Luna HSM

Product Overview

SafeNet

THE DATA PROTECTION COMPANY

Document Information

Product Version	5.4.1
Document Part Number	007-011329-006
Release Date	04 July 2014

Revision History

Revision	Date	Reason
A	26 February 2014	Initial release.
B	17 April 2014	Updates to the SFF Backup feature.
C	04 July 2014	Solaris client support.

Trademarks

All intellectual property is protected by copyright. All trademarks and product names used or referred to are the copyright of their respective owners. No part of this document may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, chemical, photocopy, recording or otherwise without the prior written permission of SafeNet, Inc.

Disclaimer

SafeNet makes no representations or warranties with respect to the contents of this document and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, SafeNet reserves the right to revise this publication and to make changes from time to time in the content hereof without the obligation upon SafeNet to notify any person or organization of any such revisions or changes.

We have attempted to make these documents complete, accurate, and useful, but we cannot guarantee them to be perfect. When we discover errors or omissions, or they are brought to our attention, we endeavor to correct them in succeeding releases of the product.

SafeNet invites constructive comments on the contents of this document. Send your comments, together with your personal and/or company details to the address below.

Contact Method	Contact Information
Mail	SafeNet, Inc. 4690 Millennium Drive Belcamp, Maryland 21017 USA
Email	techpubs@safenet-inc.com

CONTENTS

Rellease 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

 Copyright 2014 SafeNet, Inc. All rights reserved.

PREFACE

About the Product Overview

This document provides an overview of Luna HSM suite of products. It contains the following chapters:

This preface also includes the following information about this document:

For information regarding the document status and revision history, see "Document Information" on page 2

Customer release notes

The customer release notes (CRN) provide important information about this release that is not included in the customer documentation. Read the CRN to fully understand the capabilities, limitations, and known issues for this release. You can view or download the latest version of the CRN for this release at the following location:

- http://www.securedbysafenet.com/releasenotes/luna/crn_luna_hsm_5-4.pdf

Audience

This document is intended for personnel responsible for maintaining your organization's security infrastructure. This includes Luna HSM users and security officers, key manager administrators, and network administrators.

All products manufactured and distributed by SafeNet, Inc. are designed to be installed, operated, and maintained by personnel who have the knowledge, training, and qualifications required to safely perform the tasks assigned to them. The information, processes, and procedures contained in this document are intended for use by trained and qualified personnel only.

It is assumed that the users of this document are proficient with security concepts.

Document conventions

This document uses standard conventions for describing the user interface and for alerting you to important information.

Release 5.4.1 007-011329-006Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Notes

Notes are used to alert you to important or helpful information. They use the following format:

Note: Take note. Contains important or helpful information.

Cautions

Cautions are used to alert you to important information that may help prevent unexpected results or data loss. They use the following format:

CAUTION: Exercise caution. Contains important information that may help prevent unexpected results or data loss.

Warnings

Warnings are used to alert you to the potential for catastrophic data loss or personal injury. They use the following format:

WARNING! Be extremely careful and obey all safety and security measures. In this situation you might do something that could result in catastrophic data loss or personal injury.

Command syntax and typeface conventions

Format	Convention
bold	The bold attribute is used to indicate the following: • Command-line commands and options (Type dir /p.) • Button names (Click Save As.) • Check box and radio button names (Select the Print Duplex check box.) • Dialog box titles (On the Protect Document dialog box, click Yes.) • Field names (User Name: Enter the name of the user.) • Menu names (On the File menu, click Save.) (Click Menu > Go To > Folders.) • User input (In the Date box, type April 1.)
italics	In type, the italic attribute is used for emphasis or to indicate a related document. (See the *Installation Guide* for more information.)
<variable>	In command descriptions, angle brackets represent variables. You must substitute a value for command line arguments that are enclosed in angle brackets.
[**optional**] [<optional>]	Represent optional **keywords** or <variables> in a command line description. Optionally enter the keyword or <variable> that is enclosed in square brackets, if it is necessary or desirable to complete the task.

 Copyright 2014 SafeNet, Inc. All rights reserved.

Format	Convention
{**a**|**b**|**c**} {<a>|<b>|<c>}	Represent required alternate **keywords** or <variables> in a command line description. You must choose one command line argument enclosed within the braces. Choices are separated by vertical (OR) bars.
[**a**|**b**|**c**] [<a>|<b>|<c>]	Represent optional alternate keywords or variables in a command line description. Choose one command line argument enclosed within the braces, if desired. Choices are separated by vertical (OR) bars.

Support Contacts

If you encounter a problem while installing, registering or operating this product, please ensure that you have read the documentation. If you cannot resolve the issue, please contact your supplier or SafeNet support. SafeNet support operates 24 hours a day, 7 days a week. Your level of access to this service is governed by the support plan arrangements made between SafeNet and your organization. Please consult this support plan for further information about your entitlements, including the hours when telephone support is available to you.

Table 1: Technical support contacts

Contact method	Contact	
Address	SafeNet, Inc. 4690 Millennium Drive Belcamp, Maryland 21017 USA	
Phone	United States	(800) 545-6608, (410) 931-7520
	Australia and New Zealand	+1 410-931-7520
	China	(86) 10 8851 9191
	France	0825 341000
	Germany	01803 7246269
	India	+1 410-931-7520
	United Kingdom	0870 7529200, +1 410-931-7520
Web	www.safenet-inc.com	
Support and Downloads	www.safenet-inc.com/support Provides access to the SafeNet Knowledge Base and quick downloads for various products.	
Customer Technical Support Portal	https://serviceportal.safenet-inc.com Existing customers with a Customer Connection Center account, or a Service Portal account, can log in to manage incidents, get the latest software upgrades, and access the SafeNet Knowledge Base.	

 Copyright 2014 SafeNet, Inc. All rights reserved.

CHAPTER 1
The Luna HSM Product Line

This chapter provides an overview of the Luna HSM product line. It contains the following sections:

Luna HSM Products - Overview

SafeNet Luna HSMs are hardware security modules designed to protect critical cryptographic keys and to accelerate sensitive cryptographic operations across a wide range of security applications. All Luna HSMs enable separation of roles by distinguishing between the HSM Security officer space (an administrative function) and the HSM Partition or User space, where client keys and objects are secured, and where client-invoked cryptographic operations take place. Luna HSMs fall into three categories:

- Luna PCI-E is a card-type HSM that installs into the PCIe slot(s) of a host computer. Multiple Luna PCI-E HSMs can coexist in one host system. Each Luna PCI-E HSM supports one HSM partition. See "About Luna PCI-E" on page 14.
- Luna G5 is a desktop HSM unit that connects locally to a host computer via USB interface. Multiple Luna G5 HSMs can be linked via USB connection. Each Luna G5 HSM supports one HSM partition. See "About Luna G5" on page 18.
- Luna SA is a self-contained, network attached HSM appliance, containing an HSM card similar to Luna PCI-E, and normally resides in an equipment rack in a server room (often of the "lights off", unattended variety), and is accessed remotely via secure administrative and client links. Each Luna SA HSM supports multiple HSM partitions, the number governed by purchased licenses. "About Luna SA" on page 9.

HSM Basics

An HSM is a Hardware Security Module. It has storage, cryptographic, and access-control functions that allow cryptographic operations to be performed and segregated within a secure physical hardware boundary, while offloading such functions from the general-purpose pathways of the host or client. Here are basic elements common to Luna HSMs:

Volatile and non-Volatile Data Storage

Luna HSMs can contain both volatile and non-volatile data.

- Non-volatile data includes identification parameters and data objects (such as keys and certificates) that you wish to store for long-term re-use. Those objects persist on the HSM until you explicitly destroy or overwrite them.

Release 5.4.1 007-011329-006Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

- Volatile data is any data that should not persist when it is not in use. Volatile (or session) data disappears when the HSM loses power, or when a session closes.

Keys and objects are stored under multiple layers of encryption, and are decrypted within the physical bounds of the HSM, only into volatile/session storage, and only while being used.

Initialization

Luna HSMs must be initialized before you can use them for the first time (or after an event, like too many consecutive failed login attempts on the Security Officer (SO) account, that zeroizes the HSM).

Initialization establishes several HSM parameters, including identification and authentication of HSM Security Officer (SO) and HSM Partition User who then have access to create and use HSM/Partition objects (keys, certificates, encrypted data, etc.).

Many applications from PKI and other cryptographic product vendors do not include the capability to initialize a Luna HSM, so SafeNet supplies the Lunacm utility program on all supported platforms, to perform that function and other maintenance functions.

Once a Luna HSM is initialized, no one can access it unless they provide the passwords or keys that unlock that specific HSM or Partition.

You can re-initialize a Luna HSM at any time (as SO). Re-initialization destroys all data on the token.

Authentication methods

Luna HSMs are factory configured to be either:

- Password authenticated - uses typed text strings to access the HSM and authenticate to all roles on the HSM; advantage, greater convenience.
- PED authenticated - uses physical tokens, called PED Keys, mediated by a PIN Entry Device, or PED to access the HSM and authenticate to all roles on the HSM; advantage, greater security.

An HSM in the field cannot be changed from Password-authenticated to PED-authenticated, or from PED-authenticated to Password-authenticated. The only exception is the Luna Backup HSM, which configures itself at the time of a backup operation, to match the authentication scheme of the HSM being backed up - the Backup HSM performs Backup and Restore only, and has no ability to perform cryptographic operations

Historical Note

The product name "Luna" was taken from the name of the Luna moth, to conform with the originating company name "Chrysalis-ITS". The company name was derived from the hidden or secret existence of the moth as it developed within its cocoon, or the chrysalis. This was evocative of the hidden world of cryptography. Other moth names were considered for additional product lines, but the "Luna" brand very quickly achieved marketplace recognition and efforts were aligned under that brand.

After years of growing success with the Luna brand in the crypto markets, Chrysalis-ITS was acquired by SafeNet. Because the brand was well recognized and respected in the HSM marketplace, SafeNet maintained it.

Our SNMP MIB is still called CHRYSALIS.

About Luna SA

The SafeNet Luna SA is an Ethernet-attached HSM (Hardware Security Module) Server designed to protect critical cryptographic keys and to accelerate sensitive cryptographic operations across a wide range of security applications.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Luna SA includes many features that increase security, connectivity, and ease-of-administration in dedicated and shared security applications.

Luna SA comes in one of two model families, according to the level of authentication and access control. Your Luna SA was factory configured to operate as either:

- a Password Authenticated version, equivalent to FIPS 140-2 level 2, using passwords, only, for authentication and access control
- a PED (Trusted Path) Authenticated version, equivalent to FIPS 140-2 level 3, that requires Luna PED and PED Keys for authentication and access control.

Physical Features

The standard appliance is the 1U-high, rack-mount device:

Here are some of the important physical features of the Luna SA appliance.

Front View

First, the front; this illustration shows the appliance with its snap-on decorative bezel removed...

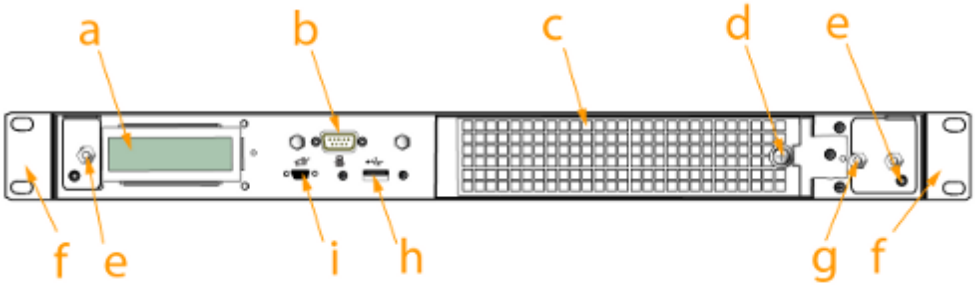

Item	Name	Description
a	LCD system status screen	Shows IP info and scrolls through system status messages
b	Serial (console) port	Local connection for initial setup, and for admin account reset (local-only action for security reasons)

 Copyright 2014 SafeNet, Inc. All rights reserved.

Item	Name	Description
c	Ventilation-fan filter cover	Removable bracket allows cleaning of air filter
d	Fan filter cover retaining screw	A captive thumb-screw (no tool needed).
e	Mounts for removable front bezel	The decorative/protective front bezel mounts on the appliance front panel. Spring clips behind the bezel engage the mounting posts at the left and right ends of the appliance front panel.
f	Rack-mount tabs (removable)	Use these on the front, and the sliding tabs toward the rear to support your Luna appliance in a compatible equipment rack
g	Securing screw for fan bay	Torx screw secures the fan bay; opening to swap fan modules triggers a tamper event on the appliance
h	USB port	Use to connect Luna Remote Backup HSM (for backup of your HSM partition contents), Luna G5 HSM, or Luna DOCK 2 (for PKI and for migration of cryptographic material from older backup token HSMs); same as USB port on back panel
i	PED port	Attach Luna PED 2, Pin Entry Device, reads the hardware (iKey) authentication devices for Trusted Path (FIPS 140 level 3) access control

Rear View

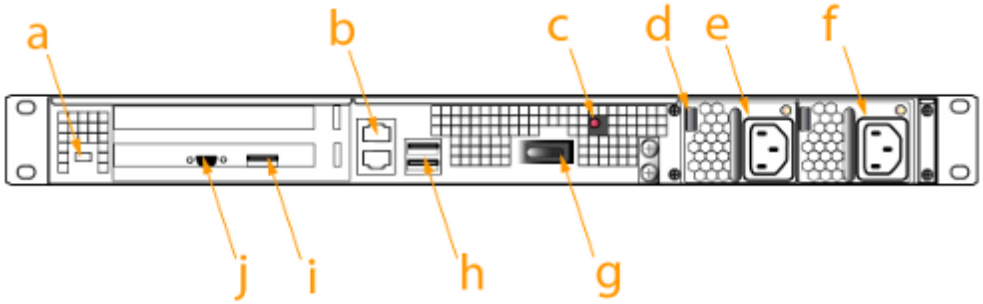

Item	Name	Description
a	Kensington Security Slot	Attach an industry-standard locking cable for additional physical security.
b	Ethernet ports	For network connection of your Luna appliance.
c	Decommissioning button	Recessed for safety; renders HSM contents unusable.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Item	Name	Description
d	Power supply release tab	Press tab to release the catch, and slide the power supply out.
e	Removable power supply	One of two redundant power supplies.
f	Second removable power supply	The other of two redundant power supplies.
g	Start/stop switch	Use to stop the system if the command-line shutdown is not available; use to restart the system if it has been switched off.
h	USB ports	Use to connect Luna Remote Backup HSM (for backup of your HSM partition contents), Luna G5 HSM, or Luna DOCK 2 (for PKI and for migration of cryptographic material from older backup token HSMs); same as USB port on front panel.
i, j	Unused ports	These ports are not used for Luna SA; we recommend that you do not remove the covers that were installed at the factory.

