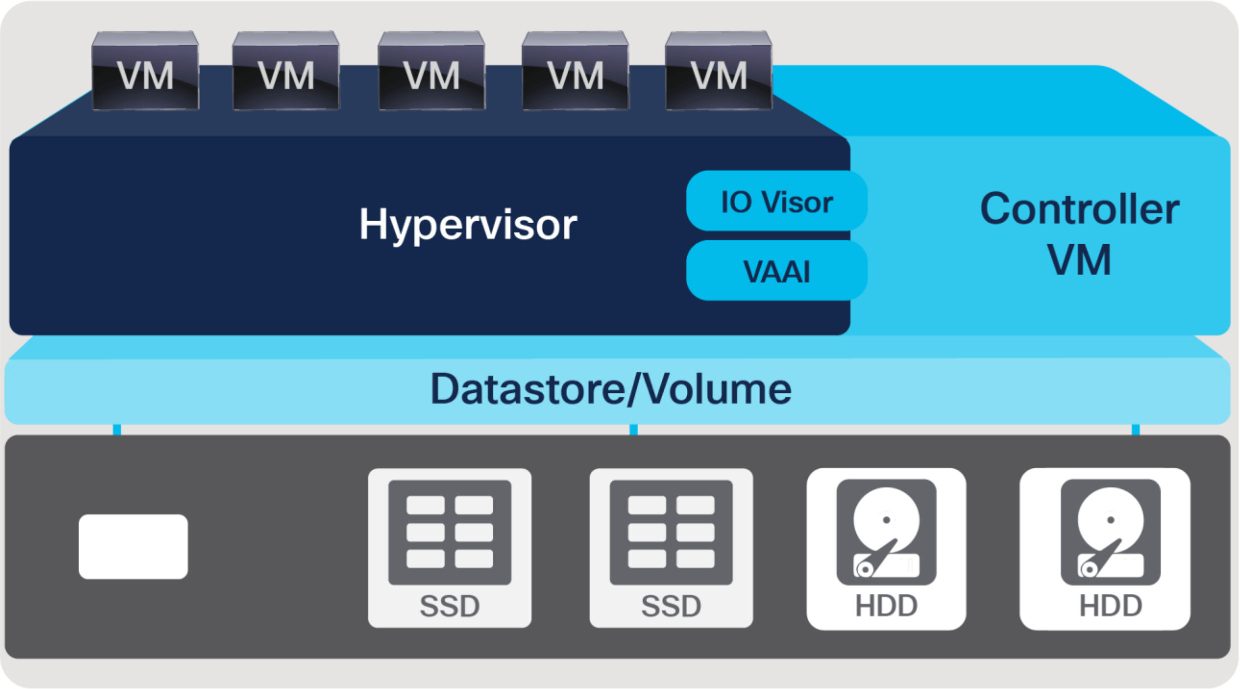

Figure 1.
HyperFlex Node Architecture

These nodes are incorporated into a distributed HyperFlex Edge cluster, as shown below. Note the conspicuous lack of Fabric Interconnects (FIs) in Edge that are otherwise present in a traditional HyperFlex cluster. The nodes themselves are directly connected to top-of-rack switches based on the user-selected topology.

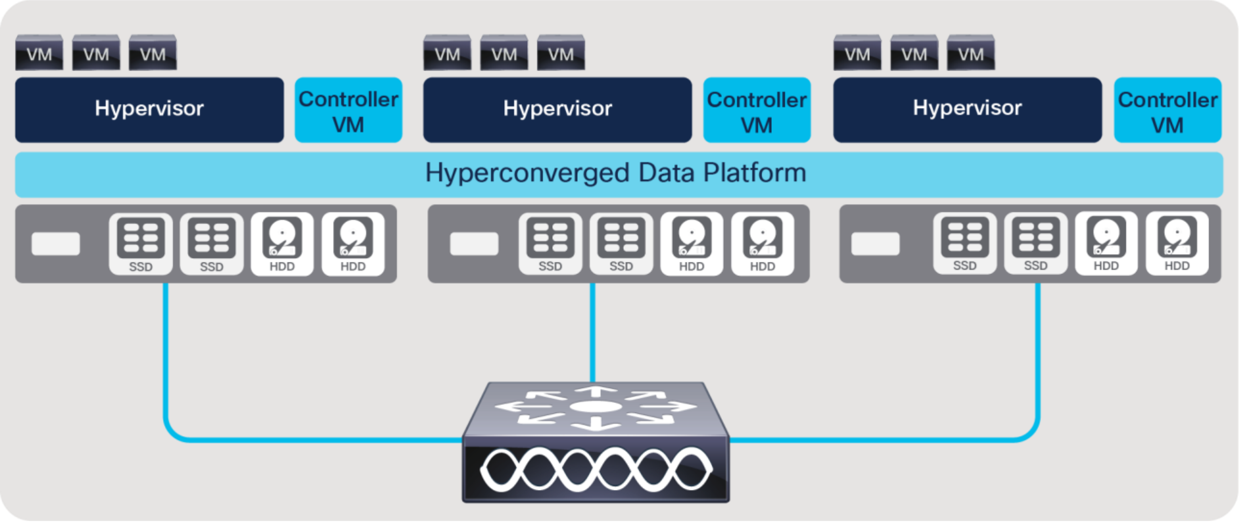

Figure 2.
HyperFlex Cluster Architectue

Cisco UCS

The physical HX Edge node is deployed on a Cisco UCS® 220 platform in either a hybrid or an all-flash configuration. A physical server configuration for the HX nodes is applied during cluster build at install time. This is performed via communications to the Cisco® Integrated Management Controller (IMC) for Edge nodes. All nodes are auto-configured by the HX Installer with the settings required for HX to operate securely and efficiently (BIOS settings, boot order, VLANs, MAC addresses, IP addressing, ESXi hypervisor configuration, etc.).

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

HX Connect is the native HTML 5.0 management interface for the HyperFlex cluster. The HX vCenter plug-in is another management interface available in the vCenter vSphere Web Client once the cluster is deployed. These are separate interfaces. Both are accessed via HTTPS in a web browser, with HX Connect providing user management (including RBAC).

VMware vCenter

The Cisco HyperFlex HX Data Platform requires VMware vCenter to be deployed to manage certain aspects of cluster creation, such as ESXi 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. VMware vCenter also serves as the KMS client for key exchange in the implementation of VMware ESXi–native encryption (VMCrypt). It is important to note that all nodes must share a single vCenter cluster object for a given cluster. This 1:1 mapping is a requirement.

VMware vSphere Hypervisor (ESXi)

VMware vSphere Hypervisor (ESXi) is the hypervisor of choice in this paper. It abstracts node compute and memory hardware for the guest VMs. HXDP integrates closely with ESXi to facilitate network and storage virtualization.

Networking

The HX networking environment is segmented and isolated to provide out-of-the-box traffic security. This section describes the traffic paths, vSwitch architecture, and VLAN best practices.

A HyperFlex Edge cluster does not use Cisco UCS Fabric Interconnects (FIs). There are two deployment models for 1GE operation. The first uses a single upstream Top-of-Rack (ToR) switch. The second deployment model uses a redundant set of two ToR switches for added availability. A 10GE option is also available. The tests conducted in this paper were performed using a single 1GE switch, although the 10GE is a preferred option for customers due to the higher performance achievable. See the References section for a link to the deployment guide for additional information on HX Edge network topologies.

East-west traffic and north-south traffic

East-west traffic describes the node-to-node communication that occurs within a cluster. On a standard HX cluster, this takes place through the Fabric Interconnects (FIs). The north-south traffic in a standard cluster takes place between the nodes and the end users by traversing the FIs and then exiting the cluster via the upstream switch(es). In an Edge cluster, the FIs are not present, so all traffic must traverse the interconnecting physical switch. As a result, east-west and north-south traffic have the same path with the only differentiator being (potentially) the number of L2 hops required to satisfy the network request.

VLANs and vSwitches

Use a separate subnet and VLAN for each of the required networks. Do not use VLAN 1, the default VLAN, because it can cause networking issues, especially if Disjoint Layer 2 configuration is used. Use a numbered VLAN higher than 1.

VLANs are created for each type of traffic and for each vSwitch. There are typically four vSwitches created during a standard install, with associated VLANs for each. For Cisco HyperFlex Edge, two vSwitches are created at build time. The vSwitches are reserved and configured for:

- Cisco HyperFlex and VMware ESXi management and guest VM network(s)
- HX Data (storage traffic between nodes for the distributed data platform) and ESXi vMotion traffic

The zones that these vSwitches 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. This zone contains the following components:
    - ESXi host management interfaces
    - Storage controller VM management interfaces
    - A roaming HX cluster management interface (The cluster management IP is called the CIP-M)
- **VM zone:** This zone comprises the connections needed to service network I/O to the guest VMs that will run inside the HyperFlex hyperconverged system. This zone typically contains multiple VLANs. These interfaces and IP addresses need to be available to all staff and other computer endpoints that 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 HyperFlex HX Data Platform software, VMware ESXi hosts, and the storage controller VMs to service the HX distributed file system. These interfaces and IP addresses need to be able to communicate with each other at all times for proper operation. This zone contains 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 (called the cluster data IP or just the CIP)
    - All of these interfaces are typically placed on a non-routed subnet and are there not exposed to end users or any devices outside the above list
- **vMotion zone:** This zone comprises the connections used by the ESXi hosts to enable vMotion of the guest VMs from host to host

These vSwitches and their associated port groups are tied to one or two physical NIC uplinks on the Edge nodes, depending on the topology choice. Each node contains a vSwitch configuration similar to the following:

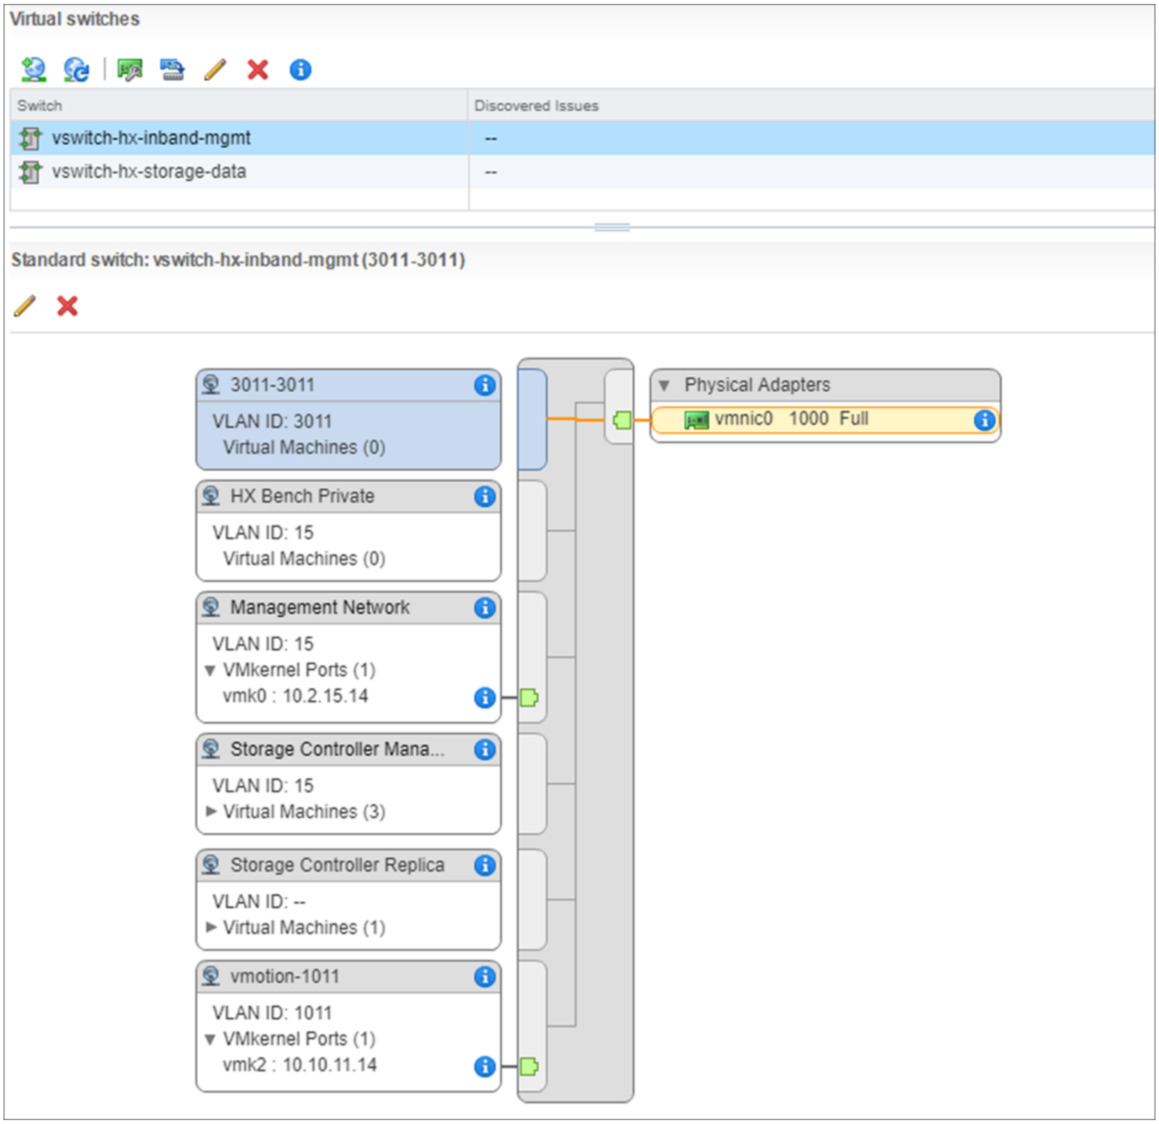

Virtual switches

Switch	Discovered Issues
vswitch-hx-inband-mgmt	--
vswitch-hx-storage-data	--

Standard switch: vswitch-hx-inband-mgmt (3011-3011)

3011-3011
VLAN ID: 3011
Virtual Machines (0)

HX Bench Private
VLAN ID: 15
Virtual Machines (0)

Management Network
VLAN ID: 15
▼ VMkernel Ports (1)
vmk0 : 10.2.15.14

Storage Controller Mana...
VLAN ID: 15
▶ Virtual Machines (3)