FIPS and Common Criteria Validations and Certifications

At any given time, a FIPS-validated version is available (except for newly introduced products that have not had time to go through the year-long evaluation and validation process), and a newer not-yet-validated version might also be available. The usual practice is to ship units pre-loaded with the firmware and software at the FIPS-validated level by default, while providing the option to update the Client software, Appliance software, and HSM firmware to the newer version. This allows customers who need FIPS validation to have that configuration from the factory, and customers who need newer features (and do not need FIPS validation) to upgrade by simply installing the newer software and following the upgrade procedure.

To check the progress of HSM versions that are submitted for FIPS 140-2 validation visit the NIST site at: (http://csrc.nist.gov/groups/STM/cmvp/validation.html).

Similarly, some versions of product are submitted for Common Criteria EAL evaluation.

You can also check SafeNet Sales or SafeNet Customer Support to inquire about certification status of SafeNet HSM products. If FIPS validation or CC EAL certification are not requirements for you, then the newest version is normally the preferred option.

SafeNet HSM Cryptographic Engine

The SafeNet HSM's integrated SafeNet-Luna Cryptographic Engine is used to perform cryptographic operations and provide secure storage for sensitive cryptographic keys.

The SafeNet-Luna Cryptographic Engine enables the Luna SA functionality by providing:

- secure cryptographic storage,
- cryptographic acceleration (up to 7000 1024-bit RSA signings per second),
- administrative access control and
- policy management.

 Copyright 2014 SafeNet, Inc. All rights reserved.

The SafeNet-Luna Cryptographic Engine can also be used in conjunction with the optional Trusted Path Authentication feature to provide FIPS 140-2 Level 3 validated HSM operation. That option is factory-configured and not subject to change in the field.

The Luna SA HSM Appliance

HSMs, in general, are designed to provide dedicated cryptographic functionality, including key generation, key storage, and digital signing, on a one-to-one basis to their host applications. For example, a database server using an HSM would require one HSM, while a secure website using SSL on the same network would require a second, separate HSM. As the number of secure applications requiring an HSM grows, so does the number of ordinary HSMs deployed.

Luna SA bypasses this limitation by implementing multiple virtual HSMs, or HSM Partitions on a single HSM server. Partitions are accessed via a Network Trust Link.

The following block diagram is a conceptual overview of the Luna SA HSM Server depicting internal systems, communications, and interaction with application servers.

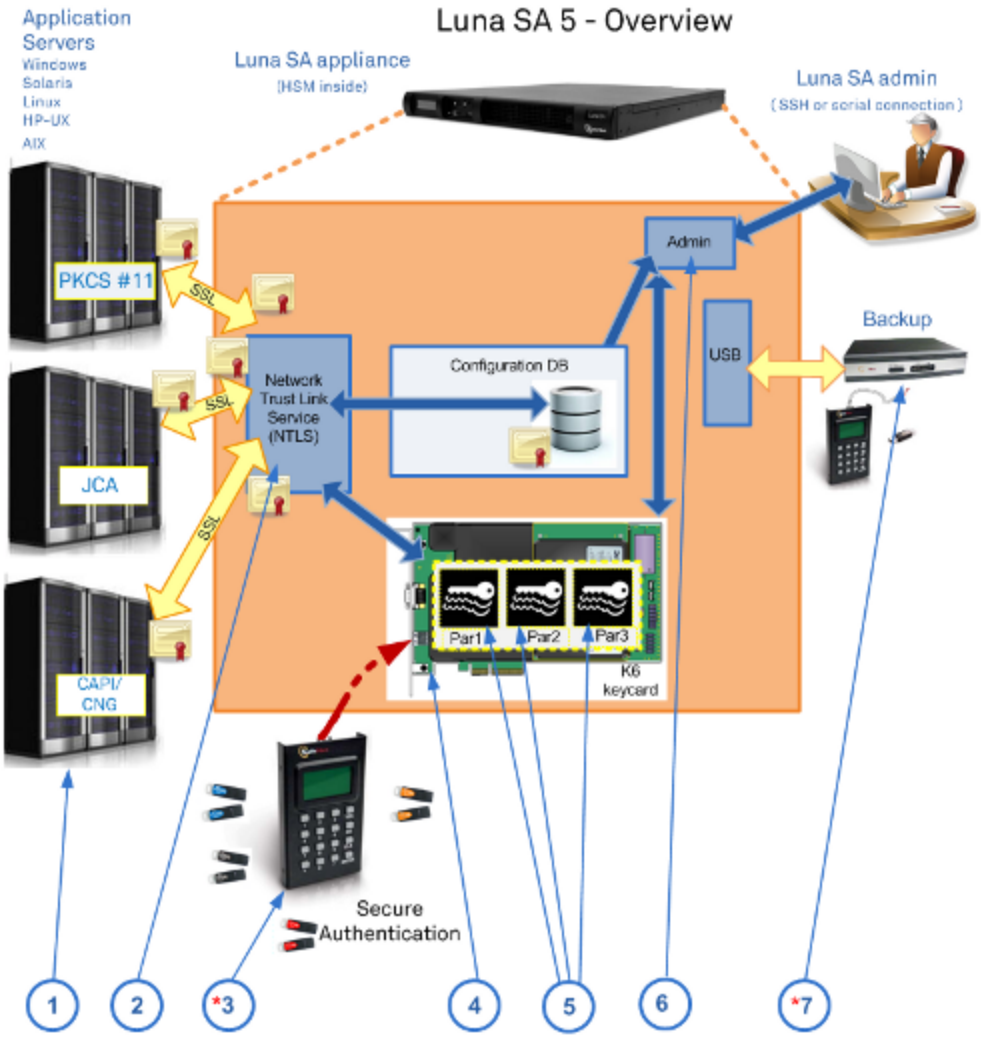

Luna SA operations encompass seven major elements. Some of these elements are optional configuration items, and might not be present in your system:

 Copyright 2014 SafeNet, Inc. All rights reserved.

1.
2. Network Trust Link
3. PED (trusted path) authentication
4. SafeNet K6 HSM Cryptographic Engine
5. HSM Partitions
6. Secure command line interface
7. Secure backup HSM

(* The Secure Backup HSM, and Luna PED (Trusted Path Authentication and Access Control) are options that might not be included with your system.)

About Luna PCI-E

The Luna HSM Customer documentation uses "Luna PCI-E" whenever it refers to either of the performance versions - Luna PCI-E 1700 or Luna PCI-E 7000, without need to specifically identify one version. Those two versions are so-named because their tested performance at repetitive RSA 1024-bit signings per second (under laboratory conditions was near one or the other of those numbers (1700 or 7000).

1024-bit RSA keys are actually outdated for most applications, due to their small size. However 1024-bit RSA signing has been an industry-standard way to convey application and HSM performance for many years and will continue to be used until an industry consensus is reached for an updated indicator.

HSM Basics

An HSM is a Hardware Security Module. An HSM stores cryptographic objects (keys, certificates, etc.), creates and destroys crypto objects, and performs cryptographic operations (encrypt, decrypt, sign, verify, wrap, unwrap) using those objects within the secure physical confines of the HSM - not exposed on a computer file system. The HSM also controls access to its contents and its functions.

The Luna PCI-E Cryptographic Module is an HSM. Here are the basic elements common to Luna HSMs:

Volatile and non-Volatile Data Storage

Luna HSMs can contain both volatile and non-volatile data:

- Non-volatile data includes identification parameters and data objects (such as keys and certificates) that you wish to store for long-term re-use. Those objects persist on the HSM until you explicitly destroy or overwrite them.
- Volatile data is any data that should not persist when it is not in use. Volatile (or session) data disappears when the HSM loses power.

The Card

The Luna PCI-E 5 [K6] HSM card is designed to the PCIe 1.1 standard, for use in PCIe x4 slots. The HSM card can be used in larger connector slots (x16).

Some x16 slots are intended by the computer motherboard manufacturer to be used for video cards, and might not work correctly with Luna PCI-E 5. The symptom is that, at startup, the system detects a card in the slot, but the card does not respond as a video card, and so the system stops booting. This could happen to any non-video PCIe card inserted in such a slot. If you encounter a problem, try another available slot. Modern motherboards tend to support PCIe 2.0 standard, which is backward compatible with 1.1, when correctly implemented.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Of the three major vendors of PCI bridge chips (including the one that we used), each has known problems either of performance, compatibility, or both. Due to the variety of systems and component combinations in the market, we are unable to test with all possible platforms. At the time that this Help was written we found greater incompatibility among server systems than among desktop/workstation systems. If you encounter a problem that is not solved by moving the Luna PCI-E 5 card, contact SafeNet Technical Support -- e-mail: support@safenet-inc.com or phone 800-545-6608 (+1 410-931-7520 International).

Partition

Luna PCI-E is a versatile HSM capable of many roles. Part of that versatility is achieved by separating HSM management (the Security Officer or HSM Admin space) from HSM operation (the User or client). This is achieved by means of the HSM partition or virtual HSM within the physical HSM.

The owner of the partition:

- can see and manage the contents of the partition, and
- can enable or disable access by client applications as desired, entirely separately from the overall HSM management performed by the SO.

The SO:

- can perform HSM updates/upgrades,
- can modify operating parameters
- can deal with tamper events,
- can create or destroy a partition, reset the authentication of an existing partition (when someone forgot his password or lost his PED Key, or someone has left the organization ... or been fired...),
- can authorize the creation of a partition challenge secret, and
- can perform other global operations without ever being able to see or touch the contents of the User/Owner's partition.

The roles are kept separate.

Initialization

Luna HSMs must be initialized before you can use them for the first time (or after an event, like too many consecutive failed login attempts on the SO account, which zeroizes the HSM). Initialization establishes several HSM parameters, including identification and authentication of HSM Security Officer (SO) and HSM Partition User who then have access to create and use HSM/Partition objects (keys, certificates, encrypted data, etc.). Once a Luna HSM is initialized, no one can access it unless they provide the passwords or keys that unlock that specific HSM or Partition. Initialization is meant to be performed only once on an HSM, and it erases any Authentication Data, and data or token objects contained on the token. Once the HSM is in use, be sure to avoid mistakenly initializing it again.

You can re-initialize a Luna HSM at any time (as SO). Re-initialization destroys all data on the token.

Note: On the other hand, until you put Luna PCI-E into service with actual production data, keys, and certificates on it, you can reinitialize it and practice with a variety of optional settings, as many times as you wish.

Many applications from PKI and other cryptographic product vendors do not include the capability to initialize a Luna HSM, so SafeNet supplies the Lunacm utility program on all supported platforms, to perform that function and other maintenance functions.

 Copyright 2014 SafeNet, Inc. All rights reserved.

Your Luna PCI-E Cryptographic Module or HSM is shipped in a pre-initialized state, as part of the factory quality assurance process. However, in that state the HSM is not associated with Security Officer [SO] or User Authentication Data, and is not ready to receive or to create and store objects. You must perform a one-time initialization procedure with the lunacm utility before the HSM can operate with an application program.

Sessions and Authentication

When you access a Luna HSM or HSM Partition, by providing the passwords (Password Authentication versions) or the PED Keys (Trusted Path Authentication versions), you open a session. That session remains open until you (or your application) explicitly close it.

Removing power from a Luna HSM immediately closes all sessions and causes all volatile data to disappear.

Your application program might not be capable of logging into Luna PCI-E, on its own. If not, then the program expects to encounter the HSM already in a logged-in state. For Luna PCI-E to operate with such an application, you must log into the token with the lunacm utility. Find it in your LunaPCI directory, following installation. Every time you reboot your computer, you are required to log into the HSM with the lunacm utility before you can resume using Luna PCI-E with your application, unless you have invoked Autoactivation.