Storage Controller Replica
VLAN ID: --
▶ Virtual Machines (1)

vmotion-1011
VLAN ID: 1011
▼ VMkernel Ports (1)
vmk2 : 10.10.11.14

▼ Physical Adapters
vmnic0 1000 Full

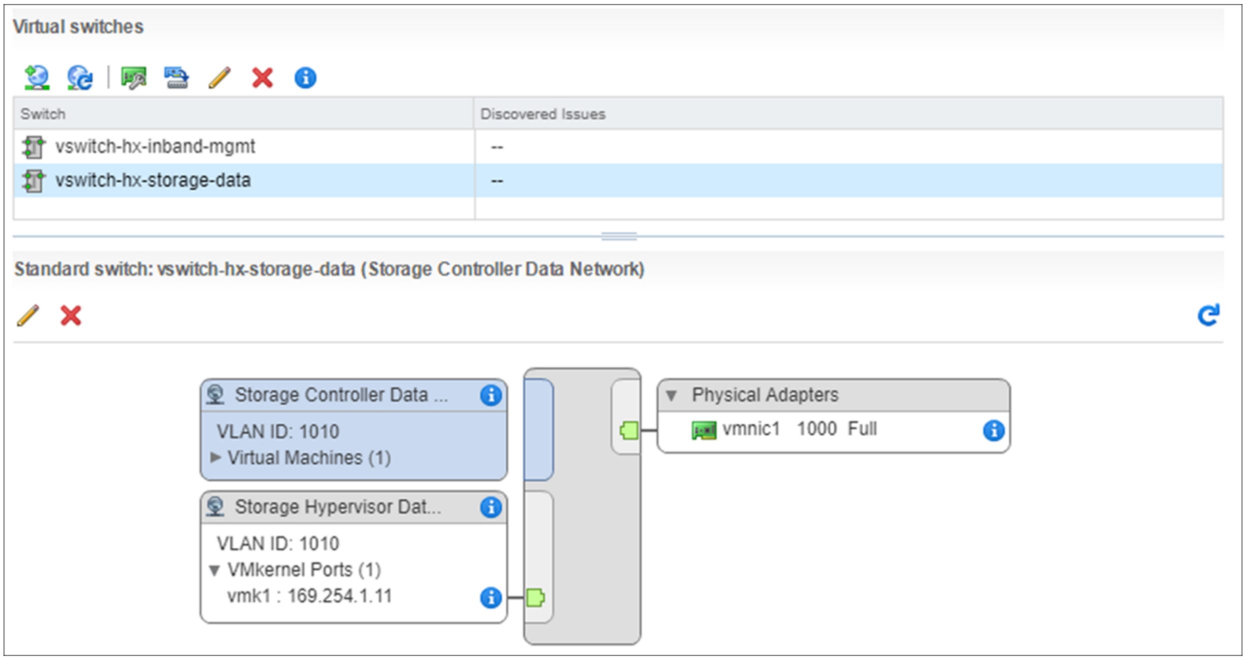

Virtual switches

Switch	Discovered Issues
vswitch-hx-inband-mgmt	--
vswitch-hx-storage-data	--

Standard switch: vswitch-hx-storage-data (Storage Controller Data Network)

Storage Controller Data ...
VLAN ID: 1010
▶ Virtual Machines (1)

Storage Hypervisor Dat...
VLAN ID: 1010
▼ VMkernel Ports (1)
vmk1 : 169.254.1.11

▼ Physical Adapters
vmnic1 1000 Full

Upstream switch

Each vSwitch in the architecture has one or more VLANs associated with it. For traffic leaving the cluster, you must configure the upstream switches to accommodate non-native VLANs. The HX installer requires all switchports to be tagged or configured in trunked mode.

HX data security

Secure communications

All communication occurring with the HX platform management interfaces is FIPS compliant using SSH or HTTPS. See the References section for a link to the HyperFlex Hardening Guide for an in-depth discussion of these points.

Encryption on HX

HX data-at-rest security in a standard cluster is accomplished via Secure Encrypted Disks (SEDs) and is managed by Cisco HX Connect in conjunction with Cisco UCS Manager (UCSM) and local or remote key stores using the Key Management Interoperability Protocol v1.1. In a HyperFlex Edge system, the lack of FIs also means that there is no UCSM. As a result, SEDs are not supported. Therefore, in order to encrypt the system VMs, a software solution is utilized to achieve the same result.

In a software solution with HyperFlex Edge, a third-party external key manager is required. Local key management is not supported since the cluster lacks facilities to manage these keys. A typical Edge deployment with an external, remote key manager is shown below.

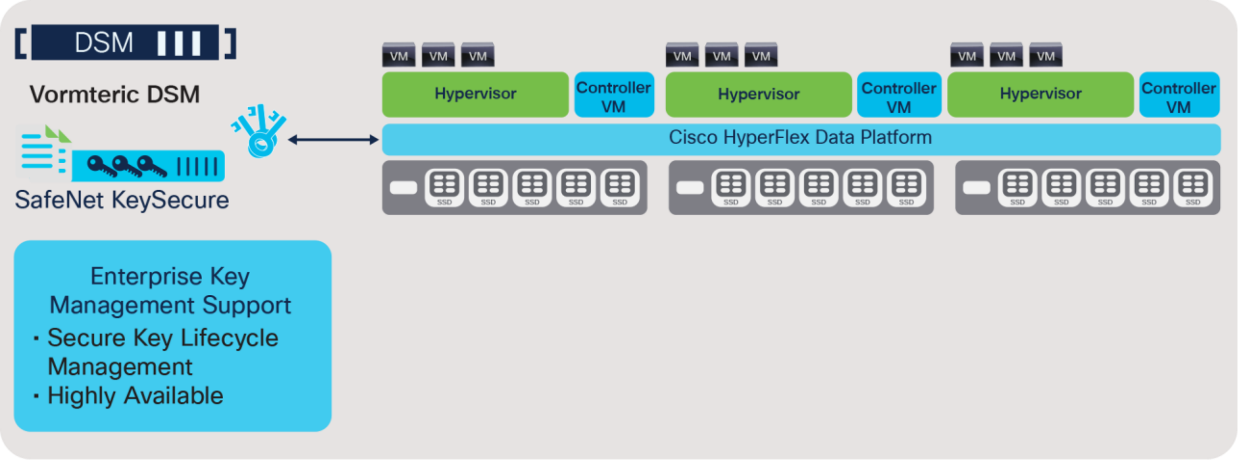
DSM
Vormteric DSM
SafeNet KeySecure
Enterprise Key Management Support
· Secure Key Lifecycle Management
· Highly Available
VM
VM
VM
Hypervisor
Controller VM
Cisco HyperFlex Data Platform
SSD

Figure 3.
Key Exchange with Vormetric DSM and a HyperFlex Cluster

- Data is only as secure as the encryption keys
- Key management involves protecting, storing, backing up and organizing keys
- Specialized vendors provide enterprise key management offerings

Key management

Key Management Interoperability Protocol (KMIP) is the communication standard for remote Key Management Servers (KMS). KMIP server key handling for HyperFlex is performed through encryption partners. These same partners will work with VMCrypt for key management. The KMIP server configuration is added to vCenter for VMware ESXi–native encryption (VMCrypt).

Key management best practices:

- Always deploy at least two KMIP servers, clustered for high availability
- Configure key backup and recovery
- Self-signed and Certificate Authority (CA)-signed certificates can be used

The advantages of a key management solution:

- Provides a single, centralized management plane for cryptographic keys and applications
- Offers high-availability and standards-based enterprise encryption key management using KMIP
- Centralizes third-party encryption keys and securely stores certificates
- Enables vaulting and an inventory of certificates
- Implements a two-factor authentication mechanism to further safeguard keys and certificates against theft

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 KMS UI and Vormetric DSM configurations. Additionally, port 5696 is required for TLS communication between vCenter and the KMS server itself for secure information exchange. See the link to the HyperFlex Hardening Guide in the References section for a comprehensive list of HX ports in use.

Encryption partners

Cisco HyperFlex partners with two industry-leading encryption and KMIP service providers, given below:

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:

- Vormetric Data Security Manager (DSM) 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 should work but require qualification. Refer to the VMware Compatibility Guide for the list of supported key managers.

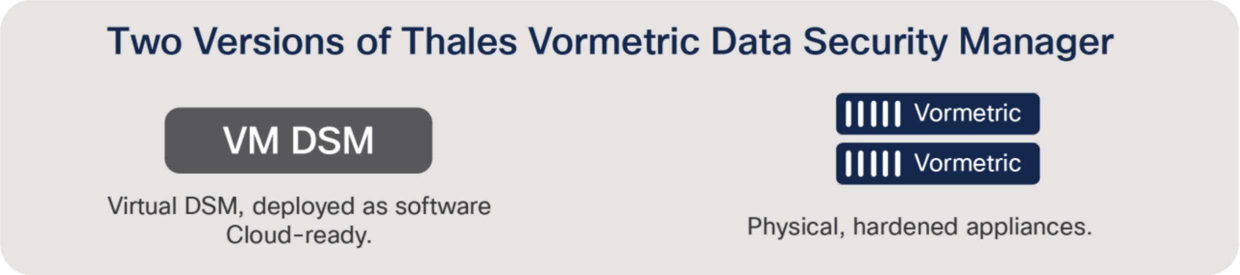
Two Versions of Thales Vormetric Data Security Manager
VM DSM
Virtual DSM, deployed as software
Cloud-ready.
Vormetric
Vormetric
Physical, hardened appliances.

Figure 4.
Vormetric Key Exchange Variants

This document specifically tested Vormetric DSM as the KMS for vCenter in the implementation of VMCrypt. This deployment is represented below.

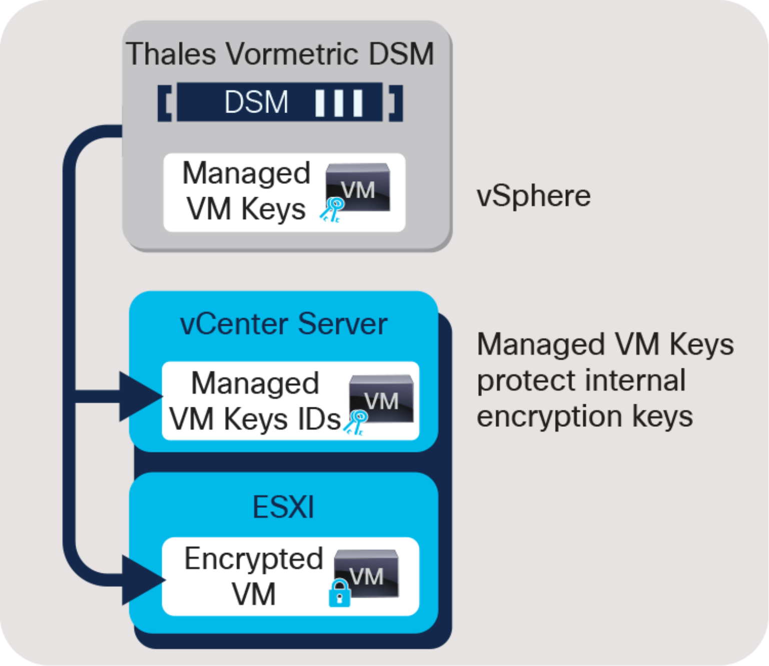

Figure 5.
Vormetric and VMware Key Architecture

VM-level encryption

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

Thales Vormetric Transparent Encryption (now called CipherTrust Transparent Encryption) https://www.thalesesecurity.com/products/data-encryption/vormetric-transparent-encryption

Gemalto:

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

VMware ESXi 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

VMCrypt characteristics

The HyperFlex environment provides storage for guest VMs deployed on VMware ESXi using VLAN-segmented networking. The VMs are available for external resources, typical of any elastic infrastructure deployment. VMCrypt works above the HXDP level and above the VM level. It operates at the hypervisor (ESXi) layer and is implemented through vCenter and the remote KMS.

- VMCrypt: vSphere-native encryption starting in ESXi 6.5 has a few requirements, given below:
    - Requires minimum vSphere Enterprise Plus
    - ESXi and vCenter 6.5
    - Requires an Encryption Storage Policy to be created on vCenter
    - Third-party remote Key Management Server (KMS)

VMCrypt has the following characteristics:

- No modification of the guest VM is required
- VMCrypt supports any guest OS
- VMCrypt supports any hardware version
- Any datastore will work
- VM policy is based through vCenter
- VMDK and VM files are encrypted
- Guest cannot access keys
- vMotion encryption is supported

How VMCrypt works

Encrypted VMs using VMCrypt rely on vCenter-based Storage Policies. The workflow for VM encryption using VMCrypt has three distinct paths, depending on the state of the VM:

- New VM
    - Apply the Encryption Storage Policy that you created when creating the VM
    - A randomly generated Disk Encryption Key (DEK) is created on the host
    - DEK is encrypted with the KMS Key Encryption Key (KEK)
    - All I/O is encrypted
- Existing unencrypted VM
    - The VM must be powered off
    - Apply the Encryption Storage Policy that you have created
    - A randomly generated Disk Encryption Key (DEK) is created on the host
    - DEK is encrypted with the KMS Key Encryption Key (KEK)
    - o The VMDK for this VM is then encrypted. (This can take some time)
- Existing encrypted VM
    - vCenter reads the key encryption key ID from the VMX file for the VM
    - vCenter retrieves the key encryption key from the KMS
    - vCenter then loads this KEK on the host

Decrypting a VM

To decrypt a VM, you simply need to change the vCenter-based Storage Policy applied to the VM. By changing the Storage Policy back to **Datastore default**, the VM's files and VMDKs will be decrypted. Like enabling encryption, this will take time, depending on the size of the decryption operation.

VMCrypt limitations – READ THIS BEFORE ENABLING VMCRYPT ENCRYPTION

There are some limitations to VMCrypt, described below:

- The default KMS must be acquired separately. The tests conducted here used Thales Vormetric DSM Key Manager
- SAN backup is not supported
- Backup data is not backed up encrypted. The backup solution may provide its own encrypted mechanism
    - After restoring a VM, it must be re-encrypted
- vCenter cannot be encrypted
- The following functions are not supported:
    - VM "Suspend and resume"
    - Encrypting a VM with existing snapshots
    - Serial and parallel ports
    - vSphere replication
    - Snapshots of encrypted VMs

Key Management Protocol (KMIP) 1.1 has to be implemented in order for the key manager to be compatible with vSphere 6.5.

Setting up vCenter with Thales Vormetric DSM

A prerequisite for VMCrypt is to have the Thales Vormetric DSM key manager installed and functional in your environment. At a high level, the following steps are required to deploy DSM:

1. Download and deploy the DSM virtual machine. Alternatively, if you have a physical DSM, you can use that
2. For the virtual DSM, power on the VM and use vCenter to launch a console.
    - In this test environment, a single non-HA deployment of the virtual DSM was used
    - Configure HA if needed
3. Log in to the DSM command line using the default login and password: cliadmin/cliadmin123
4. Type "Help" at the command prompt to see a list of available commands
    a. Using system and network commands, assign an IP and route to the system
    b. Restart the server to take effect
5. From the CLI, generate the system certificate using the genCA command
6. Log in to the web interface and continue the configuration
    a. For the webUI, use the default login and password: admin/admin123
    b. You will be prompted to change the default password
7. Apply your DSM license to unlock the KMIP features

8. Create a new administrative user for use with the VMCrypt domain
9. Create the VMCrypt domain and assign this user
10. Change the context to the VMCrypt domain

For the complete setup guidelines, see the References section for information on getting and building DSM.

Once your key manager is installed, you are ready to set up the KMS in vCenter. The complete vCenter-DSM configuration instructions are detailed in Appendix B; however, a brief overview is discussed below.

To begin the configuration, you will need to select your vCenter instance in the navigation pane on the left. Select the Configure tab and then select key management servers under More. Complete the wizard to put the KMS details into the system. For DSM in particular, leave the username and password blank. (Certificates are used for authentication; populating these fields will cause the key management to fail).

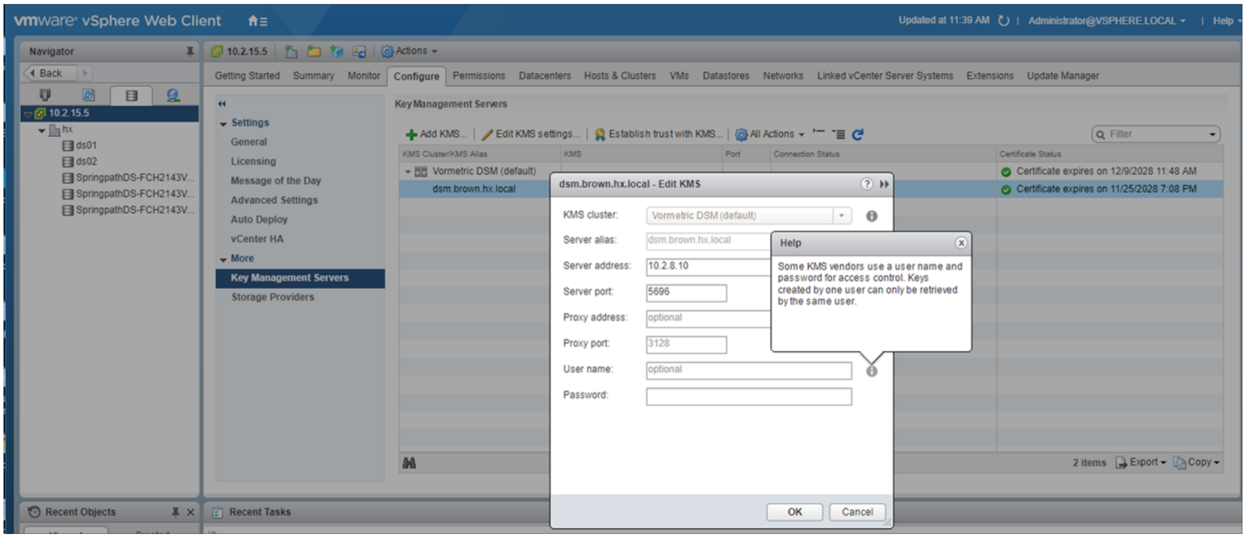
vmware vSphere Web Client
Updated at 11:39 AM | Administrator@VSPHERE.LOCAL | Help
Navigator
Back
10.2.15.5
hx
ds01
ds02
SpringpathDS-FCH2143V...
SpringpathDS-FCH2143V...
SpringpathDS-FCH2143V...
10.2.15.5 | Actions
Getting Started | Summary | Monitor | Configure | Permissions | Datacenters | Hosts & Clusters | VMs | Datastores | Networks | Linked vCenter Server Systems | Extensions | Update Manager
Settings
General
Licensing
Message of the Day
Advanced Settings
Auto Deploy
vCenter HA
More
Key Management Servers
Storage Providers
Key Management Servers
Add KMS... | Edit KMS settings... | Establish trust with KMS... | All Actions
Filter
2 items | Export | Copy
Recent Objects
Recent Tasks

KMS Cluster/KMS Alias	KMS	Port	Connection Status	Certificate Status
Vormetric DSM (default)				Certificate expires on 12/9/2028 11:48 AM
dsm.brown.hx.local				Certificate expires on 11/25/2028 7:08 PM

dsm.brown.hx.local - Edit KMS

KMS cluster: Vormetric DSM (default)
Server alias: dsm.brown.hx.local
Server address: 10.2.8.10
Server port: 5696
Proxy address: optional
Proxy port: 3128
User name: optional
Password:

Help

Some KMS vendors use a user name and password for access control. Keys created by one user can only be retrieved by the same user.

Once the KMS is configured, choose "Yes" on the "Set as default KMS cluster". This should be selected regardless of whether the KMS is a single instance or an actual HA cluster.

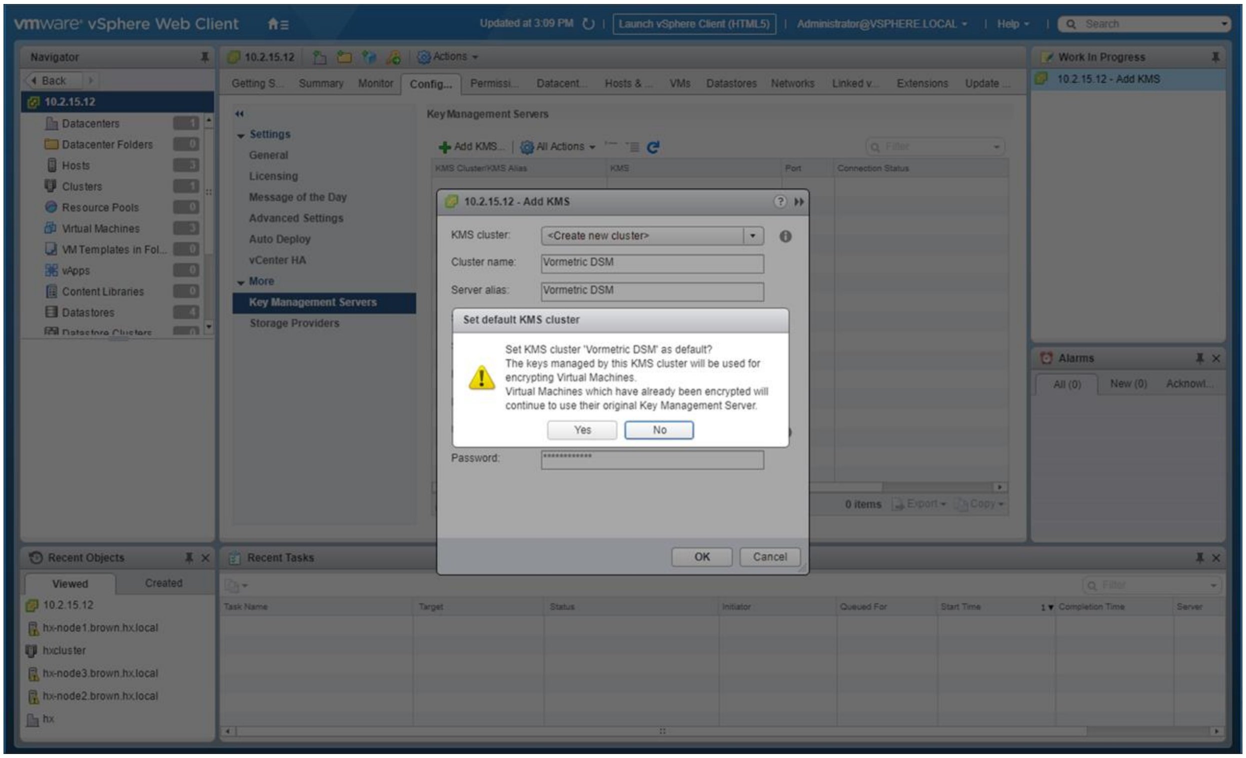

10.2.15.12 - Add KMS

KMS cluster: <Create new cluster>

Cluster name: Vormetric DSM

Server alias: Vormetric DSM

Set default KMS cluster

Set KMS cluster 'Vormetric DSM' as default?
The keys managed by this KMS cluster will be used for encrypting Virtual Machines.
Virtual Machines which have already been encrypted will continue to use their original Key Management Server.

Yes | No

Password: ************

OK | Cancel

You will then be asked to establish trust with the KMS. Select the Certificate option, and press OK.

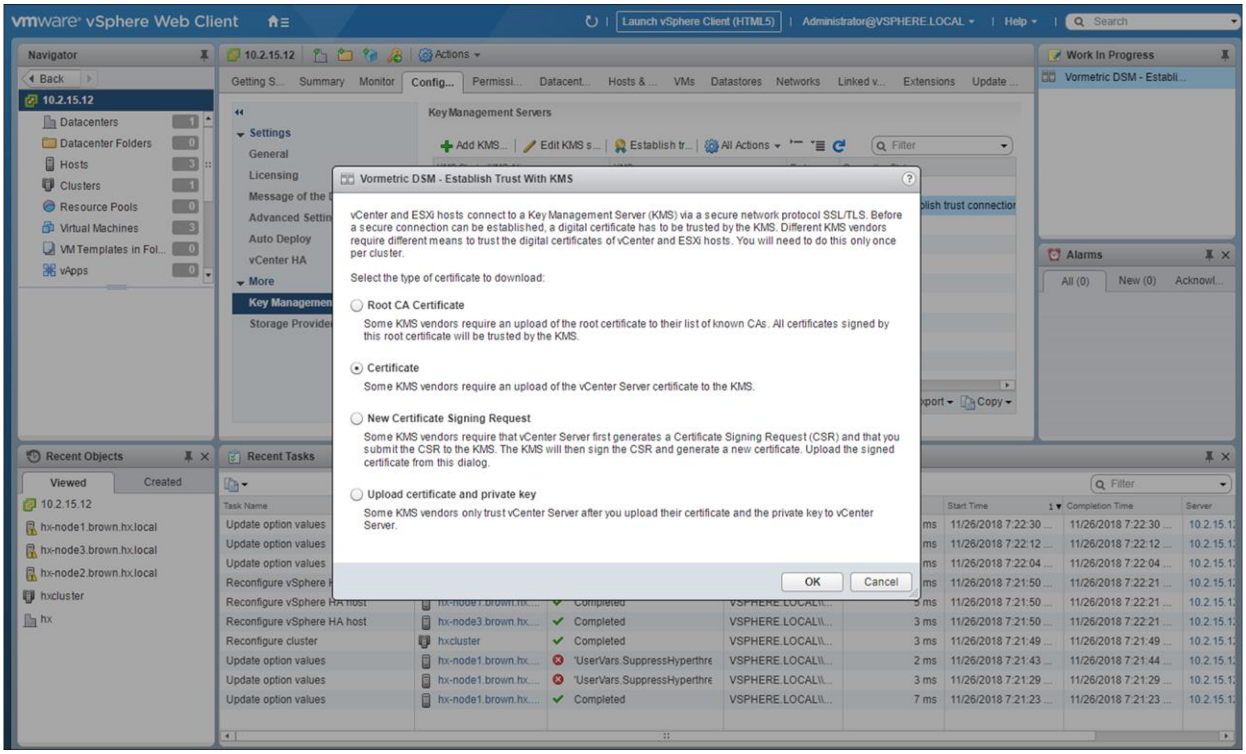

Vormetric DSM - Establish Trust With KMS

vCenter and ESXi hosts connect to a Key Management Server (KMS) via a secure network protocol SSL/TLS. Before a secure connection can be established, a digital certificate has to be trusted by the KMS. Different KMS vendors require different means to trust the digital certificates of vCenter and ESXi hosts. You will need to do this only once per cluster.

Select the type of certificate to download:

○ Root CA Certificate

Some KMS vendors require an upload of the root certificate to their list of known CAs. All certificates signed by this root certificate will be trusted by the KMS.

● Certificate

Some KMS vendors require an upload of the vCenter Server certificate to the KMS.

○ New Certificate Signing Request

Some KMS vendors require that vCenter Server first generates a Certificate Signing Request (CSR) and that you submit the CSR to the KMS. The KMS will then sign the CSR and generate a new certificate. Upload the signed certificate from this dialog.

○ Upload certificate and private key

Some KMS vendors only trust vCenter Server after you upload their certificate and the private key to vCenter Server.

OK | Cancel

You will then be presented with the certificate information. Press the Trust button.