Luna PCI-E Physical Appearance and Features

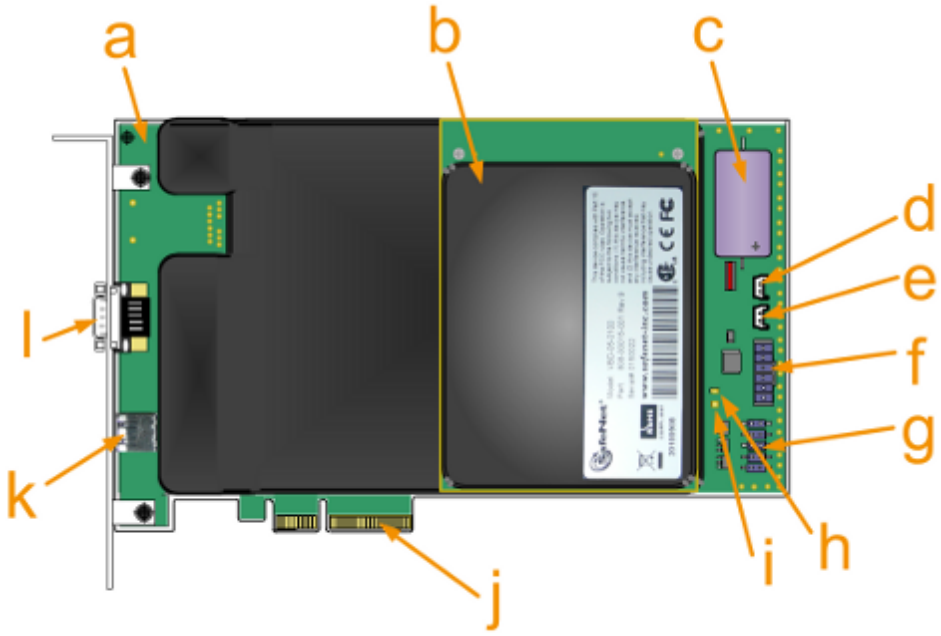

Feature	Description
a	K6 main board
b	Daughter board
c	Battery for Real Time Clock (RTC) and NVRAM
d	Header for Tamper2 (indicated as JP3 on the board), or the "decommission" circuit - closing/shunting

 Copyright 2014 SafeNet, Inc. All rights reserved.

Feature	Description
	those pins causes the KEK and any cached data to be erased. [If used, this is intended to be wired to a normally-open switch, accessible outside the host computer. Such a switch should be shielded/shrouded to prevent accidental activation. To ship the Luna PCI-E 5.x HSM to SafeNet (or other recipient) with assurance that your crypto objects cannot be recovered by anyone, just shunt this header momentarily, or touch a screw-driver blade to both pins simultaneously - the "decommission" action occurs instantly.]
e	Header for Tamper1 (indicated as JP2 on the board), the physical tamper circuit - closing/shunting those pins, or closing a connected switch causes a tamper event and destroys the MTK, the Master Tamper Key that encrypts everything on the HSM. [If used, this pin pair would usually be wired to a chassis switch that is held open when the lid or panel is in place. Opening the lid or panel would close the switch and tamper the HSM.]
f	Serial Connector - not for customer use
g	PED port - same as the externally available PED port "m", below
h	Indicator LED D1 [ERROR] - glows red when the HSM is in an error state or system HALT [when the HSM senses a tamper of any type, or upon startup if the HSM cannot initialize the dualport communication between itself and the host computer]
i	Indicator LED D2 [ACTIVE] - glows or flickers green when the HSM is active
j	PCIe x4 card-edge connector - can be inserted in any PCIe 4-channel (or larger) socket
k	USB connector (for connection to backup HSM or a Luna DOCK 2 reader)
l	PED Port - connect a SafeNet Luna PED 2 PIN Entry Device, reads and imprints iKey PED Keys (a "something you have" authentication factor) that carry primary authentication for the HSM and HSM partitions; also provides a keypad interface for PED Key operation and for additional, optional "something you know" authentication factor], Use a SafeNet-supplied PED cable

Developing a security plan and associated procedures

Not every application environment will require rigorous security and paper-trail management, with respect to HSMs and their contents. However, in high-security environments where security and process auditing is mandated, you might be required to refer to a history of any sensitive materials and any systems associated with them -- who had access, what did they do, and when did they do it. Rehearse everyday operational activities, as well as maintenance and update activities (Authentication Data [password] update cycles, personnel changes, backups, logging) before implementation in your live environment.

Have all secure physical storage sites and all the related handling procedures prepared in advance. Log your receipt of the Luna hardware and then log all storage and handling events thereafter. In an operational environment, you should be able to refer back to a complete “paper trail” – an unbroken record that tracks the existence, storage, handling, and all transitions/hand-offs experienced by each token/HSM that you ever use. Once you take possession, never allow yourself or your organization to lose track, even briefly, of any of your HSMs. If your environment includes auditing, your security auditors will require such a record.

 Copyright 2014 SafeNet, Inc. All rights reserved.

About Luna G5

Your Luna G5 Cryptographic Module or HSM is shipped in a pre-initialized state, as part of the factory quality assurance process. However, in that state the HSM is not associated with Security Officer [SO] or User Authentication Data, and is not ready to receive or to create and store objects. You must perform a one-time initialization procedure with the `lunacm` utility before the HSM can operate with an application program.

Note: Initialization is meant to be performed only once on an HSM, and it erases any Authentication Data, data or token objects contained on the HSM. Once the HSM is in use, be sure to avoid mistakenly initializing it again. On the other hand, until you put the Luna G5 into service with actual production data, keys and certificates on it, you can reinitialize it and practice with a variety of optional settings, as many times as you wish.

Sessions and Authentication

When you access a Luna HSM or HSM Partition, by providing the passwords (Password Authentication versions) or the PED Keys (PED Authentication versions), you open a session. That session remains open until you (or your application) explicitly close it.

Removing power from a Luna HSM immediately closes all sessions and causes all volatile data to disappear.

Your application program might not be capable of logging into the Luna G5, on its own. If not, then the program expects to encounter the HSM already in a logged-in state. For the Luna G5 HSM to operate with such an application, you must log into the HSM or its User partition (sometimes referred to as a "token" in some cryptography documentation and discussions) using the `lunacm` utility. Find it in your LunaClient directory, following installation. Every time you reboot your computer, you are required to log into the HSM with the `lunacm` utility before you can resume using the Luna G5 with your application, unless the application is Luna HSM-aware.

Luna G5 as Encryption/Signing HSM or as RA HSM

The Luna G5 HSM is shipped in different configurations. The Password Authenticated version can be factory configured as an Encryption and Signature HSM (token) or as a Registration Authority (RA) HSM. An RA HSM has the same capabilities as an Encryption and Signature HSM, with the additional ability to wrap private keys off the token for use by smart cards and other applications where multiple secure key generation and issuance is required.

 Copyright 2014 SafeNet, Inc. All rights reserved.

Develop a security plan and associated procedures

Not every application environment will require rigorous security and paper-trail management, with respect to HSMs and their contents. However, in high-security environments where security and process auditing is mandated, you may be required to refer to a history of any sensitive materials and any systems associated with them -- who had access, what did they do, and when. Rehearse everyday operational activities, as well as maintenance and update activities (Authentication Data [password] update cycles, personnel changes, backups, logging) before implementation in your live environment.

Have all secure physical storage sites and all the related handling procedures prepared in advance. Log your receipt of the Luna hardware and then log all storage and handling events thereafter. In an operational environment, you should be able to refer back to a complete “paper trail” – an unbroken record that tracks the existence, storage, handling, and all transitions/hand-offs experienced by each HSM that you ever use. Once you take possession, never allow yourself or your organization to lose track, even briefly, of any of your HSMs or authentication devices (PED Keys, for PED-authenticated HSMs).

If you don't know where a PED Key is, you are not in control of it. If you don't know where it has been, you cannot assert that it has not been copied. If this is ever in doubt, consider resetting or changing passwords/PED Keys. Partition authentication (password, black PED Key if applicable) can be reset with resetPw. Partition or HSM authentication can be changed with changePw. Consider exercising these options if there is any chance an HSM's authentication might have been compromised.

Password integrity is as secure as your personnel choose to keep those passwords.

Physical authentication devices (PED Keys) are as secure as your security policies and procedures and their enforcement.

About Luna Backup HSM

The Luna Backup HSM is physically similar to the Luna G5 HSM, but is used exclusively to securely backup sensitive material from Luna HSMs, and to restore backed-up material to Luna HSMs. Some important characteristics are:

- The Luna Backup HSM can be connected locally, by USB cable, to the primary HSM, or it can be connected to a server and used to backup from, and restore to, remotely located primary HSMs.
- The Luna Backup HSM takes on the authentication type of the primary HSM with which it is paired for backup - so it becomes a Password-authenticated Backup HSM (sometimes called the FIPS 140-2 level 2 version) when backing up a Password-authenticated primary HSM, and the same Luna Backup HSM becomes a PED-authenticated Backup HSM (sometimes called the FIPS 140-2 level 3 version) when backing up a PED-authenticated primary HSM.
- The Luna Backup HSM performs backup and restore operations only; it is not capable of cryptographic operations, and cannot (for example) be substituted for a Luna G5 HSM.

Note: When the Luna Backup HSM contains backup data, and has therefore taken on the authentication characteristics of either a Password-authenticated or a PED authenticated HSM, it cannot restore to the other type. This is a security feature. PED-authenticated-to-Password-authenticated is prevented, because keys and objects that were created on a PED-authenticated HSM are more secure, and moving them to a less-secure type of HSM would be considered a breach of security. Password-authenticated-to-PED-authenticated is prevented because anyone seeing keys and objects on a PED-authenticated HSM is entitled to assume that those keys and objects have always had that level of security throughout their existence.

 Copyright 2014 SafeNet, Inc. All rights reserved.

CHAPTER 2

Luna HSM Authentication Types

This chapter describes the types of authentication available on Luna HSMs. Each Luna HSM comes in one of two authentication types – Password authenticated or PED authenticated. The authentication type is configured at the factory and cannot be modified in the field. See the following sections for more information:

- "About Password Authentication" on page 20.
- "About PED Authentication" on page 22.
- "Comparing Password and PED Authentication" on page 25
- "About Remote PED" on page 26

About Password Authentication

This section applies to versions of Luna HSM that control access via typed text-string authentication, or passwords, at all authentication levels. For Luna HSMs, this is sometimes referred to as "FIPS 140-2 Level 2" or simply "FIPS Level 2" or "FIPS 2" authentication.

If you received a Luna PED and PED Keys, then your Luna appliance's HSM probably uses Trusted Path Authentication, and **not Password Authentication** (verify with the `hsm displayLicenses` command), and this page does not apply to you. We also can refer to that version as "FIPS 140-2 Level 3" authentication. See "About Trusted Path Authentication", instead.

In general, there are two paths to access the Luna appliance and its HSM:

- the administrative path, via SSH or via local serial link, which uses the `lunash` command-line interface
- the Client path, via SSL, by which client applications use the Luna SA API to perform cryptographic functions within pre-assigned virtual HSMs (called Partitions) on the Luna system.

For Luna HSMs with Password Authentication, the various, layered roles are protected by passwords:

Role	Description
Appliance Admin	When you login to the Luna appliance via **lunash** the only accepted ID is "admin" which requires the admin password. As the appliance admin, you can connect and login locally, via a serial terminal, or remotely via SSH. With no other authentication, admin can perform general, appliance-level administration.
HSM Admin	To access the HSM to perform HSM-specific administration tasks (set HSM-wide policies, update firmware and capabilities, backup and restore the HSM, create and remove HSM Partitions, etc.), you must be logged in to **lunash** as admin, then you must further be logged in as HSM Admin (of which there can be only one per Luna HSM) . Good security practices suggest that the HSM Admin password should be different from the appliance admin password. However, your corporate policies may differ. As the HSM Admin, you can connect locally, via a serial terminal, or remotely via SSH – you must first be logged in as admin to have access to **lunash** commands.

Release 5.4.1 007-011329-006Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Role	Description
Partition Owner	To access HSM Partitions, in order to perform Partition-specific administration tasks (set Partition-specific policies, assign Partition to Clients, revoke Clients, etc.), you must be logged in to **lunash** as admin, then you must further be logged in as Partition Owner (of which there can be several -- one for each Partition in the HSM) , using the Partition Password. Good security practices suggest that the Partition Password should be different from the appliance admin password, different than the HSM Admin password, and different than other Partition Passwords (for other Partitions). However, your corporate policies may differ. As the Partition Owner, you can connect locally, via a serial terminal, or remotely via SSH – you must first be logged in as admin to have access to **lunash** commands.
Client	To access HSM Partitions with an application to perform cryptographic operations on data, you must connect remotely via SSL (called NTLS in our implementation) as a Client (one that has been registered by certificate exchange and assigned by the Partition Owner to this Partition) , then pass a User-type (this is done invisibly by your client application), and present the Partition Password (also done automatically by your application). The password used by a Client is the same Partition Password that is used by the Partition Owner for the particular Partition. What limits the scope of operations that a registered, authenticated Client can perform on a Partition is the fact that Partition administrative commands can be issued only via **lunash**. Thus, for security, Clients must not be allowed to learn the appliance admin password that gives access to **lunash**.

Summary

Objects on the HSM are encrypted by the owner of the HSM Admin space or of the User space (partition), and can be decrypted and accessed only by means of the specific secret (password) imparted by the HSM Admin or the partition User respectively.

If you cannot present the secret (the password) that encrypted the objects, then the HSM is just a secure storage device to which you have no access, and those objects might as well not exist.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

About PED Authentication

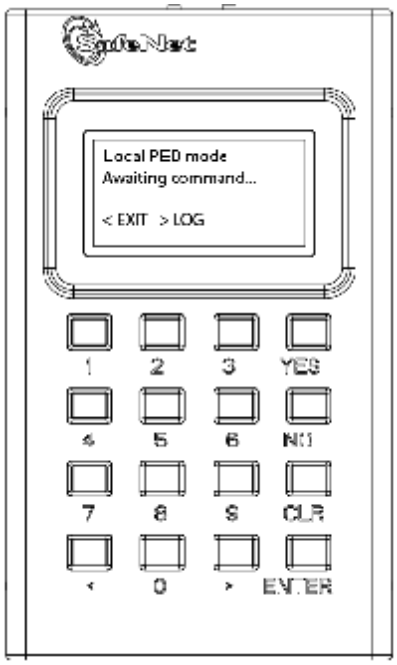

This section applies to versions of SafeNet HSM that control access via Trusted Path Authentication - that is, HSMs that control access by means of the PED and PED Keys, rather than by typed-in text strings. For Luna HSMs, this is sometimes referred to as "FIPS 140-2 Level 3" or simply "FIPS Level 3" or "FIPS 3" authentication.