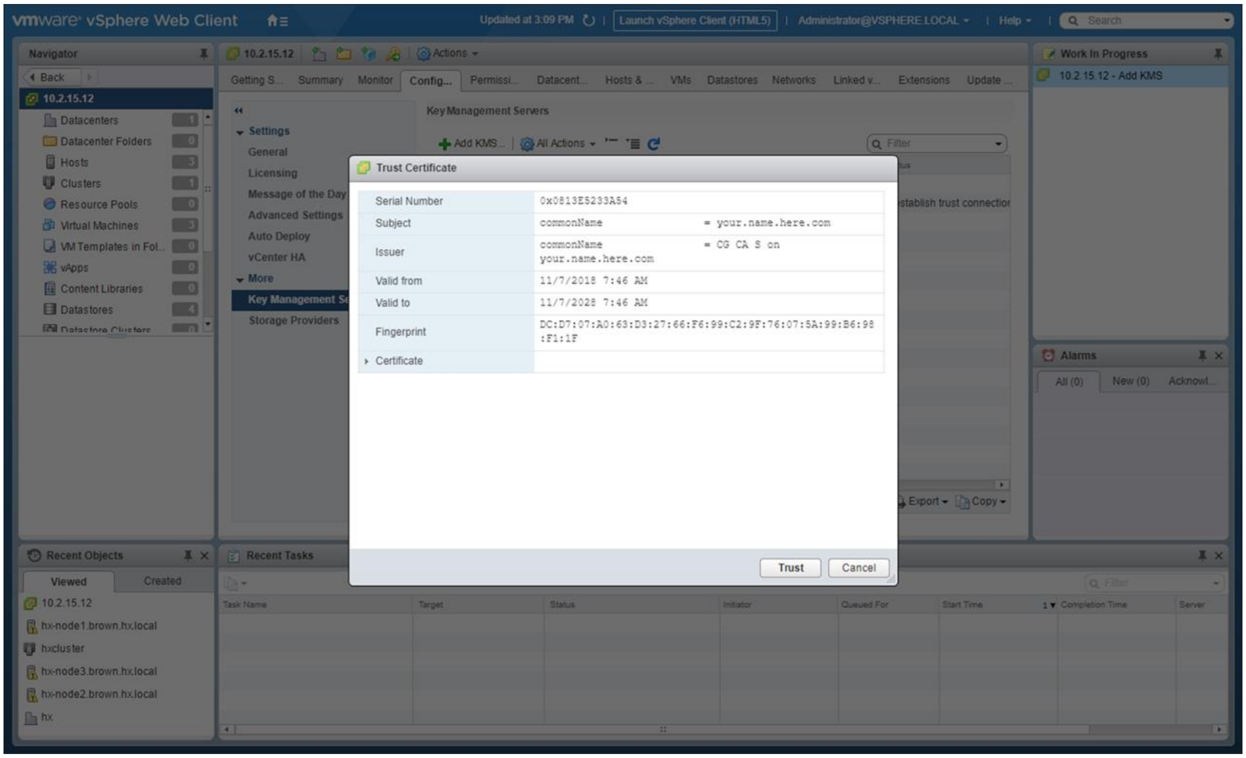

Serial Number	0x0813E5233A54
Subject	commonName = your.name.here.com
Issuer	commonName = CG CA S on your.name.here.com
Valid from	11/7/2018 7:46 AM
Valid to	11/7/2028 7:46 AM
Fingerprint	DC:D7:07:A0:63:D3:27:66:F6:99:C2:9F:76:07:5A:99:B6:98:F1:1F
▸ Certificate	

Download the certificate file and change the extension to “crt”. You will need to open the certificate in a reader (see Appendix B) and copy the common name. Next, import the certificate file into the KMIP section of your VMCrypt domain using the DSM UI.

Go back to vCenter and refresh the connection status for the KMS. It should appear with a green check as Normal.

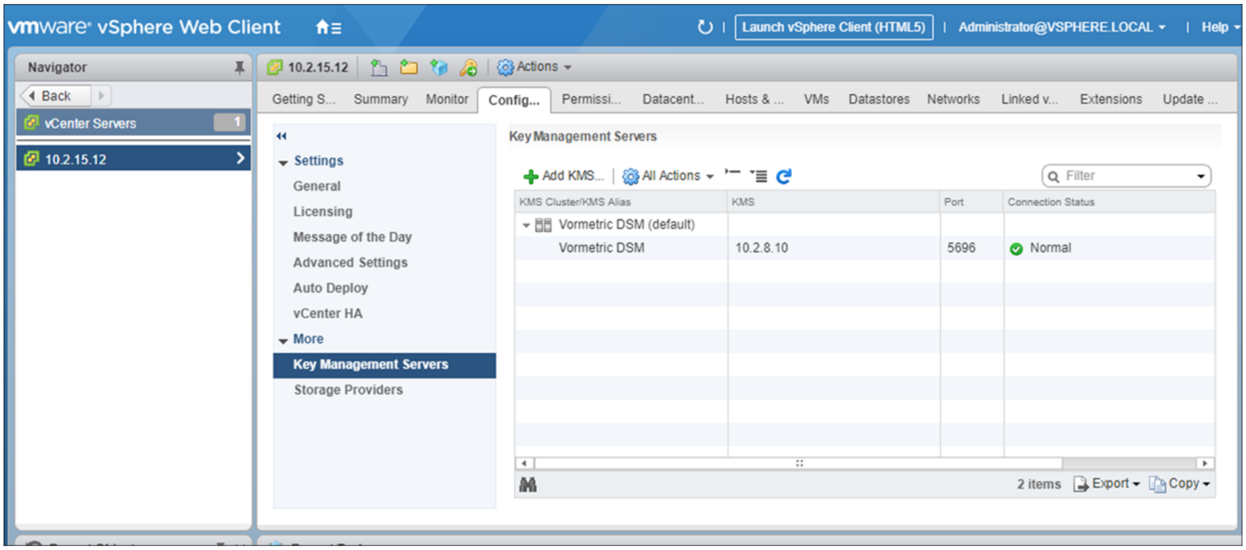

Key Management Servers

KMS Cluster/KMS Alias	KMS	Port	Connection Status
Vormetric DSM (default)			
Vormetric DSM	10.2.8.10	5696	✓ Normal

Turning on encryption

Once vCenter and your KMS (DSM) are in proper communication, you will need to create an encryption storage policy in vCenter in order to enable VM encryption. In vCenter, select the storage policies in the administration interface and create a new encryption policy. Follow the policy creation wizard.

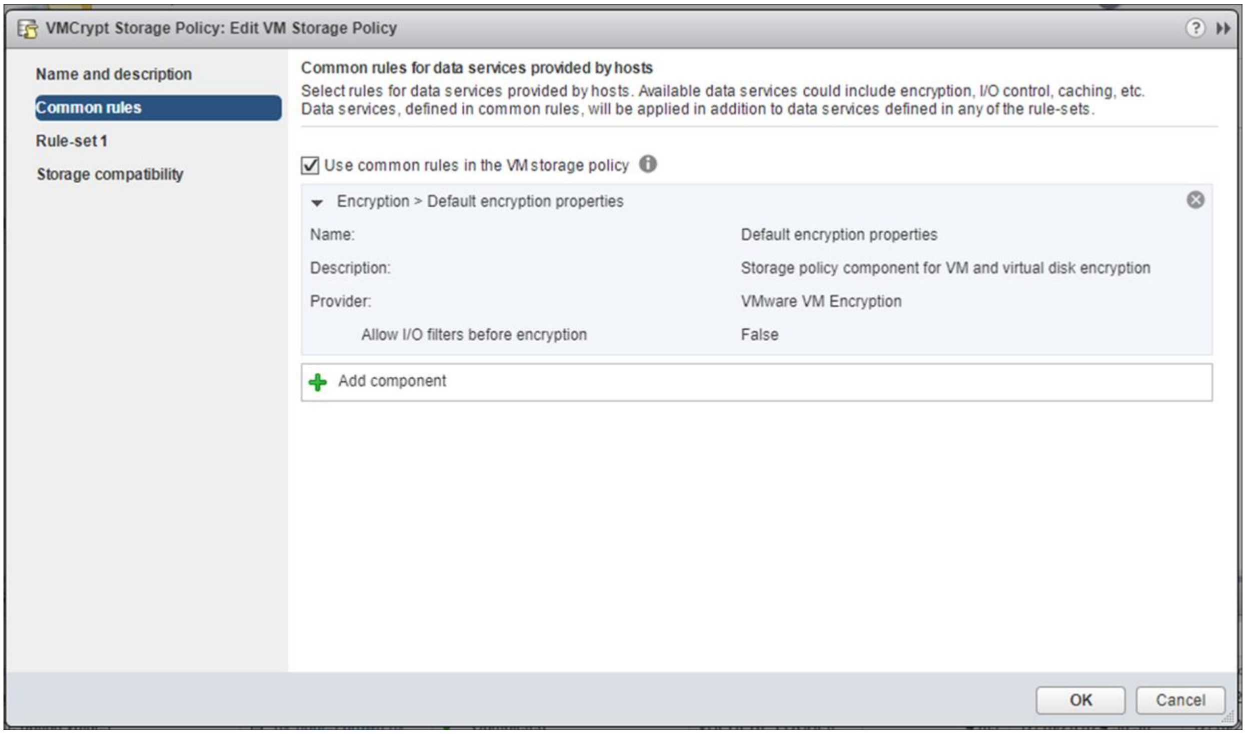

VMCrypt Storage Policy: Edit VM Storage Policy

Common rules for data services provided by hosts

Select rules for data services provided by hosts. Available data services could include encryption, I/O control, caching, etc. Data services, defined in common rules, will be applied in addition to data services defined in any of the rule-sets.

☑ Use common rules in the VM storage policy

Encryption > Default encryption properties

Name:	Default encryption properties
Description:	Storage policy component for VM and virtual disk encryption
Provider:	VMware VM Encryption
Allow I/O filters before encryption	False

Once your encryption policy is created, you are ready to apply it to the VM(s) that need encryption. If you are running a comparison benchmark, as was done in this paper, you will want to run your unencrypted benchmark before you enable encryption. While it is possible to decrypt an encrypted VM (see the section on Decrypting a VM, above), it makes the most sense to run your unencrypted tests first, then encrypt and rerun your benchmark with the now-encrypted Cisco HxBench subordinate VMs.

With VMCrypt, for existing unencrypted VMs, you will need to ensure that the VM is shut down and powered off before you can apply the encryption storage policy. Right click on the VM, and select Edit VM Storage Policies...

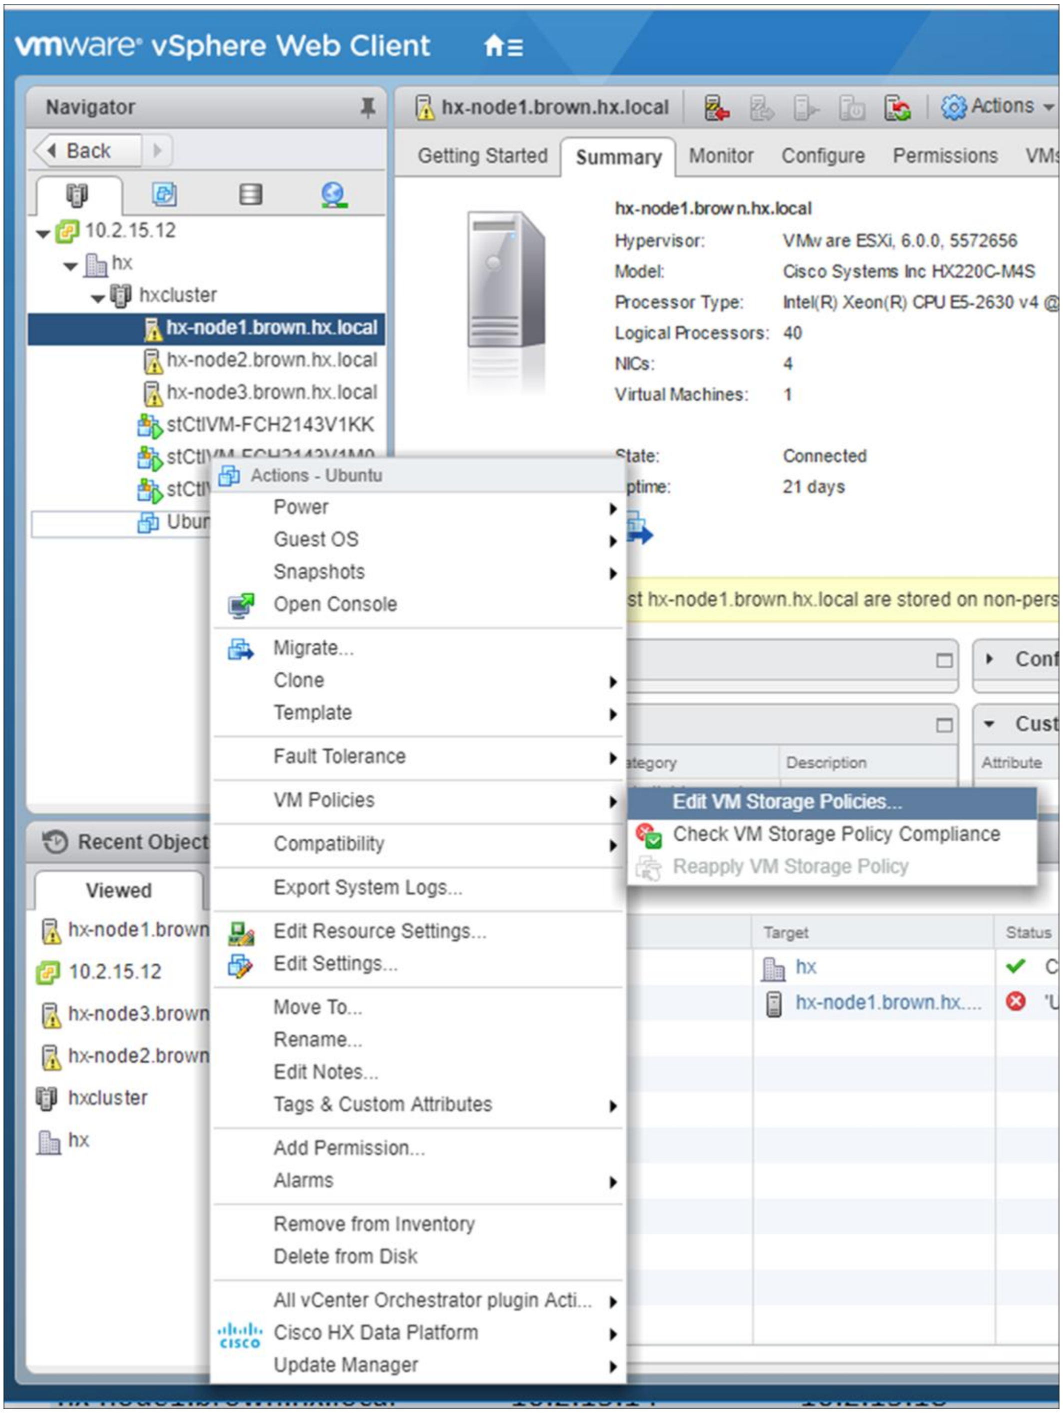

hx-node1.brown.hx.local

Hypervisor: VMware ESXi, 6.0.0, 5572656

Model: Cisco Systems Inc HX220C-M4S

Processor Type: Intel(R) Xeon(R) CPU E5-2630 v4 @

Logical Processors: 40

NICs: 4

Virtual Machines: 1

State: Connected

ptime: 21 days

st hx-node1.brown.hx.local are stored on non-pers

Target	Status
hx	✓ C
hx-node1.brown.hx....	✕ 'U

In the follow-on dialogue, select the VM Encryption storage policy from the drop-down menu for each component of the VM (VM home and Hard disk 1). Click on OK.