Note: If you did not receive a Luna PED and PED Keys, then your Luna HSM probably uses Password Authentication, and not Trusted Path Authentication (verify with the `hsm displayLicenses` command), and the pages in this section do not apply to you. See "About Password Authentication" on page 20, instead.

You can also verify the type of a Luna HSM by running the `hsm showPolicies` command. The output includes these lines near the top:

```
Description                          Value
===========                          =====
Enable PIN-based authentication Disallowed
Enable PED-based authentication     Allowed
```

The above result is from a PED-authenticated HSM.
A Password-authenticated HSM would show:

```
Description                          Value
===========                          =====
Enable PIN-based authentication      Allowed
Enable PED-based authentication      Disallowed
```

 Copyright 2014 SafeNet, Inc. All rights reserved.

PED Connections

The Trusted Path is the connection between the Luna PED and the Luna HSM.

- For Luna SA, the PED connection is on the appliance front panel.
- For Luna PCI-E, the PED connection is a slot-edge connector, directly on the HSM card, accessible at the exterior of a tower or server computer (not through the host computer).
- For Luna G5, the PED connection is an external connection to the device (not through the host computer).

For local PED, the connection is a secure physical link, directly to the HSM, bypassing the computer memory and bus. For Remote PED, the connection is a cryptographically secured link across the network - when credentials travel between PED and HSM, they are encrypted throughout the journey. At no time does an authentication secret exist in-clear, anywhere in computer memory or on any computer bus.

In general, there are three paths to access the Luna HSM:

- the administrative path, via SSH or via local serial link, which uses the `lunash` command-line interface
- the Client path, via TLS (our implementation is called NTLS), by which client applications use the SafeNet HSM API to perform cryptographic functions within pre-assigned virtual HSMs (called Partitions) on the HSM
- the Trusted Path, used for authentication data passed from the PED and PED Keys - this path ensures that HSM authentication data does not pass unencrypted through a host or terminal computer, where it might be subject to attack.

Roles

For SafeNet HSM with Trusted Path Authentication, the various layered roles are protected by a combination of PED Keys and passwords:

Appliance Admin (Luna SA only)

When you login to the Luna appliance via lunash the accepted IDs are "admin" which requires the admin password, "operator", which requires the operator password, or "monitor" which requires the monitor password. (Named users can later be added with admin, operator, or monitor authority.) The password is typed at the command line (operator and monitor are restricted identities that have access to subsets of the lunash command set used by admin).

As the appliance admin, you can connect and log in locally, via a serial terminal, or remotely via SSH. With no further authentication, admin can perform general, appliance-level administration (not accessing the HSM), and can run view/list/show/display commands on the HSM that do not make changes.

Admin sees the full available command set, while operator- and monitor-level users see only subsets.

If any administrative user attempts an HSM command that needs authentication, the interface prompts for that authentication. On PED-authenticated systems, you are directed to the PED, which prompts for PED Keys and keypad actions.

EXCEPTION: You can also log in through the local (serial link) console connection as an identity called "recover" (password "PASSWORD").

HSM Admin or Security Officer

To access the HSM to perform HSM-specific administration tasks (set HSM-wide policies, update firmware and capabilities, backup and restore the HSM, create and remove HSM Partitions, etc.), you must first be authenticated as SO or HSM Admin (of which there can be only one per Luna HSM).

The authentication data for SO/HSM Admin is not a password. It is a secret carried on a blue PED Key.

 Copyright 2014 SafeNet, Inc. All rights reserved.

For the SO to login and issue HSM commands, someone must be present at the connected local Luna PED, or at the configured Remote Luna PED, to insert the required blue PED Key, when prompted. Otherwise, HSM commands cannot be used.

Thus, anyone wishing to issue HSM-wide administrative commands to the Luna appliance must be present in the room with the Luna PED, and must have the cooperation of the SO/HSM Admin blue PED Key holder (who, in turn, needs physical access to the connected Luna PED).

The options are to perform such authentication via a PED connected physically to the HSM appliance, or to perform authentication via a PED connecting through a secured Remote PED connection.

Partition User (or Crypto Officer)

To access HSM Partitions to perform Partition-specific administration tasks, such as

- set Partition-specific policies
- assign Partition to Clients
- revoke Clients, etc.

You must be authenticated as Partition User (of which there can be one per HSM on Luna PCI-E or Luna G5, or there can be several on Luna SA -- one for each Partition in the HSM), and for that you use the Partition User black PED Key.

The authentication data for Partition User (also known as Crypto Officer in some security and authentication schemes) is both a password and a secret carried on a black PED Key. As the Partition User/Crypto Officer, you can connect locally, via a serial terminal, or remotely via SSH. To perform Partition administration on Luna SA, you must first be logged in as admin to have access to **lunash** commands.

For Luna PCI-E and Luna G5, you simply need access to the host computer, where you can use **lunacm** commands. For the Partition User/Crypto Officer to login and issue Partition administration commands, someone must be present at the connected Luna PED (or the configured and validated Remote PED) to insert the required black PED Key, when prompted or the Partition must have been left in Activated state. Otherwise, Partition administration commands cannot be used.

If you have invoked the Crypto Officer/Crypto User distinction, then there are two Partition Passwords, but only the Crypto Officer password allows you to run **lunash** or **lunacm** commands to administer the Partition. The Crypto User password allows only a limited set of cryptographic activities via a Client application.

For Luna SA, good security practices suggest that the Partition Password should be different than the appliance admin password and different than other Partition Passwords (for other Partitions). If Crypto Officer/Crypto User are in force, then their passwords should differ as well. However, your corporate policies might vary.

Client (or Crypto User)

To access HSM Partitions with an application to perform cryptographic operations on data, (for Luna SA only, requires that you connect remotely via SSL as a Client that has been registered by certificate exchange and assigned by the Partition User to this Partition), you must pass a User-type (this is done invisibly by your client application), and present the Partition Password (also done automatically by your application).

At this point, the two models diverge:

- For a standard "Client", the password is the same Partition Password that is used by the Partition User for the particular Partition. What limits the scope of operations that a registered, authenticated Client can perform on a Partition on Luna SA is the fact that Partition administrative commands can be issued only via `lunash`. Thus, for security, Clients should not be allowed to learn the appliance admin password (for Luna SA) that gives access to `lunash` command line. For Luna PCI-E and Luna G5, the password or other authentication that gives access the client application (that uses the HSM for crypto operations) is often the same authentication that gives access to

 Copyright 2014 SafeNet, Inc. All rights reserved.

`lunacm` for partition administration, so the ability to keep roles separate is more dependent on control of PED Keys.

- For a Crypto User client, the password is different from the Crypto Officer password, offering another layer of protection for the Partition and its contents.

Summary

Objects on the HSM are encrypted by the owner of the HSM Admin space [rarely] or of the User space (partition), and can be decrypted and accessed only by means of the specific secret injected from the blue PED Key (HSM Admin) or the black PED Key (User) respectively.

If you cannot present the secret (the PED Key) that encrypted the objects, then the HSM is just a secure storage device to which you have no access, and those objects might as well not exist.

Audit

Not mentioned above is the Auditor. This role combines a special, limited-access appliance account, and a special HSM role (authenticated by the white PED Key), for the purpose of managing HSM audit logs. These roles are distinct and separate from other roles on the appliance and the HSM, conforming to the requirements of auditing standards.

Using Luna PED Remotely

By default, Luna PED is connected locally, and powered by the HSM using one cable. However, Luna PED can also be used remotely from the HSM or HSMs for which it manages access control. See "About Remote PED" on page 26.

Comparing Password and PED Authentication

The following table outlines the key differences between PED and password authentication.

Feature	Password-authenticated HSM	PED-authenticated HSM
Ability to restrict access to cryptographic keys	• knowledge of Partition Password is sufficient • for backup/restore, knowledge of partition domain password is sufficient	• ownership of the black PED Key is mandatory • for backup/restore, ownership of both black and red PED Keys is necessary • the Crypto User role is available to restrict access to usage of keys, with no key management • option to associate a PED PIN (something-you-know) with any PED Key (something you have), imposing a two-factor authentication requirement on any role
Dual Control	• not available	• Mof N (split-knowledge secret sharing) requires "M" different holders of portions of the role secret, in order to authenticate to an HSM role - can be applied to any, all, or none of the administrative and management operations required on the HSM
Key-custodian	• linked to password	• linked to partition password knowledge,

 Copyright 2014 SafeNet, Inc. All rights reserved.

Feature	Password-authenticated HSM	PED-authenticated HSM
responsibility	knowledge, only	• linked to black PED Key(s) ownership
	Roles limited to: • Appliance admin • HSM Admin (SO) • Partition Owner	Available roles: • Appliance admin • HSM Admin (Security Officer) • Domain (Cloning / Token-Backup) • Secure Recovery • Remote PED • Partition Owner (or Crypto Officer) • Crypto User (usage of keys only, no key management) for all roles, two-factor authentication (selectable option) and MofN (selectable option)
Two-factor authentication for remote access	• not available	• Remote PED and orange (Remote PED Vector) PED Key deliver highly secure remote management of HSM, including remote backup

About Remote PED

When it is not convenient to be physically near the host computer that contains a Luna HSM, in order to connect a Luna PED and present required PED Keys, you can operate remotely and securely.

The PED-Authenticated Luna HSM, and one-or-more orange PED Keys are imprinted with a Remote PED Vector (RPV). This can occur at any time before the HSM is deployed, and requires a locally connected PED. All future PED and PED Key interaction can then be accomplished with Luna PED and PED Keys that are physically distant from the HSM, as follows:

- One computer, running a supported OS, hosts the HSM - this could be:
  - a server or tower containing a Luna PCI-E HSM, or
  - a server or other computer with a USB-connected Luna G5 HSM, or
  - a Luna SA HSM appliance
- The HSM host computer must be network attached. HSM administration commands can be input locally, or via remote connection, but the network connection is essential for Remote PED operation
- A second computer (laptop, workstation, server running a supported Windows version) has a Luna PED (Remote Capable) attached via USB, and powered via its included power block.
- The Remote PED host computer must be network attached. The administration of the distant HSM host does not have to come from this Remote PED host computer, but it is usually done that way, since the person handling the PED must coordinate with the person giving commands to the HSM.
- The Remote PED host computer and PED must have the orange Remote PED Key (RPK) available, along with:
  - either blue, black and red (optionally, white and purple, as well) PED Keys that were imprinted with the HSM previously,

 Copyright 2014 SafeNet, Inc. All rights reserved.

  - or blank blue, black. and red (optionally, white and purple) PED Keys that are about to be imprinted along with the HSM.
- The HSM is told to look to a remote PED for its authentication requests.
- The PED host computer has the LunaPED driver installed, and runs the pedserver utility.
- The HSM host computer runs the pedclient utility, and the HSM is told to connect to the Remote PED.
- The Remote PED (via the pedserver) receives the request and prompts for the orange PED Key.
- The Remote PED and the HSM (via the pedclient/pedserver connection) agree that the provided orange PED Key contains the same Remote PED Vector as is imprinted on the HSM, and the secure Remote PED link is established.
- The HSM administrator runs commands on the HSM (on the host computer) via remote desktop or ssh connection.

All future authentication for the HSM can be performed at the Remote PED, with no need for personnel to visit the HSM host, which could be locked away in a lights-off facility on the other side of the world..

 Copyright 2014 SafeNet, Inc. All rights reserved.

CHAPTER 3
Configurations

This chapter introduces some configurations that you can perform with your Luna HSM products, including some that are mandatory in order to make use of your Luna HSM, and some that are optional (and might, or might not, require additional equipment or software) that can enhance the utility and usability of your Luna HSMs.

The various configurations are introduced in the following sections:

High Availability (HA) Configurations

Luna HSM products include availability and scalability capabilities for mission critical applications that require uninterrupted up-time. These features allow the grouping of multiple devices into a single logical group – known as an HA (High Availability) group. When an HA group is defined, cryptographic services remain available to the consuming applications as long as at least one member in the group remains functional and connected to the application server. In addition many cryptographic commands are automatically distributed across the HA group to enable linear performance gains for many applications. The following sections describe these features and the available configuration options in detail to help you understand how best to configure the HA groups for their application and environment.

Overview

The Luna high-availability (HA) and load balancing (LB) functionality is implemented in the HSM client libraries. The HSMs and appliances are not involved and are unaware that they might be configured in an HA group. This allows you to configure HA on a per-application basis. On each application server, define an HA group by first registering the server as normal clients to all the desired HSMs, then use client-side administration commands to define the HA group and set any desired configuration options. You can configure several options including:

- setting automatic or manual recovery mode;
- setting some HSMs as standby members; and
- performing various manual synchronization and recovery operations.

Once defined, the library presents to the application a virtual HSM that is a consolidation of all the physical HSMs in the HA group. From this point on the library distributes operations and automatically synchronizes key material transparently to the application.

Release 5.4.1 007-011329-006Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

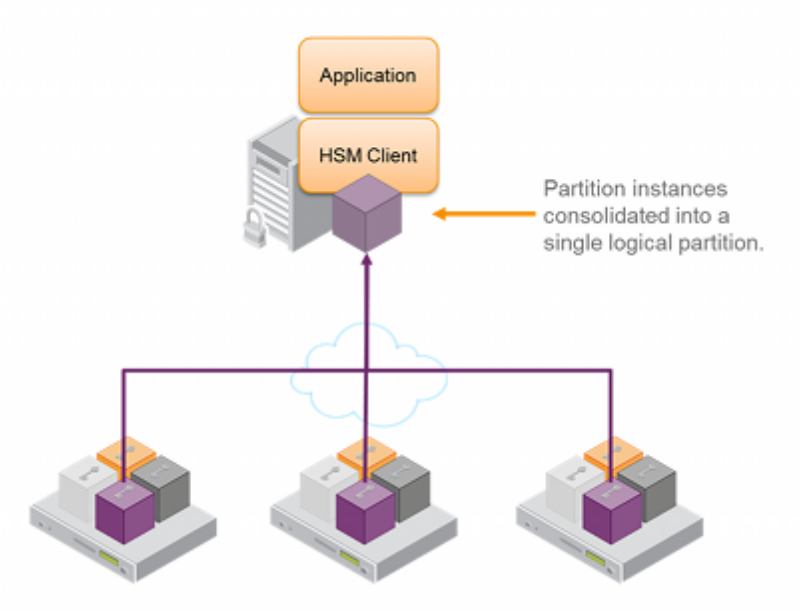

High Availability

The Luna high availability function supports the grouping of up to sixteen members. However, the maximum practical group size for your application is driven by a trade-off between performance and the cost of replicating key material across the entire group. A common practice is to set the group size to N+1 where N is defined by the desired performance per application server(s). As depicted below, this solution gives the desired performance with a single extra HSM providing the availability requirement. The number of HSMs per group of application servers varies based on the application use case but, as depicted, groups of three are typical.