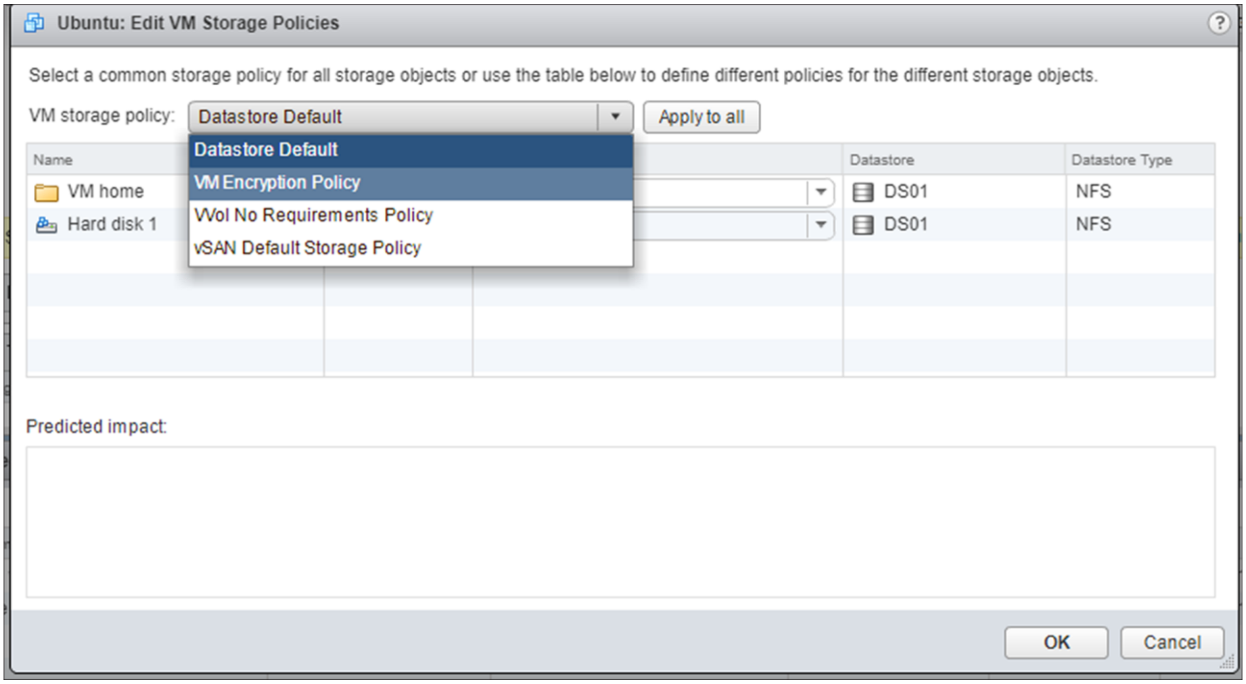

Ubuntu: Edit VM Storage Policies

Select a common storage policy for all storage objects or use the table below to define different policies for the different storage objects.

VM storage policy: Datastore Default

Name			Datastore	Datastore Type
VM home			DS01	NFS
Hard disk 1			DS01	NFS

Predicted impact:

For a typical Cisco HxBench subordinate VM with a 100 GB working set and an 8 GB OS disk, the initial encryption on an M5 Edge system will take about 47 minutes. You can kick off all of the encryption operations for each subordinate VM at the same time.

HXDP test bed

Cisco HyperFlex Edge

The Cisco HyperFlex (HX) Edge deployment in the encryption test bed consisted of 3 HX Edge nodes. Each node is dual connected to a top-of-rack 1 GBE switch (Cisco Nexus® 5000 series). The lab is segmented with a services subnet. This services subnet contains the following:

- NTP
- DNS
- AD
- KMS
- DHCP

A separate infrastructure segment was also present for the deployment of the cluster's vCenter instance and the HyperFlex installer used to build the Edge cluster. These were all routable to the management vSwitch on the Edge cluster. The Edge cluster and test setup were installed with the following software versions:

- Cisco HyperFlex HX Data Platform 3.5(1a) with RF2
- VMware ESXi v6.5U2
- vCenter 6.5

- Cisco HxBench v 1.3.7
  - 60% read
  - Full curve test (VM load 10%/50%/75%/90%/100%) at 60 mins per interval
  - Block size 8k
  - Threads/VM: 32
  - Load VM/node: 1
  - Deduplication 0% and 50%
  - Compression 0% and 50%
  - Workload size: 100GB and 300GB
  - OS VMDK per subordinate VM: 8GB
- Vdbench 5.0.6
- Vormetric DSM v 6.0.1 single instance VM (No HA clustering was used).

Cisco HyperFlex Edge testbed hardware:

- HXAF 220M5: 3 hybrid nodes
  - 2 onboard 1GBE interfaces per node
  - 1 cache SSD at 447 GB per node
  - 6 rotational disks at 1.1 TB per node
  - 2 Intel® Xeon® CPU E5-2630 v4 processors (40 logical cores total) per node
- Cisco Nexus 5000 Series Switches

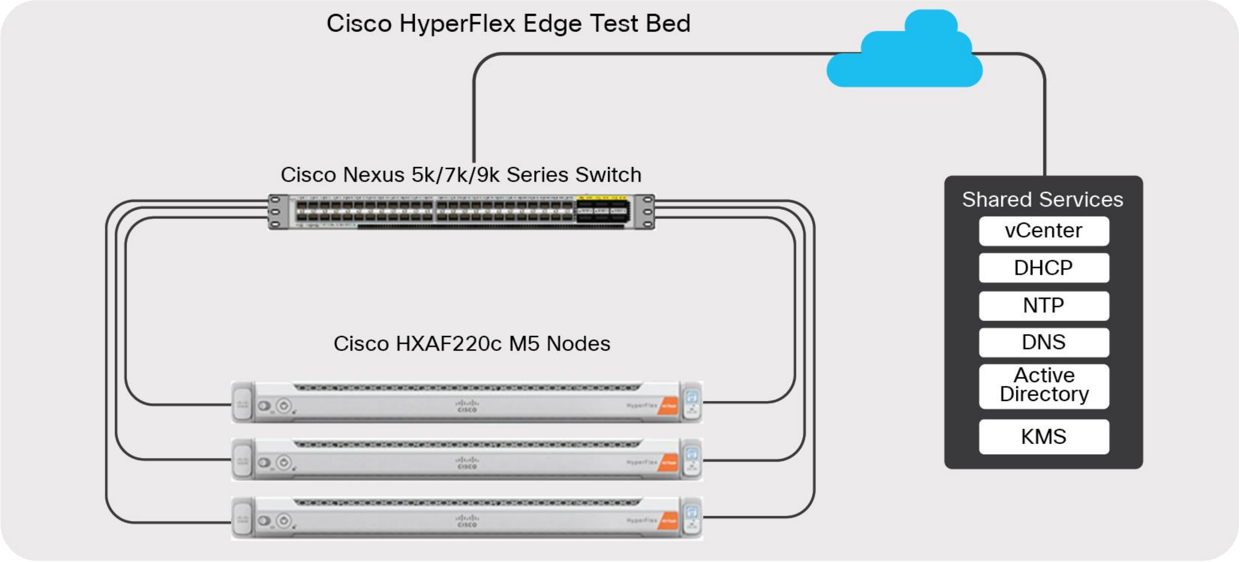

Figure 6.
VMCrypt and HyperFlex Test Bed

This test bed was used with Cisco HxBench to deploy one load (subordinate) VM per node. Cisco HxBench was configured to run many different tests using small and large working sets along with different deduplication and compression ratios. A representative run is presented in the results section. The tests were used for a unencrypted set of runs followed by an encrypted set of benchmarks. The tests take approximately six hours each to run, after the test VM workload disks are freshly primed.

The Cisco HxBench subordinate VMs were encrypted after the first set of tests were run using the encryption storage policy set up in vCenter. The VM encryption took approximately 47 minutes per small (100 GB) VM and all three were run concurrently. It took correspondingly longer to encrypt the large (300 GB) VMs. Note that the subordinate VMs needed to be powered off for the application of the encryption storage policy.

Each Cisco HxBench test run was configured to use the existing subordinate VMs from the previous run(s), but each VMDK data set was primed for each run to ensure identical initial conditions for each baseline.

Test results: Cisco HxBench VM unencrypted vs. encrypted

See Appendix A for a comprehensive overview of the HX Bench test configuration parameters.

In general, CPU utilization increased across the board for the various test scenarios as expected. The performance impact ranged from approximately 10 percent to 50 percent, depending on the specifics of the test. For a smaller working set with no deduplication or compression on the system, turning on encryption had a negligible impact on overall VM performance. For a large working set with 2:1 deduplication and compression the impact can be close to 50 percent. The latency for all tests remained under 10ms even when the impact was most pronounced in IOPS. This indicates that it may be possible to mitigate the impact of encryption on those workloads by allocating more vCPU resources to the VM.

The following test results are for a large working set (that is, it overfills the cache) where the data set is not dedupable or compressible. A workload where deduplication and compression of the HxBench workload is set to 0 percent and 0 percent, respectively, most closely “levels” the playing field in a CPU-impact evaluation for encryption because the encrypted working set is not compressible or dedupable.

Unencrypted:

Table 2. Unecrypted Performance

Load Point	10%	50%	75%	90%	100%
IOPS	1,700	8,100	12,093	14,068	12,185
Latency (ms)	2.2	4.0	6.0	6.8	8.4
Throughput (MBps)	13.3	63.3	94.5	109.9	95.2

Encrypted:

Table 3. Encrypted Performance

Load Point	10%	50%	75%	90%	100%
IOPS	1,501	7,499	11,200	12,909	10,813
Latency (ms)	2.1	3.9	6.7	7.7	9.4

Load Point	10%	50%	75%	90%	100%
Throughput (MBps)	11.7	58.6	87.5	100.8	84.5

Cisco HxBench unencrypted vs. encrypted: 300 GB, no deduplication or compression

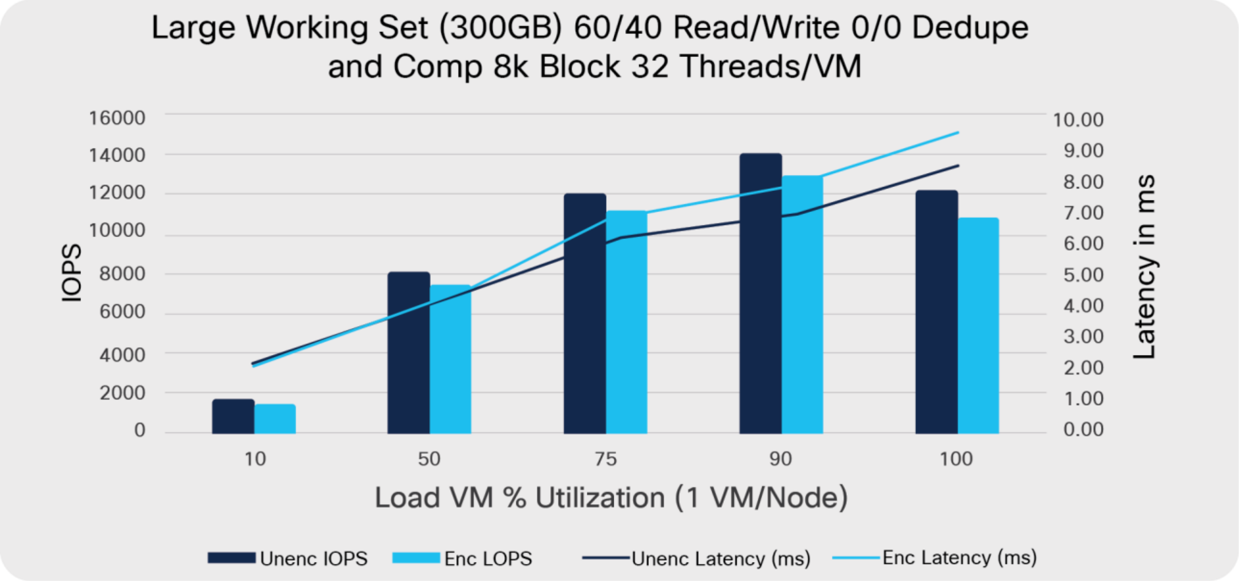

Figure 7.
Comprehensive Performance Comparison

Latencies remain within 1ms across the load spectrum.

- @10% encrypted IOPS are ~88% of unencrypted IOPS
- @50% encrypted IOPS are ~93% of unencrypted IOPS
- @75% encrypted IOPS are ~93% of unencrypted IOPS
- @90% encrypted IOPS are ~92% of unencrypted IOPS
- @100% encrypted IOPS are ~89% of unencrypted IOPS

Note the IOP inflection point from 90 percent to 100 percent indicating that something is probably going on with cache misses to the persistent tier (which, in this test bed, is rotational media).

There is a consistent ~10 percent IOP penalty due to CPU overhead in encryption operations for this particular test case. Looking at these results, we can see the following:

- There is a definite ~10 percent CPU overhead on the encryption process up until the cache overflowed (somewhere after 90 percent)
- Even though the cache is overfilled, there is still about a 10 percent penalty to IOPs in the encryption process, so it is cache-fill-independent
- The cache overflow has an effect regardless of encryption or plain text, so you should size properly

Cache note: Cache population in a hybrid system follows a Zipf distribution where 10 percent of the content is accessed 90 percent of the time. The read cache is approximately 10 percent of the persistent tier on average in terms of disk capacity. This means that, when fully warmed, the cache size matches the Zipf caching algorithm in use. Since a large portion of the read cache is also metadata (~50 to 60 percent), and in this case, the cache in the testbed is ~1300 GB vs. 900 GB of the working set, we can be guaranteed that, with this set, we will overflow the cache at some load point.