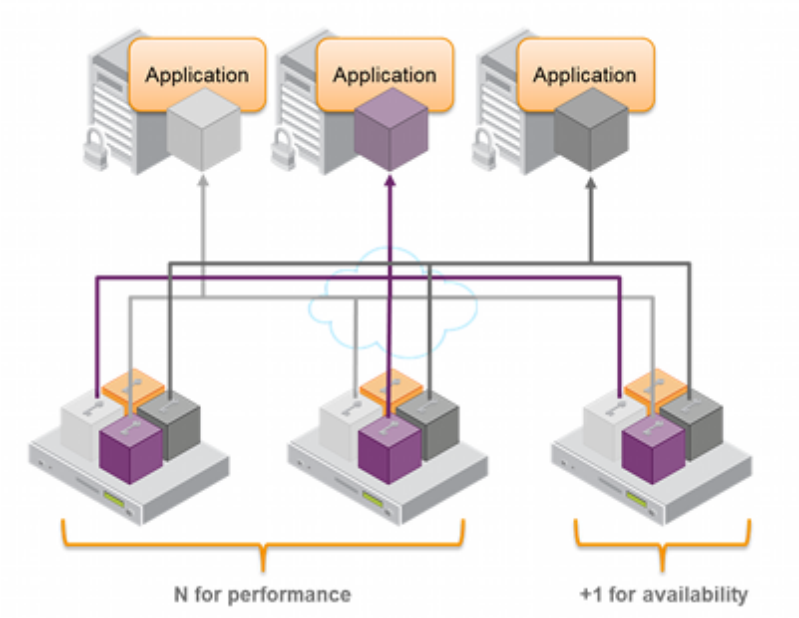

 Copyright 2014 SafeNet, Inc. All rights reserved.

As performance needs grow beyond the performance capacity of three HSMs, it often makes sense to define a second independent group of application servers and HSMs to further isolate applications from any single point of failure. This has the added advantage of facilitating the distribution of HSM and application sets in different data centers.

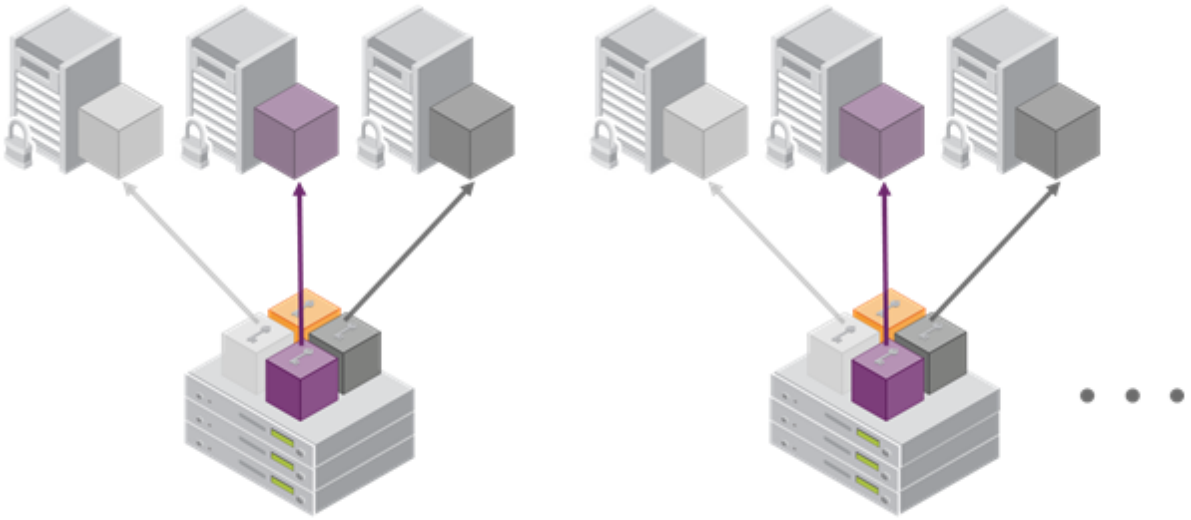

Whenever an application creates key material, the HA functionality transparently replicates the key material to all members of the HA group before reporting back to the application that the new key is ready. The HA library always starts with what it considers its primary HSM (initially the first member defined in an HA group). Once the key is created on the primary it is automatically replicated to each member in the group. If a member fails during this process the key replication to the failed member is aborted after the fail-over time out . If any member is unavailable during the replication process (that is, the unit failed before or during the operation), the HA library keeps track of this and automatically replicates the key when that member rejoins the group . Once the key is replicated on all active members of the HA group a success code is returned to the application.

Load Balancing

The default behavior of the client library is to attempt to load-balance the application's cryptographic requests across the entire set of devices in the HA group. The top level algorithm is a round-robin scheme that is modified to favor the least busy device in the set. As each new command is processed the Luna client looks at how many commands it has scheduled on every device in the group. If all devices have an equal number of outstanding commands the new command is scheduled on the next device in the list – creating a round-robin behavior. However, if the devices have a different number of commands outstanding on them, the new command is scheduled on the device with the fewest commands queued – creating a least-busy behavior. This modified round-robin has the advantage of biasing load away from any device currently performing a lengthy-command. In addition to this least-busy bias, the type of command also affects the scheduling algorithm.

Single-part (stateless) cryptographic operations are load-balanced. However, multi-part (stateful) and key management commands are not load-balanced. Multi-part operations carry cryptographic context across individual commands. The cost of distributing this context to different HA group members is generally greater than the benefit. For this reason multi-part commands are all targeted at the primary . Multi-part operations either are not used or are infrequent actions,

 Copyright 2014 SafeNet, Inc. All rights reserved.

so most applications are not affected by this restriction. Key management commands affect the state of the keys stored in the HSM. As such, these commands are targeted at all HSMs in the group. That is the command is performed on the primary HSM and then the result is replicated to all members in the HA group. Key management operations are also an infrequent occurrence for most applications .

It is important to understand that the least-busy algorithm uses the number of commands outstanding on each device as the indication of its busyness. When an application performs a repeated command set, this method works very well. However, when the pattern is interrupted, the type of command can have an impact. For example, when the HSM is performing signing and an atypical asymmetric key generation request is issued, some number of the application's signing commands are scheduled on the same device (behind the key generation). Commands queued behind the key generation therefore have a large latency driven by the key generation. However, the least-busy characteristic automatically schedules more commands to other devices in the HA group, minimizing the impact of the key generation.

It is also important to note that the load-balancing algorithm operates independently in each application process. Multiple processes on the same client or on different clients do not share their "busyness" information while making their scheduling choice. In most cases this is reasonable, but some mixed use cases might cause certain applications to hog the HSMs.

Finally, when an HA group is shared across many servers, different initial members can be selected while the HA group is being defined on each server. The member first assigned to each group becomes the primary. This approach optimizes an HA group to distribute the key management and/or multi-part cryptographic operation load more equally.

In summary, the load-balancing scheme used by Luna is a combination of round-robin and least-busy for most operations. However, as required, the algorithm adapts to various conditions and use cases so it might not always emulate a round-robin approach.

Failover

When an HA group is running normally the client library continues to schedule commands across all members as described above. The client continuously monitors the health of each member at two different levels. First, the connectivity with the member is monitored at the networking layer. Disruption of the network connection invokes a fail-over event within a twenty second timeout . Second, every command sent to a device is continuously monitored for completion. Any command that fails to complete within twenty seconds also invokes a fail-over event. Most commands are completed within milliseconds. However, some commands can take extended periods to complete – either because the command itself is time-consuming (for example, key generation); or because the device is under extreme load. To cover these events the HSM automatically sends "heartbeats" every two seconds for all commands that have not completed within the first two seconds. The twenty second timer is extended every time one of these heartbeats arrives at client, thus preventing false fail-over events.

A fail-over event involves dropping a device from the available members in the HA group. All commands that were pending on the failed device are transparently rescheduled on the remaining members of the group. So when a failure occurs, the application experiences a latency stall on some of the commands in process (on the failing unit) but otherwise sees no impact on the transaction flow . Note that the least-busy scheduling algorithm automatically minimizes the number of commands that stall on a failing unit during the twenty second timeout.

When the primary unit fails, clients automatically select the next member in the group as the new primary. Any key management or single-part cryptographic operation are transparently restarted on a new group member. In the event that the primary unit fails, any in-progress, multi-part, cryptographic operations must be restarted by the application, as the operation returns an error code.

As long as one HA group member remains functional, cryptographic service is maintained to an application no matter how many other group members fail. As discussed in the Recovery section below, members can also be put back into service without restarting the application.

 Copyright 2014 SafeNet, Inc. All rights reserved.

Recovery

After a failure, the recovery process is typically straight-forward. Depending on the deployment, an automated or manual recovery process might be appropriate. In either case there is no need to restart an application!

With automatic recovery, the client library automatically performs periodic recovery attempts while a member is failed. The frequency of these checks is adjustable and can be limited on the number of re-tries. Each time a reconnection is attempted, one application command experiences a slight delay while the library attempts to recovery. As such, the retry frequency cannot be set any faster than once per minute. Even if a manual recovery process is selected the application does not need to be restarted. Simply run the client recovery command and the recovery logic inside the client makes a recovery attempt the next time the application uses the HSM. As part of recovery any key material created while the member was offline is automatically replicated to the recovered unit .

Sometimes a failure of a device is permanent. In this event, the only solution is to deploy a new member to the group. In this case, remove the failed unit from the HA group, add a new device to the group and then kick the recovery process. The running clients automatically resynchronize keys to the new member and start scheduling operations to it.

Finally, sometimes both an HSM and application fail at the same time. If, while an HSM was offline, no new key material was created the recovery is still straightforward: simply return the HSM to service and then restart the application. However, if new key material was created after an HSM failed but before the application failed, a manual re-synchronization might be required . Confirm which member or members have the current key material (normally the unit (s) that was online at the time the application failed). Put them back in service with the application. Then, for each member that has stale key material (a copy of an object that was deleted; or an old copy of an object who's attributes were changed), delete all their key material after first making sure they are not part of the HA group. Be particularly careful that the member is not part of the HA group or the action might destroy active key material by causing an accidental synchronization during the delete operation! After the HSM is cleared of key material, rejoin it to the group and the synchronization logic automatically repopulates the device's key material from the active units.

Standby Mode

By default all members in an HA group are treated as active. That is, they are kept both current with key material and used to load-balance cryptographic services. In some deployment scenarios it makes sense to define some members as standby. Standby members are registered just like active members except, after they are added to the HA group, they are defined as "standby". As depicted below, applications can be deployed in geographically dispersed locations. In this scenario, use Luna's standby capability to use the HSMs in the remote datacenter to cost effectively improve availability. In this mode, only the local units (non-standby) are used for active load-balancing. However, as key material is created they are automatically replicated to both the active (local) units and standby (remote) unit. In the event of a failure of all local members the standby unit is automatically promoted to active status.. The primary reason for using this feature is to reduce costs while improving reliability and this approach allows remote HSMs that have high latency to be avoided when not needed. However, in the worst case scenario where all the local HSMs fail, the remote member automatically activates itself and keeps the application running.

 Copyright 2014 SafeNet, Inc. All rights reserved.

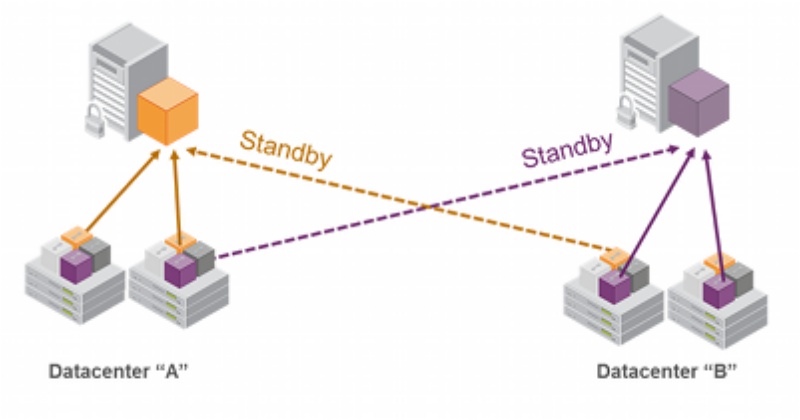

Notes and More

It is important that all members in an HA group have the same configuration and version. Running HA groups with different versions is unsupported. Ensure that HSMs are configured identically to ensure smooth high availability and load balancing operation. Luna HSMs come with various key management configurations: cloning mode, key-export mode or SIM -mode. HA functionality is supported with both cloning and SIM variants – provided all members in the group have the same configuration. Clients automatically and transparently use the correct secure key replication method based on the group's configuration.

It is also critical that all members in an HA group share the same Security Domain role (Red PED key for Trusted Path authentication devices and security domain password for password authenticated devices). The Security Domain defines which HSMs are allowed to share key material. Because HA group members are, by definition, intended to be peers they need to be in the same Security Domain.

By default the client library present both physical slots and virtual slots for the HA group. Directing applications at the physical slots bypasses the high availability and load balancing functionality. An application must be directed at the virtual slots to activate the high availability and load balancing functionality. A configuration setting referred to as HAonly hides the physical slots. SafeNet recommends using this setting to prevent incorrect application configurations. Doing so also simplifies the PKCS #11 slot ordering given a dynamic HA group

Application developers should be aware that the PKCS #11 object handle model is fully virtualized with the Luna HA logic. As such, the application must not assume fixed handle numbers across instances of an application. A handle's value remains consistent for the life of a process; but it might be a different value the next time the application is executed.

The network topography of the HA group is generally not important to the proper functioning of the group. As long as the client has a network path to each member the HA logic will function. Keep in mind that having a varying range of latencies between the client and each HA member causes a command scheduling bias towards the low-latency members. It also implies that commands scheduled on the long-latency devices have a larger overall latency associated with each command. In this case, the command latency is a characteristic of the network; to achieve uniform load distribution ensure that latencies to each device in the group are similar (or use standby mode).

The Luna HA and load-balancing feature works on a per-client and per-partition bases. This provides a lot of flexibility. For example, it is possible to define a different sub-set of HSMs in each client and even in each client's partitions (in the event that a single client uses multiple partitions). SafeNet recommends to avoid these complex configurations and to keep the HA topography uniform for an entire HSM. That is, treat HSM members at the HSM level as atomic and whole. This simplifies the configuration management associated with the HA feature.

When a client is configured to use auto recovery the manual recovery commands must not be used. Invoking them can cause multiple concurrent recovery processes which result in error codes and possible key corruption .

 Copyright 2014 SafeNet, Inc. All rights reserved.

Most customers should enable auto-recovery in all configurations. We anticipate that the only reason you might wish to choose manual recovery is if you do not want to impart the retry time to periodic transactions. That is, each time a recovery is attempted a single application thread experiences an increased latency while the library uses that thread to attempt the re-connection (the latency impact is a few hundred milliseconds).

Example: Database Encryption

This section walks through a specific sample use case of some of the HA logic with a specific application – namely a transparent database encryption.

Typical Database Encryption Key Architecture

Database engines typically use a two-layered key architecture. At the top layer is a master encryption key that is the root of data protection. Losing this key is equivalent to losing the database, so it obviously needs to be highly durable. At the second layer are table keys used to protect table-spaces and/or columns. These table keys are stored with the database as blobs encrypted by the master encryption key. This architecture maps to the following operations on the HSM:

1. Initial generation of master key for each database.
2. Generation and encryption of table keys with the master key.
3. Decryption of table keys when the database needs to access encrypted elements.
4. Generation of new master keys during a re-key and then re-encrypting all table keys with it.
5. Generation and encryption of new table keys for storage in the database (often done in a software module).

The HSM is not involved in the use of table keys. Instead it provides the strong protection of the MEK which is used to protect the table keys. Users must follow backup procedures to ensure their MEK is as durable as the database itself. Refer to the backup section of this manual for proper backup procedures.

HSM High Availability with Database Encryption

When the HSMs are configured as an HA group the database's master key is automatically and transparently replicated to all the members when the key is created; and each time it is re-keyed. If an HSM group member was offline or fails during the replication it does not immediately receive a copy of the key. Instead the HA group proceeds after replicating to all of the active members. Once a member is re-joined to the group the HSM client automatically replicates the new master keys to the recovered member.

With this in mind, before every re-key event the user should ensure the HA group has sufficient redundancy. A re-key will succeed so long as one HA group member exists, but proceeding with too few HSMs will result in an availability risk. For example, proceeding with only one HSM means the new master key will be at risk since it exists only on a single HSM. Even with sufficient redundancy, SafeNet recommends maintaining an offline backup of a database's master key.

HSM Load Balancing with Database Encryption

While a database is up and running the master key exists on all members in the HA group. As such, requests to encrypt or decrypt table keys are distributed across the entire group. So the load-balancing feature is able to deliver improved performance and scalability when the database requires a large number of accesses to the table keys. With that said, most deployments will not need much load-balancing as the typical database deployment results in a small number of table keys.

 Copyright 2014 SafeNet, Inc. All rights reserved.

While the table keys are re-keyed, new keys are generated in the HSM and encrypted for storage in the database. Within an HA group, these keys are generated on the primary HSM and then, even though they exist on the HSM for only a moment, they are replicated to the entire HSM group as part of the availability logic. These events are infrequent enough that this extra replication has minimal impact.

Conclusion

The Luna high availability and load balancing features provide an excellent set of tools to scale applications and manage availability of cryptographic services without compromising the integrity of cryptographic keys. They do not need to be copied out of an HSM and stored in a file to achieve high levels of availability. Indeed, recovery from many failures is much more rapid with Luna's keys-in-hardware approach since each HSM maintains its own copy of all keys directly inside it. A broad range of deployment options are supported that allow solution architects to achieve the availability needed in a manner that optimizes the cost and performance without compromising the assurance of the solution.

Backup and Restore Configurations

While some applications might deal in ephemeral objects (keys, certs, other) that are erased after using, in many Luna HSM applications, the keys and objects within the HSM and partition have value and are meant to persist. For such valuable data, any security regime requires that the data be backed up in secure fashion, and stored securely.

For Luna SA, the backup option is the Luna [Remote] Backup HSM, which can be connected directly to the Luna SA HSM to perform backup or restore operations on the spot. The Backup HSM can also be connected to a host computer, located at a distance from the source HSM, and can perform backup and restore operations over secure network connection. This is normally the case when the source HSM is kept in a secure server room or a lights-out facility. The Backup HSM is not able to perform cryptographic operations; it functions only in its secure backup/restore role. The Backup HSM configures itself to be Password Authenticated or PED Authenticated, according to the HSM that it backs up. This is negotiated at backup time. See the Administration Guide for more detailed information and instructions.

For Luna PCI-E, the backup option is the Luna [Remote] Backup HSM, which can be connected directly to the Luna PCI-E HSM to perform backup or restore operations on the spot. The Backup HSM can also be connected to a host computer, located at a distance from the source HSM, and can perform backup and restore operations over secure network connection. This is normally the case when the source HSM is kept in a secure server room or a lights-out facility. The Backup HSM is not able to perform cryptographic operations; it functions only in its secure backup/restore role. The Backup HSM configures itself to be Password Authenticated or PED Authenticated, according to the HSM that it backs up. This is negotiated at backup time. See the Administration Guide for more detailed information and instructions.

 Copyright 2014 SafeNet, Inc. All rights reserved.

For Luna G5, the backup option is cloning of HSM or partition contents to another Luna G5 HSM, which must be of the same authentication type (Password authenticated, or PED authenticated). See the Administration Guide for more detailed information and instructions.

Host Trust Link (HTL) Configurations

The traditional model had an application server acting as a client engaging an HSM server so that together they could provide secured application and crypto services to end-users. The application server (a computer in a server room, acting as a client to the HSM, and acting as a server to your users), the HSM server (in that same server room, or another, providing secure cryptographic services and/or acceleration for your client-server transactions), and the end-user consumer of services were all individual computers in the physical possession and control of their various owners.

That model is going away, replaced by scenarios where application servers can be Virtual Machine instances, rather than individual specific computers.

Virtualization brings a number of benefits. Among those, a virtual client is:

- flexible
- portable
- not tied to a specific hardware platform.

Virtual Machines are being deployed in:

- Private Cloud – an enterprise creates its own virtual-machine environment to serve the enterprise's constituent departments or business units; the private cloud remains invisible and inaccessible to outsiders
- Hybrid Cloud – an enterprise creates its own virtual-machine environment that it makes available for internal use and also provides as a service to its external customers
- Public Cloud – an enterprise creates a virtual-machine environment that it makes available as its primary service to businesses and individuals

Both the traditional and virtual-machine environments rely on HSMs and HSM servers to secure data and transactions, and accelerate the cryptographic aspects of transactions, as well as to secure important keys and certificates.

What Threats Come with Advances in Virtual Technology?

The threat of a stolen Server has always been a security concern for Enterprises. Traditionally this form of attack was relatively difficult, as walking out of a Data Center with a server should be rather difficult. Historically, the Enterprise supplied their primary data product, such as database or application access, supported by back-room cryptographic services from SafeNet HSMs. The Enterprise provided physical security for their application/database servers and for their SafeNet HSMs and HSM Servers, while SafeNet products provided the link security via the Network Trust Link (NTL) service. This threat paradigm shifted significantly with the introduction of virtualized server instances.

 Copyright 2014 SafeNet, Inc. All rights reserved.

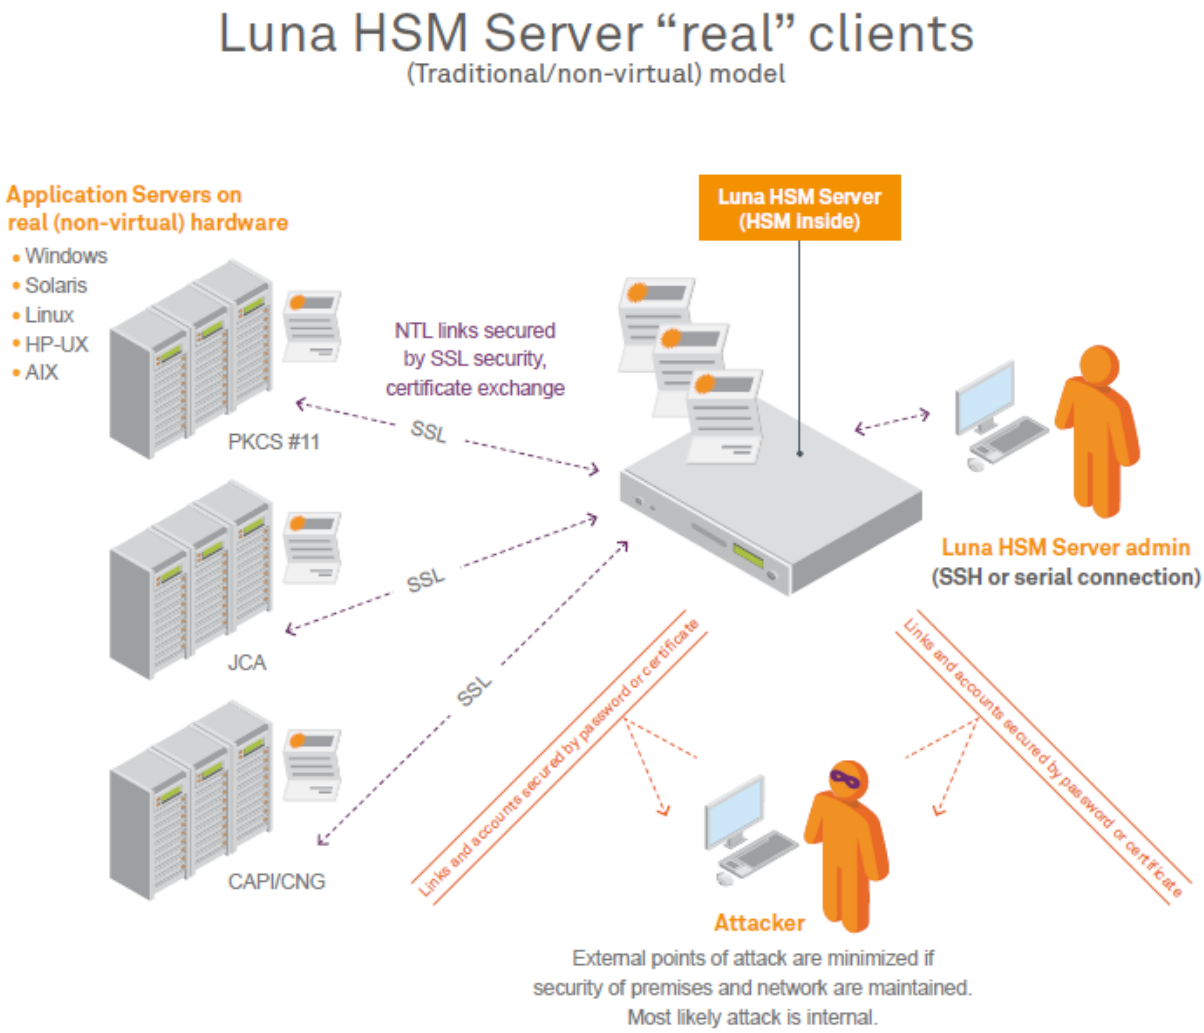

The threat has now evolved from traditional (steal the server) to virtual (steal the server Virtual Machine instance).

Virtual Clients

When the client employs the leverage of virtual hosting, that client could exist potentially anywhere in the world. The location could change at any time, unannounced. Importantly, multiple instances of a VM can be launched at any time, and control over the VM image is not fully under its nominal "owner's" control. The fact that a Luna HSM server does not notice when a VM moves allows smooth operation of the client that is using the resources of that HSM. The HSM server neither knows, nor cares, that its assigned client VM is moving (or not); the crucial concern is that that HSM server knows it is always talking to the right VM. The possibility of an unauthorized VM clone arises out of the portability and reproducibility of VMs in general. This is the virtual-world equivalent of an unauthorized person walking out of a server room with an application server.

The Unique Challenge

The challenge for an HSM serving virtual clients is to know that the HSM server is talking to the authorized instance of a virtual client, and no other. Virtual Machines in many cloud environments are 100% isolated from their physical

 Copyright 2014 SafeNet, Inc. All rights reserved.

environment, therefore no physical attribute (not a TPM, nor a CPU ID, nor a MAC address) can be used to lock down the VM. Similarly, cloud-service metadata, like any data, is easily copied and manipulated, and therefore is not suitable as a link-securing characteristic.