Conclusion

Use of the ESXi-native encryption capability, VMCrypt, is supported with Cisco HyperFlex Edge systems. Proper configuration of vCenter with a third-party remote key management server is required. Subsequent creation of an encryption storage policy allows granular encryption of selected VMs. This works in conjunction with system hardening to ensure a secure Cisco HX Edge deployment. The performance results vary, depending on VM size, workload read/write breakdown, and transitory versus steady-state conditions. Based on Cisco HxBench results for this type of environment and these test scenarios, you can expect the encrypted VMs to perform at approximately 55 to 90 percent of the IOPS compared to the same workload on unencrypted VMs at 100 percent benchmark load.

References

Port requirements for communication

- See the HyperFlex Hardening Guide link below.

Cisco HyperFlex Edge Deployment Guide

- https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/Edge_Deployment_Guide/b_HyperFlex_Edge_Deployment_Guide_3_5.html

Thales Vormetric DSM installation and configuration

- Contact your Thales Vormetric representative for the DSM installation and configuration documents, the virtual DSM OVA, and a temporary license if needed.

Cisco HxBench (Cisco account required)

- https://HyperFlexsizer.cloudapps.cisco.com/ui/index.html#/hx_bench

Oracle Vdbench

- https://www.oracle.com/technetwork/server-storage/vdbench-downloads-1901681.html

HX Hardening Guide

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

ESXi Hardening Guide

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

Appendix A. Cisco HxBench configuration information

Cisco HxBench is a storage performance testing tool that can be used to measure the storage infrastructure performance using defined test scenarios and workloads. HxBench uses the industry standard Vdbench storage performance benchmarking tool. HxBench is provided as a virtual appliance to deploy and run tests. HxBench simplifies performance benchmarking for virtualized environments.

See the references section for a link to download the tool and the required Vdbench file from Oracle. While you can deploy HxBench using static IPs, it is easiest to use a DHCP server for the network segment where you are deploying. This was done in the test bed used for this paper.

Deploy the HxBench VM in vCenter using the Deploy OVF functionality, and complete the deployment wizard. The summary will look like the following screenshot.

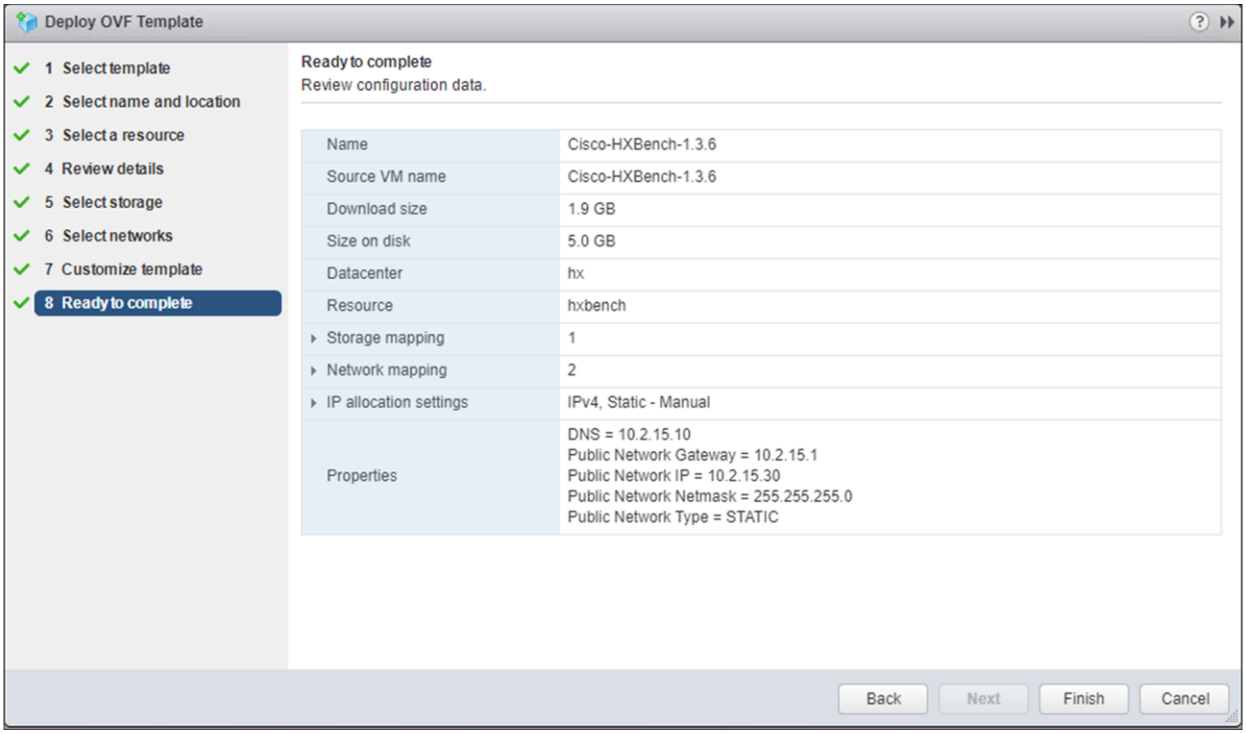

Deploy OVF Template

1. Select template ✓
2. Select name and location ✓
3. Select a resource ✓
4. Review details ✓
5. Select storage ✓
6. Select networks ✓
7. Customize template ✓
8. Ready to complete ✓

Ready to complete

Review configuration data.

Name	Cisco-HXBench-1.3.6
Source VM name	Cisco-HXBench-1.3.6
Download size	1.9 GB
Size on disk	5.0 GB
Datacenter	hx
Resource	hxbench
Storage mapping	1
Network mapping	2
IP allocation settings	IPv4, Static - Manual
Properties	DNS = 10.2.15.10 Public Network Gateway = 10.2.15.1 Public Network IP = 10.2.15.30 Public Network Netmask = 255.255.255.0 Public Network Type = STATIC

Once the HxBench VM is deployed, you can access the HxBench Web Interface by entering one of the following into your browser:

- https://<IPaddressofHxBenchControllerVM>:8000 or
- https://<IPaddressofHxBenchControllerVM>:8000/hxbench/index.html

Log in to the HxBench Web Interface using administrator credentials:

- Username: appadmin
- Password: <admin> password

Upon first login, you will be prompted to upload Vdbench. This is a one-time activity for the HxBench controller. Upload the Vdbench software to the HxBench controller.

1. Download Vdbench software version 5.04.06 from the Oracle website. Download the vdbench50406.zip file from: https://www.oracle.com/technetwork/server-storage/vdbench-downloads-1901681.html
2. Click on Start. Upload the vdbench50406.zip file to the HxBench controller using the Upload button

Upon successful completion of the Vdbench software upload, click on Next. Provide your server details for the vCenter host where the tests should run:

- Host name: vCenter host name (or IP)
- User name: <admin> username (typically administrator@vsphere.local)
- Password: <admin> password.

When you are ready to create and run a test:

- Click on the Run Test button
- Choose a predefined test. Short and long baselines, respectively, were used in this paper
- Leave the defaults for the predefined tests

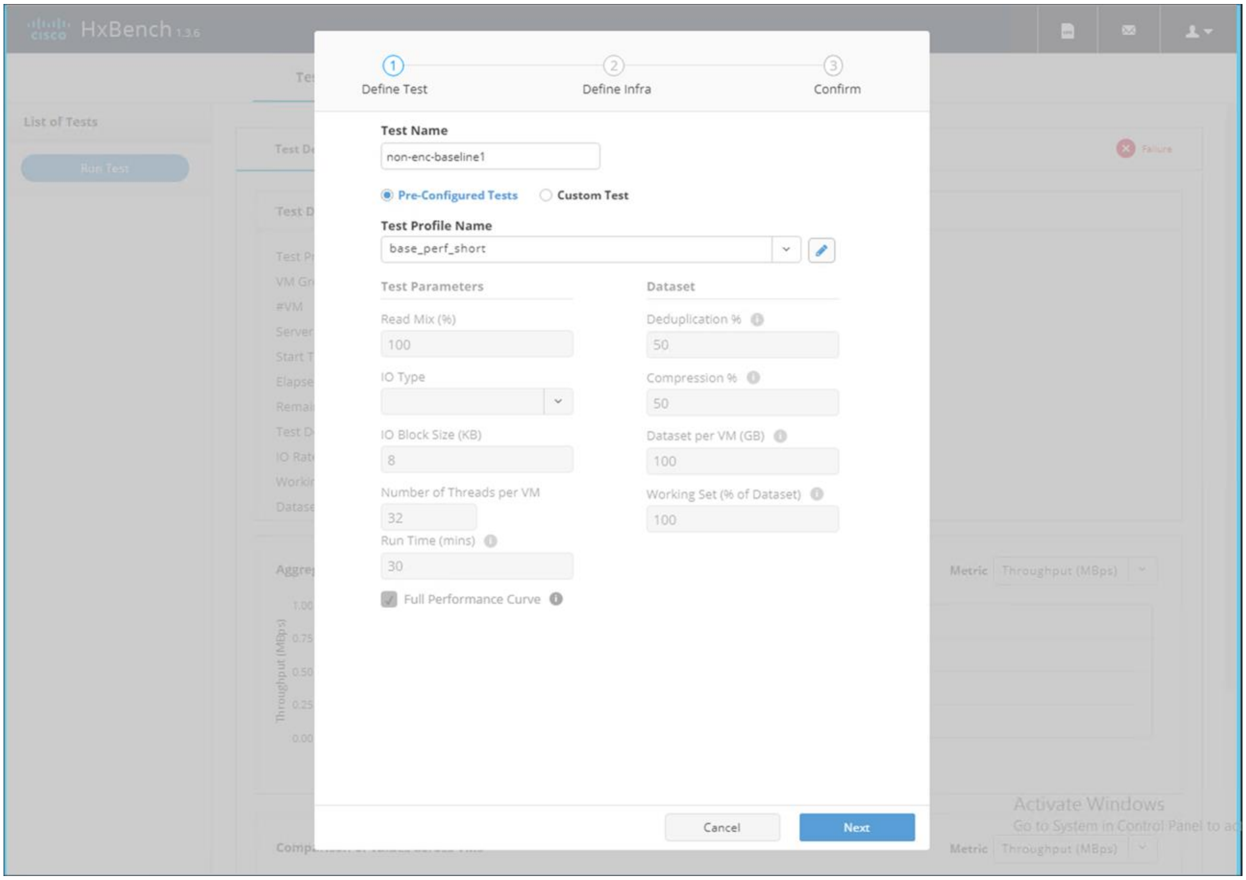

Test Name: non-enc-baseline1

● Pre-Configured Tests ○ Custom Test

Test Profile Name: base_perf_short

Test Parameters

Read Mix (%): 100

IO Type:

IO Block Size (KB): 8

Number of Threads per VM: 32

Run Time (mins): 30

☑ Full Performance Curve

Dataset

Deduplication %: 50

Compression %: 50

Dataset per VM (GB): 100

Working Set (% of Dataset): 100

Match the subordinate VM count to your node count. Since the HyperFlex Edge system being tested has 3 nodes, enter “3” in the “Total VMs across all nodes” field.

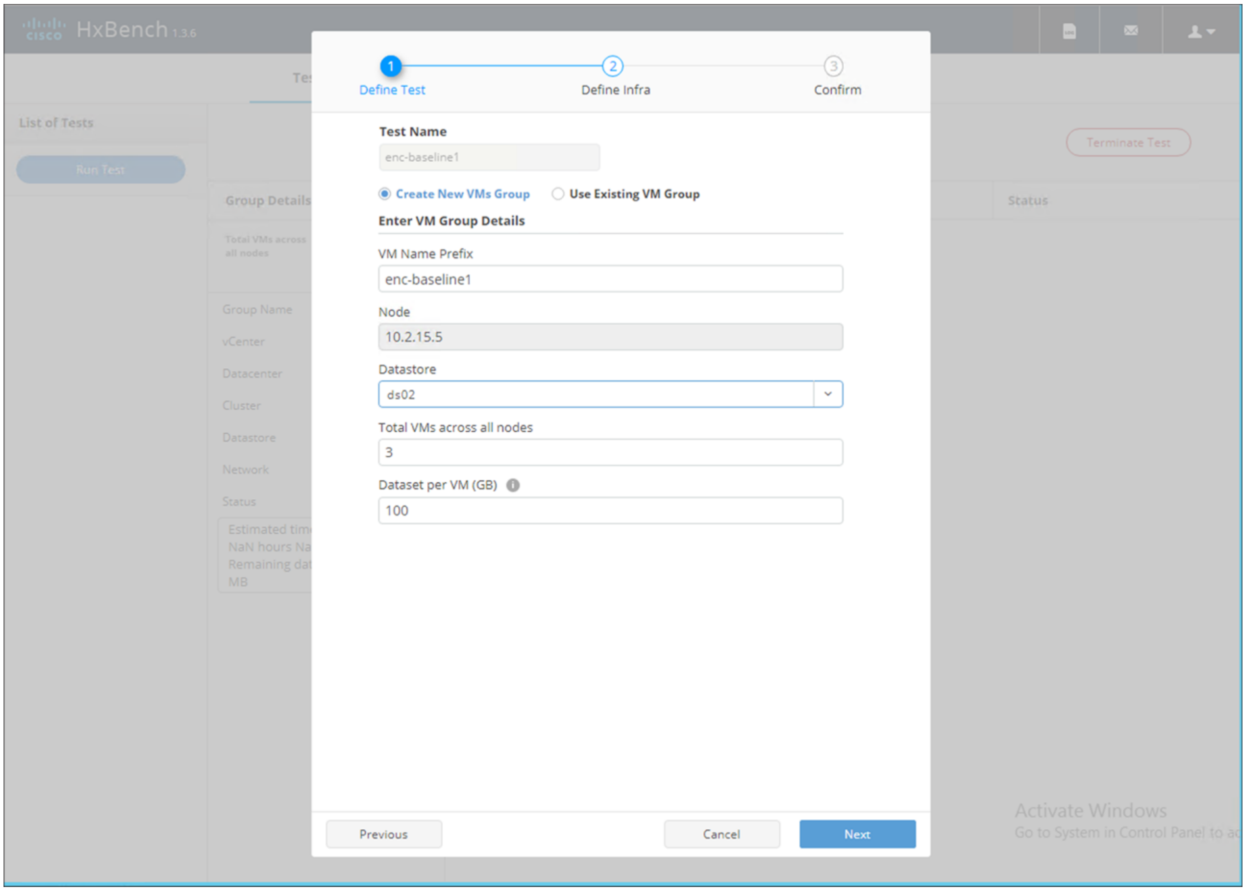

Test Name

enc-baseline1

● Create New VMs Group ○ Use Existing VM Group

Enter VM Group Details

VM Name Prefix: enc-baseline1

Node: 10.2.15.5

Datastore: ds02

Total VMs across all nodes: 3

Dataset per VM (GB): 100

Verify the test summary, and confirm the configuration.

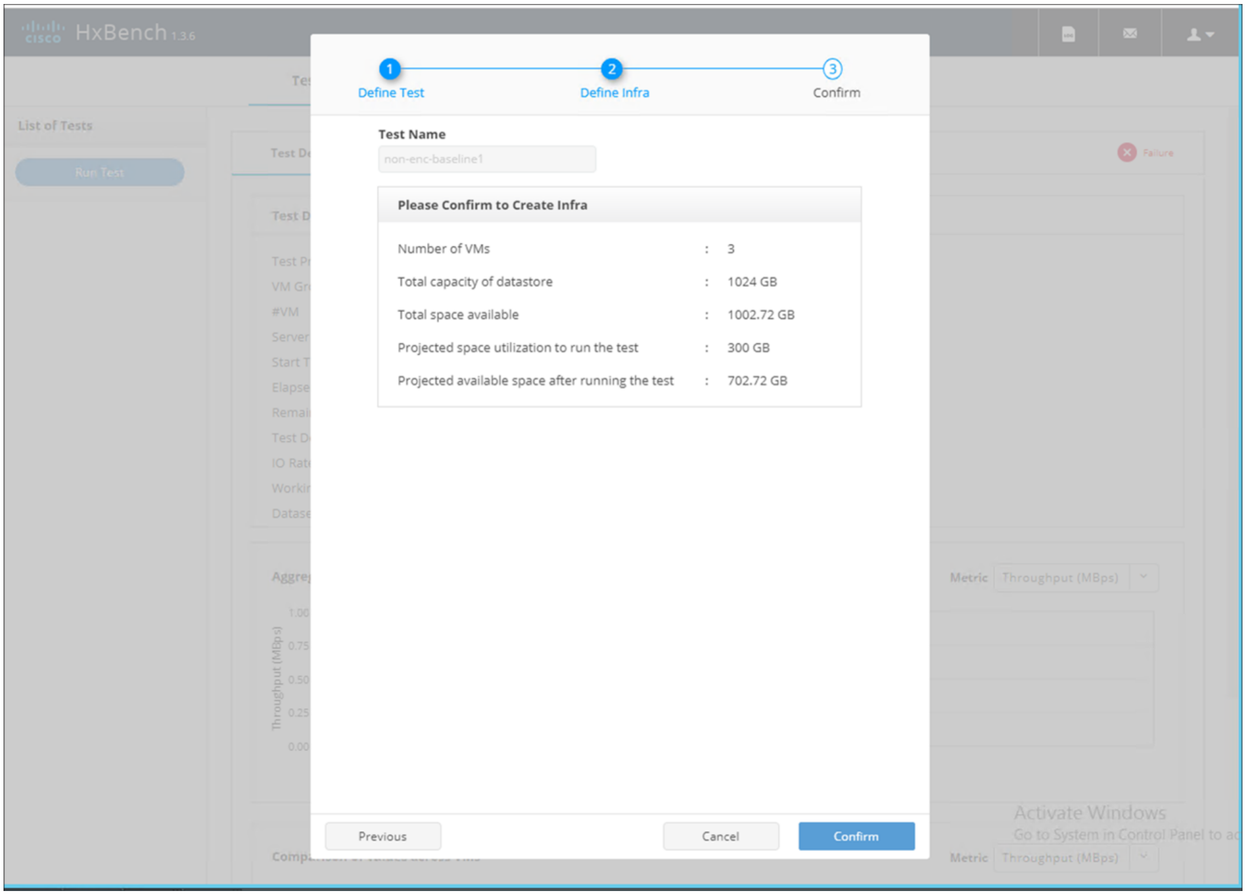

Test Name

non-enc-baseline1

Please Confirm to Create Infra

Number of VMs	:	3
Total capacity of datastore	:	1024 GB
Total space available	:	1002.72 GB
Projected space utilization to run the test	:	300 GB
Projected available space after running the test	:	702.72 GB

After confirming, you’ll see VMs being deployed as follows:

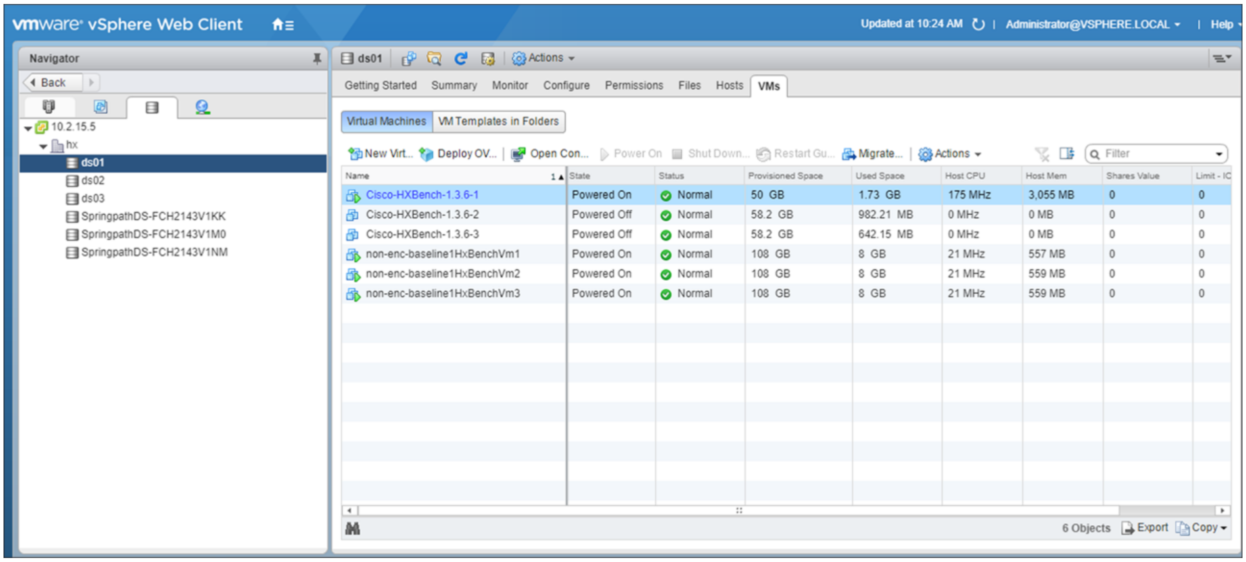

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem	Shares Value	Limit - IC
Cisco-HXBench-1.3.6-1	Powered On	Normal	50 GB	1.73 GB	175 MHz	3,055 MB	0	0
Cisco-HXBench-1.3.6-2	Powered Off	Normal	58.2 GB	982.21 MB	0 MHz	0 MB	0	0
Cisco-HXBench-1.3.6-3	Powered Off	Normal	58.2 GB	642.15 MB	0 MHz	0 MB	0	0
non-enc-baseline1HxBenchVm1	Powered On	Normal	108 GB	8 GB	21 MHz	557 MB	0	0
non-enc-baseline1HxBenchVm2	Powered On	Normal	108 GB	8 GB	21 MHz	559 MB	0	0
non-enc-baseline1HxBenchVm3	Powered On	Normal	108 GB	8 GB	21 MHz	559 MB	0	0

Once the VMs are deployed, the test will run. Subsequent tests can reuse the same VMs to reduce operational time. Be sure not to select “Don’t prime the data” so each run will get a fresh set of new data. This will help ensure that all test runs experience identical initial conditions.

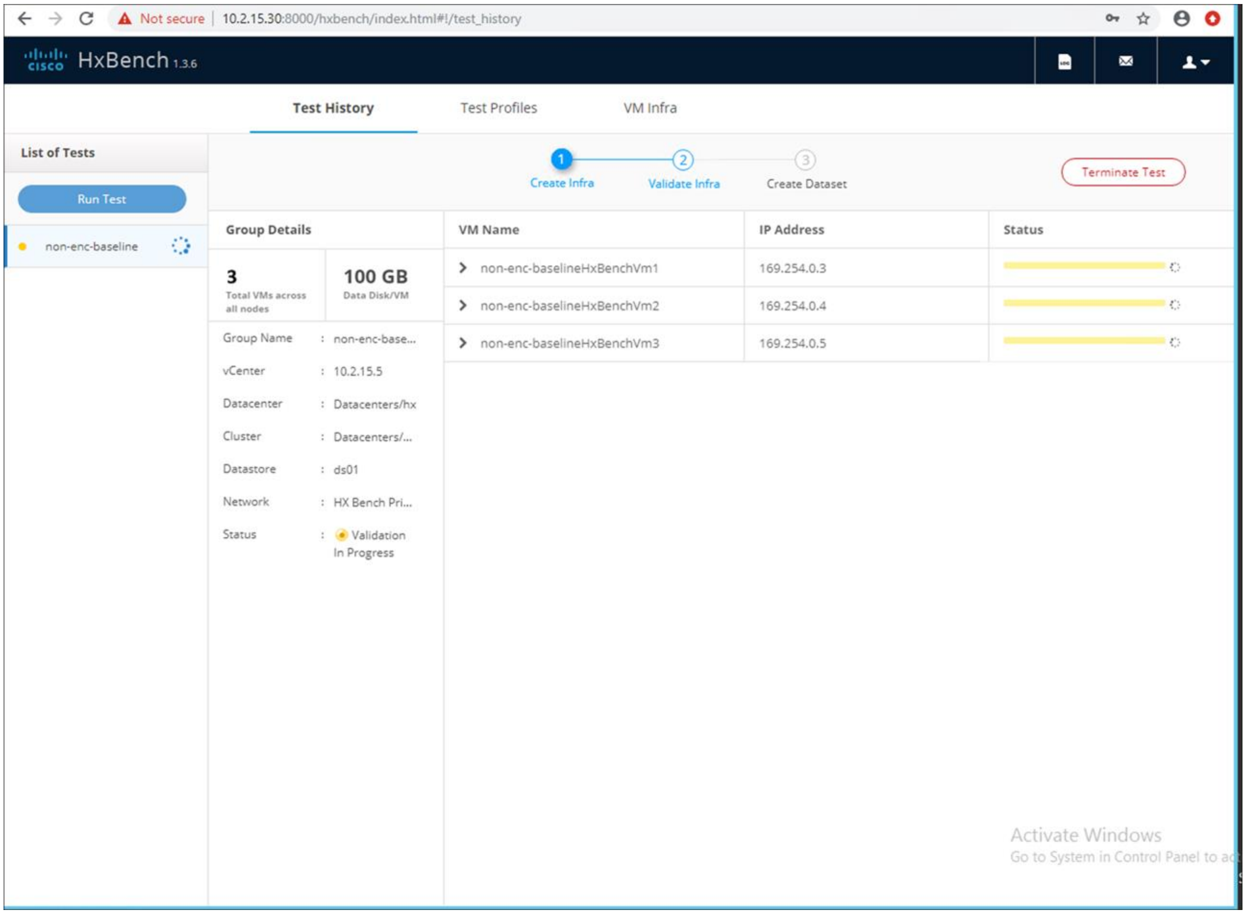

Group Details

3 Total VMs across all nodes | 100 GB Data Disk/VM

Group Name : non-enc-base...
vCenter : 10.2.15.5
Datacenter : Datacenters/hx
Cluster : Datacenters/...
Datastore : ds01
Network : HX Bench Pri...
Status : Validation In Progress

VM Name	IP Address	Status
non-enc-baselineHxBenchVm1	169.254.0.3	
non-enc-baselineHxBenchVm2	169.254.0.4	
non-enc-baselineHxBenchVm3	169.254.0.5	

Appendix B. vCenter and DSM configurations

The following table lists the ports required for component communication for the HyperFlex solution. Special thanks to Paul Cleary at Thales Vormetric for the procedure.

B.1 Introduction

This document outlines the steps necessary to enable the Vormetric Data Security Manager (DSM) to be the Key Management Server (KMS) for VMware-encrypted disks.

B.1. High-level overview

1. Create the KMS cluster in vSphere.
2. Establish trust between the vCenter Server and DSM cluster.
3. Verify that trust has been established.
4. Enable encryption for VMs via VMware Storage Policies.

B.2. Create the KMS cluster in vSphere

The first step is to configure VMware to use a Key Management Server (KMS).

1. In vSphere, select the cluster that encryption will be enabled on and click on the tab **Configure** and then on **Key Management Servers**.

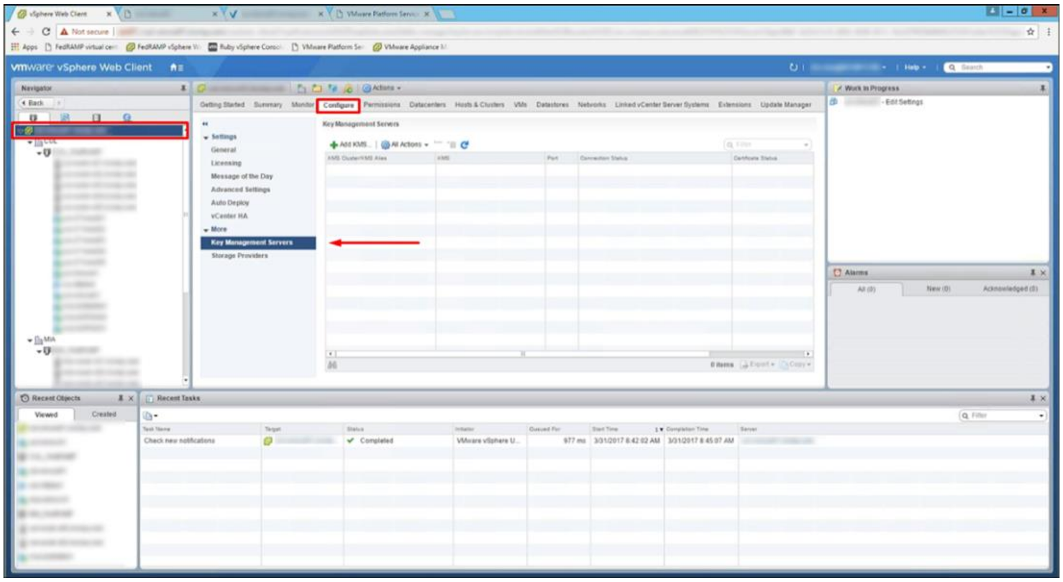

2. Click on **Add KMS** and enter the following information relevant to your organization:

Cluster name	Name of the KMS cluster that you want to create.
Server alias	Use this alias to connect to the KMS if your vCenter Server instance becomes unavailable.
Server address and port	IP address or FQDN of the DSM, and port on which the vCenter Server connects to the DSM (**The KMIP port is 5696**).