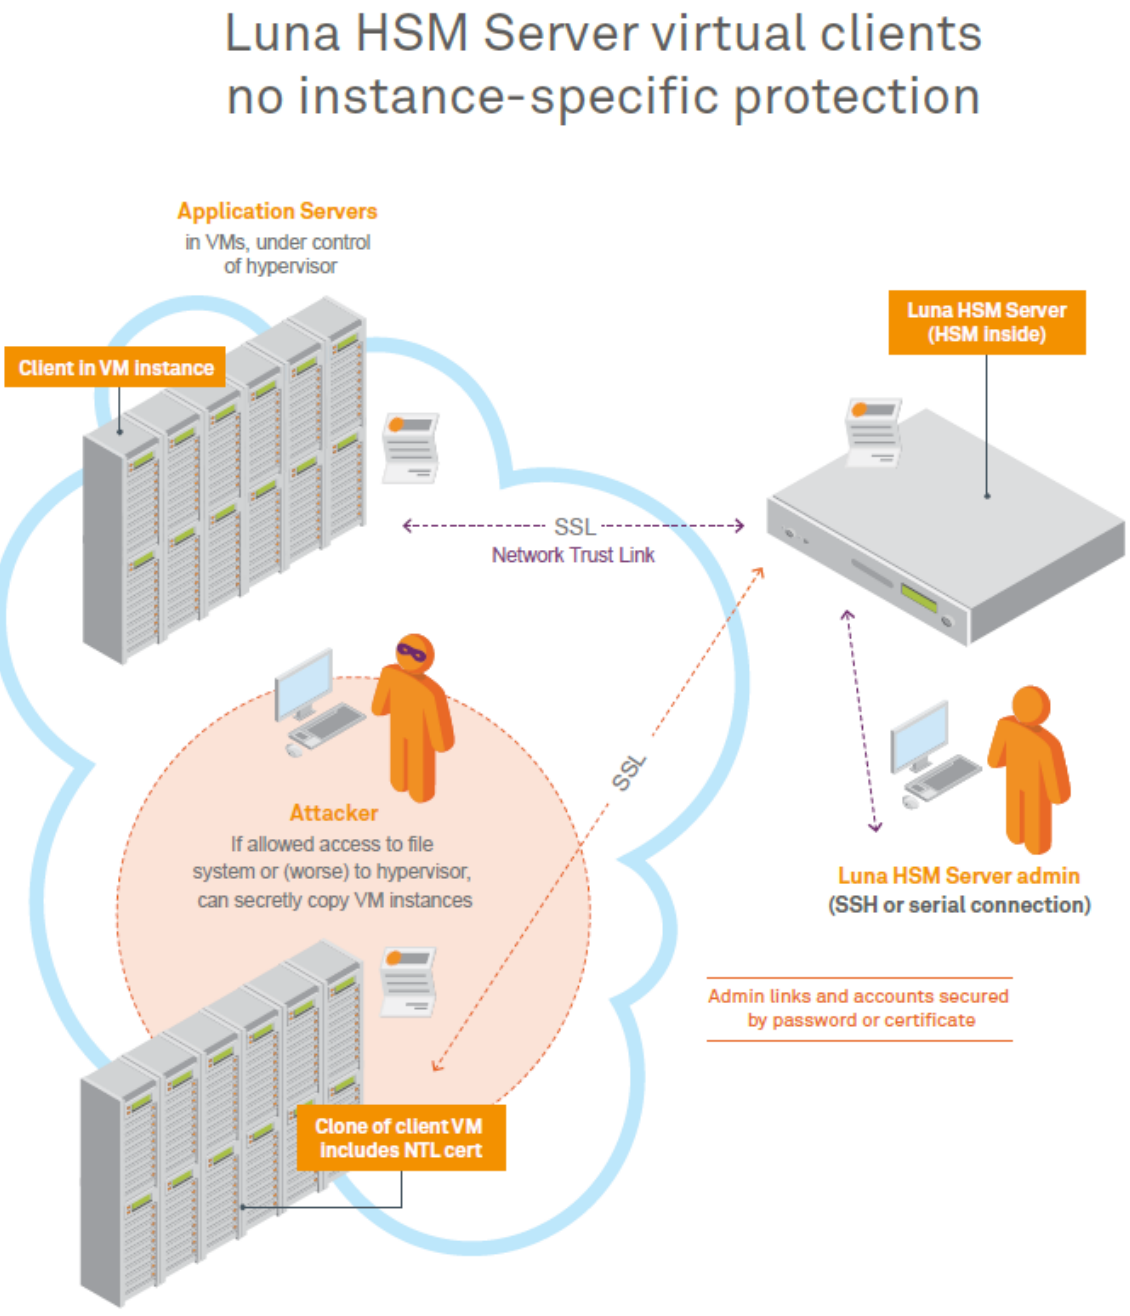

 Copyright 2014 SafeNet, Inc. All rights reserved.

What Is SafeNet Doing?

To protect against potential attacks, such as illustrated above, and to continue to offer “defense in depth”, SafeNet developed HTL, the Host Trust Link. HTL with its One-Time-Token solution is SafeNet’s built-in, HSM-based protection of HSM/Client registrations for cloud solutions.

With the VM decoupled from any specific piece of hardware or physical location, HTL uses a proprietary binding protocol to maintain the connection's association with a given VM regardless where that physical VM instantiation resides. The NTL service is still used, as before, but the new verification layer is added.

HTL supports two objectives:

- Ensure that a stored VM image, containing NTL credentials, cannot be cloned to establish an unauthorized NTLS connection to the Luna HSM server.
- Provide protection against cloning attacks after the VM binding has been established in a running VM.

The secret data used to protect the HTL link and ensure it cannot be spoofed or re-used is maintained only in RAM, which greatly increases the difficulty of an attack.

The Problem

When deploying clients in VM/Cloud environments, it is possible to pre-configure each VM with NTLS credentials (assuming that a unique IP address can be supplied for each set) and to provision both the client and an HSM partition when convenient. A virtual client can then be launched and begin interacting with a Luna HSM server over an NTLS link, without specific setup steps required for that VM. A clone of that pre-configured client VM, in the wrong hands, could work as well.

Our Solution

The HSM with HTL enabled will not allow an NTLS connection with a client instance until the Host Trust Link establishes that the client requesting NTLS is the correct VM instance of that client.

Once the VM is started and the HTL link is active, it might be possible for an internal attacker to make a complete copy of a running VM, in an attempt to impersonate the original client. The following layered protections mitigate potential concerns with respect to the provider security:

- **Binding to IP:** NTL binds the original VM to one IP address. If a clone of this VM is made with a different IP address, it will be unable to use the HSM. If a clone is made and assigned the same IP address, either the original VM would have to be killed (a noticeable event) or there would be network collisions (also detectable).
- **TLS encrypted communications:** All HTL counter values and synchronization packets are sent over a TLS link encrypted with a dynamically generated secret. This secret is in turn derived from a private key and certificate that are generated specifically for that VM instance during the HTL setup sequence. This arrangement makes it extremely unlikely that an attacker could use a cloned VM to “take over” an existing HTL connection as they would confront the hijacking protections of the TLS protocol.
- The binding protocol requires a **One Time Token (OTT)** from the Luna HSM appliance, generated specifically for that client instance. This prevents an attacker, cloning a VM at rest, from using the cloned image to connect to the Luna HSM.
- **Random data used in generating One Time Tokens** is derived from the HSM’s hardware Random Number Generator (RNG complying with NIST SP 800-90), assuring maximum randomness, and therefore highest quality input to the process.
- **One Time Token auto-refresh:** The HTL maintains a constantly changing synchronization code with the HSM

 Copyright 2014 SafeNet, Inc. All rights reserved.

server, based on a random initial counter value and step interval assigned by the HSM, which allows the authorized instance to re-establish its HTL after brief periods. The length of this period is configurable by the HSM administrator and it defaults to 2 minutes. Administrators can lengthen the time for improved reliability if the network links are unreliable, or shorten the time to increase the overall security of the HTL.

When the HTL mode is active, then ANY way an attacker manages to obtain an unauthorized copy of the VM it will be rejected by the HSM (until it receives a valid One Time Token). For such an attack to succeed, the counter would need to be re-synchronized to match the original VM by manipulating its value in RAM. This attack might be possible, but the use of a random initial counter value, a random step interval, and the ongoing synchronization, presents a significant barrier.

If a client requires VM binding, and an existing HTL link for that client goes down, the HSM server kills all existing NTLS connections from that client. This action occurs immediately, and is independent of the grace period (if any).

A client user, using a supplied GUI tool, can check the status of the HTL link for every configured, registered appliance.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

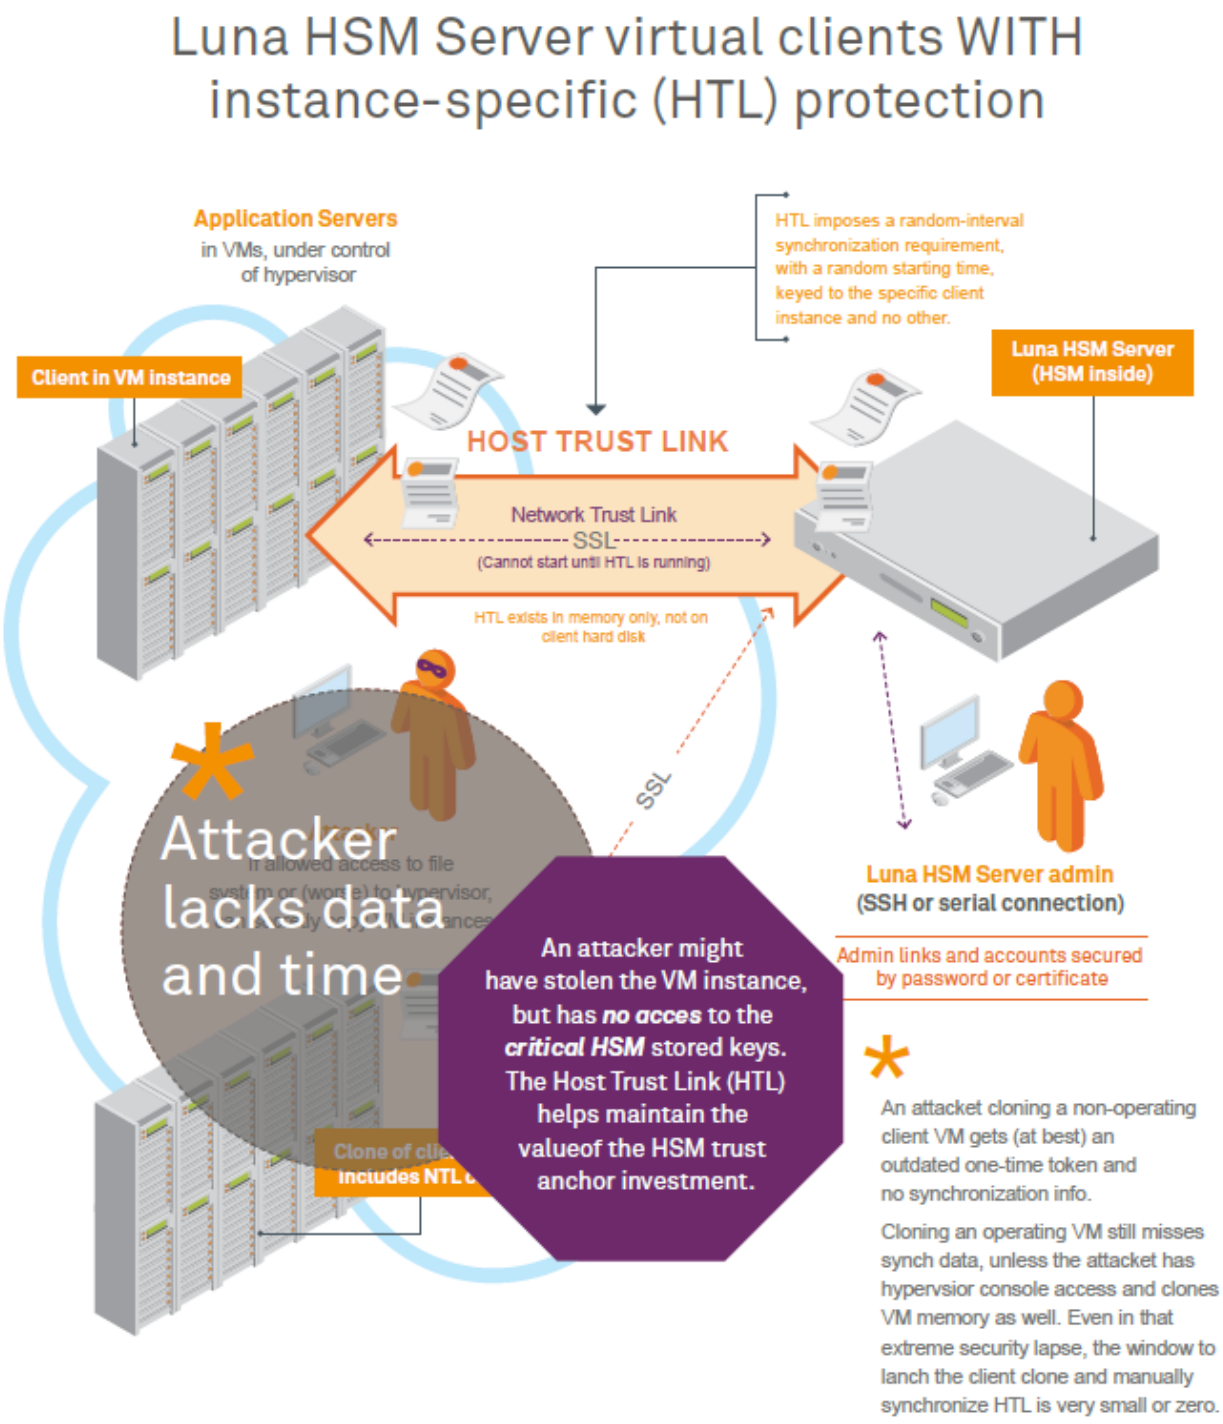

New opportunities, new threats – evolved protection

Why did we choose to create our own trust anchors to achieve the desired security when moving to a virtual environment, rather than relying on attributes available in the VM?