3. Click on **Trust** in the Trust Certificates dialogue box to trust the DSM.

B.3. Establish trust between the vCenter Server and DSM cluster

Next, a certificate is exported from vCenter and imported into the DSM.

1. Select the KMS instance that was just created and click on **Establish Trust with KMS**.

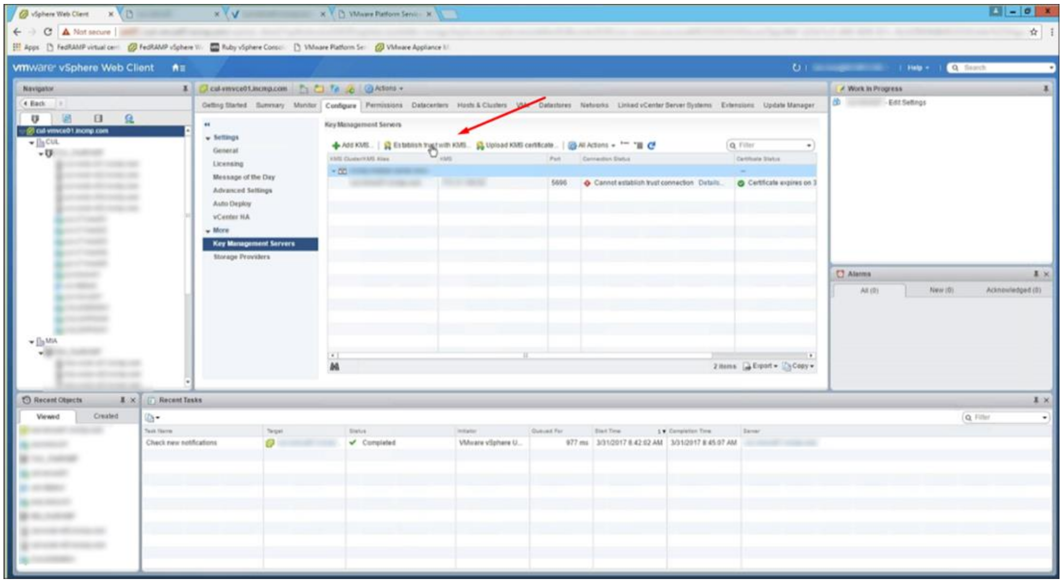

5696 | Cannot establish trust connection Details... | Certificate expires on 3

Check new notifications | Completed | VMware vSphere U... | 977 ms | 3/31/2017 8:42:02 AM | 3/31/2017 8:45:07 AM

2. In the dialogue box that appears, select **Certificate** and click on **OK**.
3. Select **Download as File** (The file will download in the ".pem" format).
4. Change the file extension to **crt**.
5. Display information about the downloaded certificate to get the common name of the vCenter Server.
   a. Use the command: openssl x509 -in certificate.crt -text -noout.

   OR

   b. Use an online certificate decoder: https://www.sslshopper.com/certificate-decoder.html

6. Copy the common name to the clipboard.

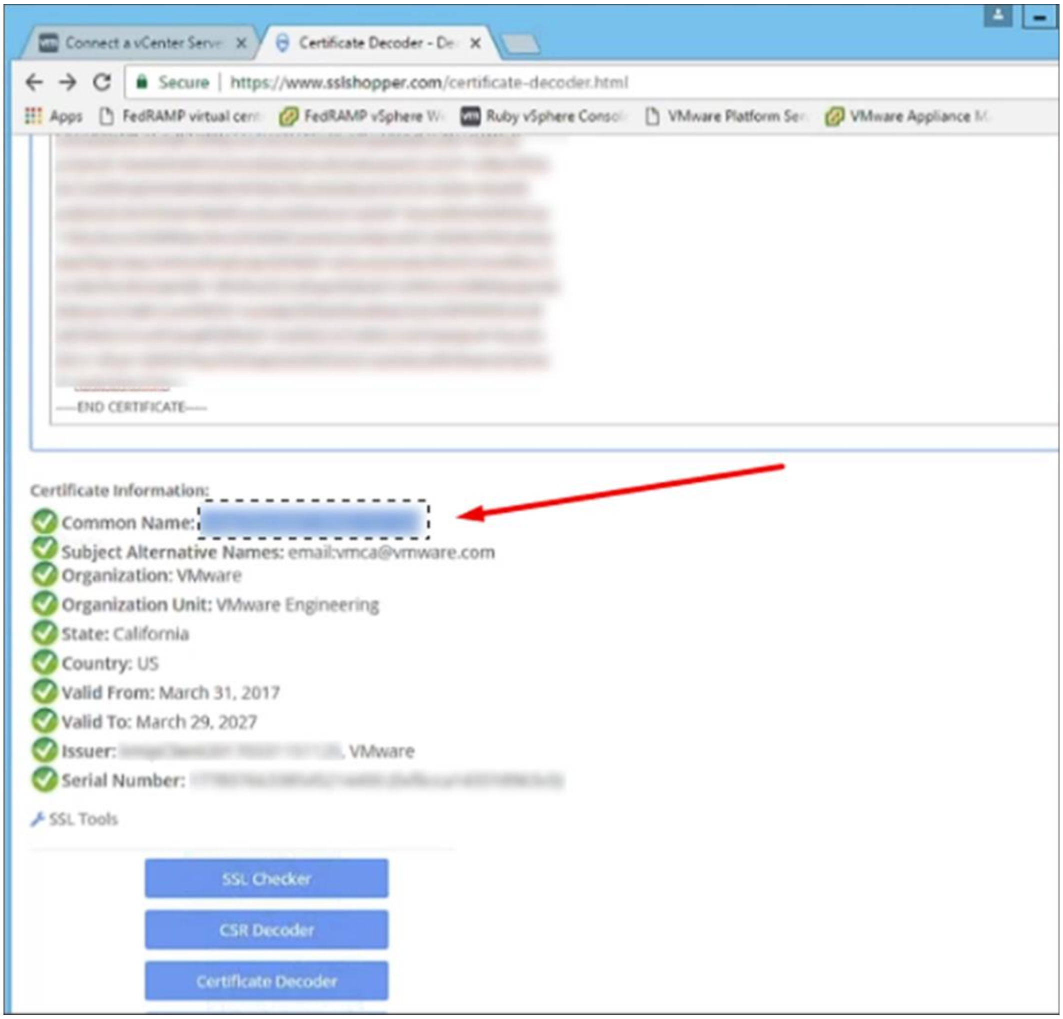

[redacted]

-----END CERTIFICATE-----

Certificate Information:

Common Name: [redacted]

Subject Alternative Names: email:vmca@vmware.com

Organization: VMware

Organization Unit: VMware Engineering

State: California

Country: US

Valid From: March 31, 2017

Valid To: March 29, 2027

Issuer: [redacted], VMware

Serial Number: [redacted]

7. Log in to the DSM and switch to the domain that will be used by KMIP devices.
8. Click on **Hosts** and then on **Add**, to add a new host.
9. Enter the common name from step 5 into the Host Name box.
10. Check on the box next to Registration Allowed Agents: KMIP.
11. Check on the box next to Communication Enabled, and click on **OK**.
12. Click on the host name that was just created.

13. Click on **Import KMIP Certificate** in the bottom right-hand corner.

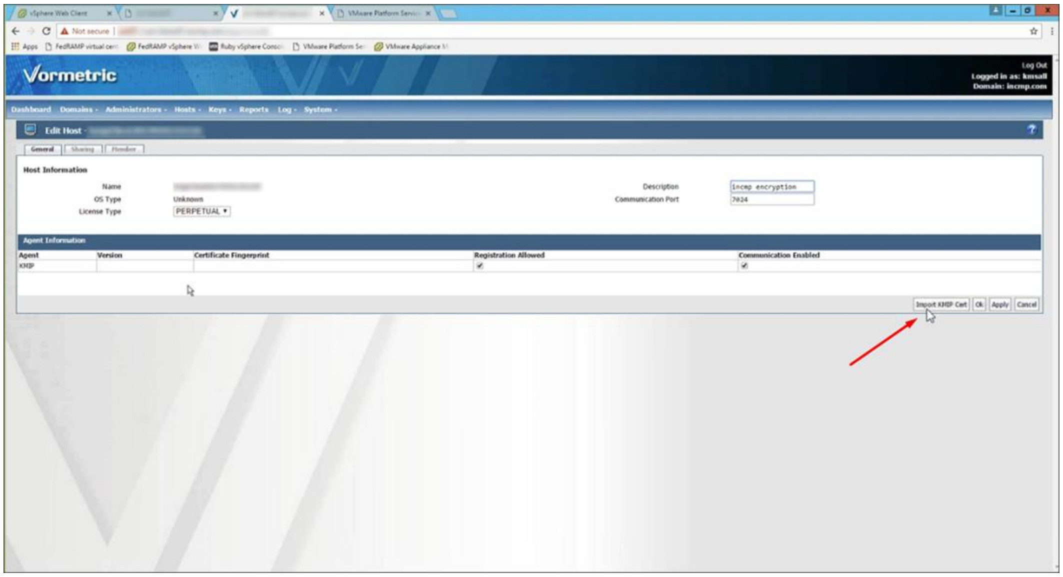

Host Information

Name	[redacted]
OS Type	Unknown
License Type	PERPETUAL ▾
Description	incmp encryption
Communication Port	7024

Agent Information

Agent	Version	Certificate Fingerprint	Registration Allowed	Communication Enabled
KMIP			☑	☑

14. After importing the certificate, the fingerprint from the vCenter Server will be listed in the KMIP Fingerprint field.

B.4. Verify trust between the vCenter Server and DSM cluster

Ensure that the vCenter Server is able to connect securely.

1. Refresh the vSphere Web Client page and look for a green check next to Normal under the **Connection Status column**.

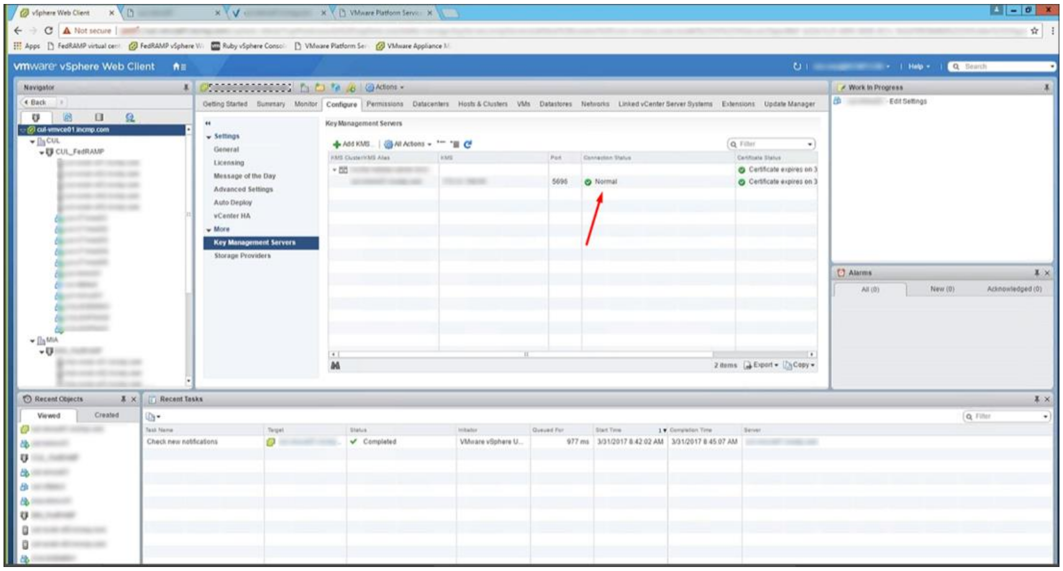

Key Management Servers

KMS Cluster/KMS Alias	KMS	Port	Connection Status	Certificate Status
[redacted]				✓ Certificate expires on 3
[redacted]	[redacted]	5696	✓ Normal	✓ Certificate expires on 3

Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time	Server
Check new notifications	[redacted]	✓ Completed	VMware vSphere U...	977 ms	3/31/2017 8:42:02 AM	3/31/2017 8:45:07 AM	[redacted]

B.5. Enable encryption for target servers

Finally, enable encryption by using VMWare Storage Policies.

1. Right-click on the VM that you would like to enable encryption for, and select **VM Policies>Edit VM Storage Policies**.
2. Select the storage policy that utilizes the Vormetric DSM as the key management server, and click on **Apply to All** and then on **OK**.
3. This will cause a reconfiguration of the VM.

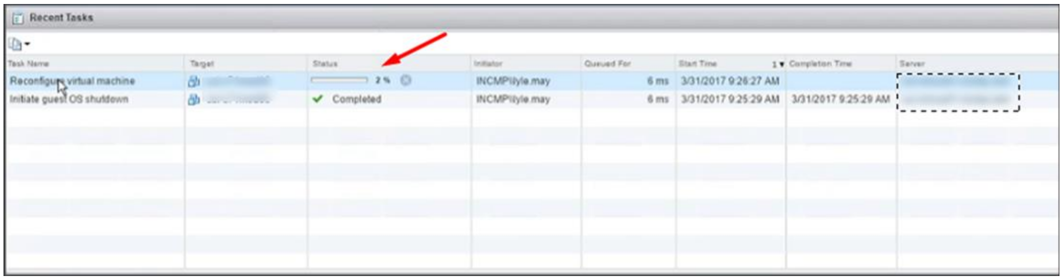

Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time	Server
Reconfigure virtual machine	[redacted]	2 %	INCMP\lyle.may	6 ms	3/31/2017 9:26:27 AM		[redacted]
Initiate guest OS shutdown	[redacted]	✓ Completed	INCMP\lyle.may	6 ms	3/31/2017 9:25:29 AM	3/31/2017 9:25:29 AM	[redacted]

4. When it is complete, the disks are encrypted and the keys are being managed by the Vormetric DSM.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C11-742309-02 03/22