 Copyright 2014 SafeNet, Inc. All rights reserved.

VMs are intended to be completely isolated from their physical environments. VM attributes are already difficult to find that

- are not static
- can reliably distinguish between VM instances, and
- are not easy to spoof.

The trend is toward greater isolation. HTL is a SafeNet-generated and controlled link-authentication protocol, independent of VM attributes. SafeNet OTT technology provides enhanced security for future clouds without giving up the benefits of the cloud.

In Which Environments Does SafeNet's HTL Protect?

HTL is introduced for the virtual environment because there is a pressing need to control a VM's ability to connect. However, HTL can also help in the non-virtual world. Some customers are concerned that an attacker could grab the NTL private key file from a legitimate physical server, move it to a rogue server, and connect from there -- a physical world version of the malicious VM clone. HTL can address that concern.

 Copyright 2014 SafeNet, Inc. All rights reserved.

CHAPTER 4

Luna HSM Product Security Features

Luna HSM products include a number of features that enhance security and allow you to configure aspects of security to fit your situation.

Some are decided at purchase time (example: does your HSM require Password authentication, or PED authentication). Others are determined during setup and configuration (example: "SO can reset Partition PIN" and "Force user PIN change after set/reset", both of which are HSM policy settings).

Further, certain policy changes in the HSM or in a Partition are destructive - meaning that any current contents are lost when the policy changes. This is considered a necessary security measure because those changes represent a modification of the security level of the HSM.

Another aspect of security is the manner in which different roles are separated - a given user or administrator can perform only a limited set of operations that fit within a defined role. Other roles have other responsibilities that do not overlap. The compartmentalization limits the scope of action of any one person, thus limiting possible damage if the holder of a single role is compromised. Of course, you can give all the passwords or all the PED Keys to just one person, if you like, but that would be a matter for your organization's security policy. If your security policy is silent on the matter, then it should be updated to address your use of HSMs.

The Luna HSM security features are described in the following sections:

Roles and Users

A basic concept for cryptographic operations with HSMs is the separation of roles. For security and oversight, it is desirable to separate administrative functions from operational cryptographic functions. To that end, Luna HSM products support a variety of roles and users. The different types of HSM, and the options available to them, support a variety of operational and security regimes.

The following diagram summarizes the Cryptoki roles.

Release 5.4.1 007-011329-006Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

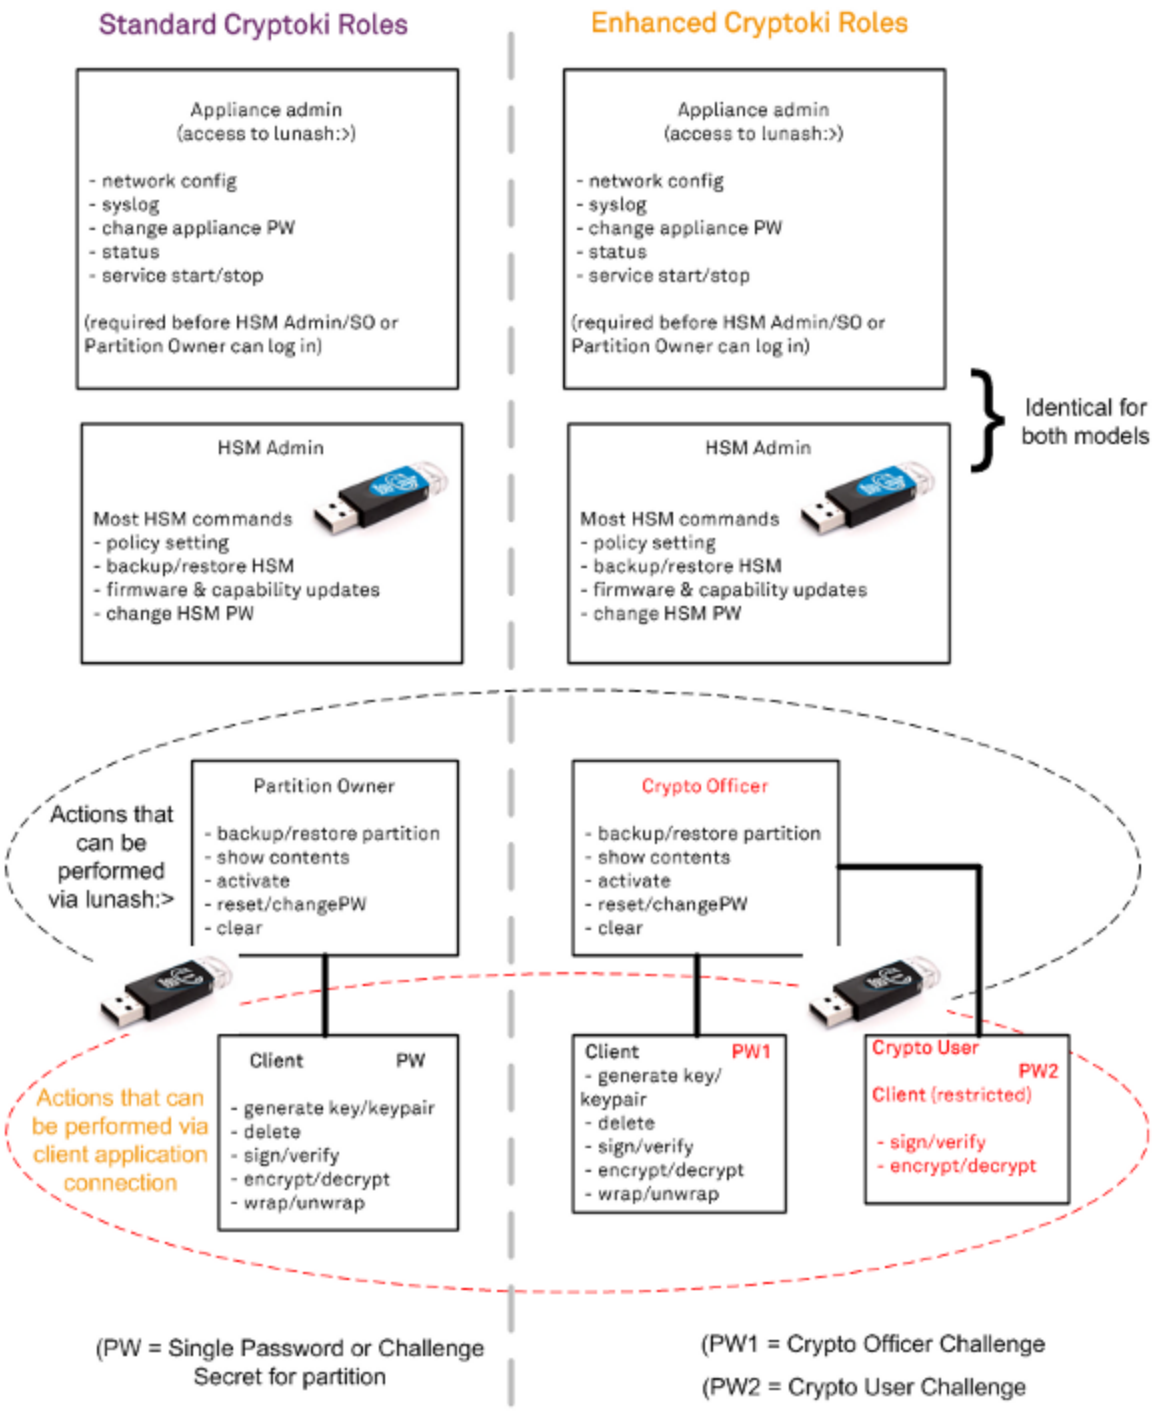

The Crypto Officer and Crypto User roles, described on the right-hand side of the diagram (above) exist only for Luna HSM with Trusted Path Authentication. They don't exist for a Luna HSM with Password Authentication.

In addition to providing the Crypto User password, a Client application must also pass the user type CKU_ RESTRICTED_USER (or the alias CKU_CRYPTO_USER).

To work with a Partition as Crypto Officer, **OR** for applications that use the existing standard, your application must pass the user type CKU_USER (along with the Crypto Officer / Partition Owner password). However, this type now

 Copyright 2014 SafeNet, Inc. All rights reserved.

has an alias CKU_CRYPTO_OFFICER, which you might prefer to use for reasons of clarity. (This concerns you only if you are an application developer.)

About Capabilities and Polices

SafeNet Luna HSMs are built on one of our general-purpose HSM platforms (hardware plus firmware), and then are loaded with what we call "personality", to make them into specific types of HSM with specific abilities and constraints, to suit different markets and applications. The built-in attributes are called "Capabilities" and describe what the HSM can do as it comes to you from the factory. Some capabilities are unalterable, except by re-manufacturing the HSM. Many HSM capabilities can be altered by means of HSM Policies, which coincide one-for-one with the capabilities that they alter.You can view the current HSM capabilities and policies with the `hsm showpolicies` command:You can change a current HSM policy with the `hsm changepolicy` command.

Similarly, capabilities and policies for each HSM partition control the behavior and the security parameters of the partition.

If a capability governs a security parameter, then the respective policy can set the HSM or the HSM partition to be more restrictive than the base capability allows, but never less restrictive.

Policy change actions that materially affect the cryptographic security of the HSM or of a partition are "destructive", meaning that if you invoke a change to such a policy, all contents of the HSM (or of the partition) are destroyed. In such an event, you can create new versions of objects that were formerly on the HSM or in the partition, or you can restore from backup.

Refer to the Configuration Guide and the Administration Guide for further discussion and instruction around capabilities and policies.

About M of N

The M of N feature provides a means by which organizations employing cryptographic modules for sensitive operations can enforce multi-person control over access to the cryptographic module. The feature is available in all Luna SAs configured to use Trusted Path authentication – using the PIN Entry Device (PED) and PED Keys.

M of N involves a splitting of the authentication secret into multiple parts or splits. The shared secret is distributed (or “split”) among several PED Keys (“split-knowledge access control”). Every type of PED-administered HSM secret can be split when it is created: blue SO PED Key, black User/Partition Owner PED Key, red Cloning Domain PED Key, orange Remote PED Vector Key, purple Secure Recover Key, white Audit PED Key.

Without M of N, you can initialize an HSM such that you must produce a single blue HSM Admin/SO PED Key in order to login and perform HSM management functions, and you must produce a single black Partition Owner/User PED Key in order to activate a Partition to receive Client connections and allow Client applications to perform operations within the Partition, and so on. And that can be the extent of your security and oversight. If that is sufficient, you can stop reading.

With M of N, the authentication secret on one blue SO PED Key or one black Partition Owner PED Key (or red Domain key or orange Remote PED key or purple Secure Recovery key) is still necessary, but is no longer sufficient for authentication. Access now requires additional authentication by an overseer, or several overseers. That additional oversight is the M of N "split knowledge shared secret". What that means is that the SO secret, or partition User/Owner secret, or cloning Domain (as well as the Remote PED secret and the Secure Recovery secret and the Audit secret) can be split into portions (over several PED Keys of the current color, rather than just one), and those must be brought together in order to re-create the complete secret. At initialization time, you get to specify into how many splits or shares each authentication secret is divided - this is quantity N (which can be any number from 1 to 16). You also

 Copyright 2014 SafeNet, Inc. All rights reserved.

specify how many of those splits or shares must be joined together by Luna PED in order to re-create the secret - this is the quantity M. M can be less than or equal to N.

Where and When to Use M of N

Use M of N when you want a particular type of HSM access to require the presence of more than one person. M of N is invoked per authentication secret. That is, it applies to only those secrets where you deliberately choose to invoke MofN as the secret is being created/imprinted. Thus you could have MofN multi-person control imposed for SO and Domain, but not for Partition owner/user, nor SRK, nor RPV... or any other combination that made sense in your environment.

During initialization of the HSM, the HSM Admin or Security Officer [SO] invokes M of N if desired as the procedure reaches the point of creating/imprinting each authentication secret. The SO specifies how many shares (also sometimes called “splits”) will make up the shared secret. This total number is N and may be any number up to 16. The SO then specifies how many of that total number of (current color) PED Keys are to be required at each login. This second number, M, can be any number up to N. From that point on, any future login or invocation of that particular authentication (blue key, black, red, orange, purple) to the HSM requires that quantity M of that-color share keys be provided. The result is that no single person can operate that aspect of the HSM. One holder of the Owner key or the HSM Admin/SO key must bring together M different share-holders, each with one of the black or blue keys, as appropriate, before the HSM can be unlocked.

M of N is not a splitting of the private signing key; it is splitting of the Luna HSM's individual authentication/access secrets. That is, M of N is a splitting of the secret that lets you into the HSM, but not a split of the working (encrypting, decrypting, signing, verifying) secrets - your keys and certificates - contained inside the HSM.

Do not use M of N unless you will be giving each split-containing PED Key to a different person. We recommend that you not use M of N unless you have established a definite need for it. The additional security of split-knowledge shared-secret multi-person access control comes at the cost of additional administrative overhead, and increased possibility of making an administrative or handling error that could leave you unable to access your keys and certificates.

Historical Note

In previous versions of Luna HSM, M of N was a selection made at the command-line (either lunash:> or lunacm:>) via the hsm init command. You could elect to use M of N or not, by means of options to the hsm init command. M of N, was a separate secret, spread across N green keys. If you invoked M of N, then it was always in force for that HSM (until the HSM was re-initialized). If you invoked M of N, it was in force HSM-wide.

Current Practice

Beginning with Luna HSM 5.0, the green keys no longer exist. Each standard authentication secret (SO, User, Domain, RPK, SRK, Auditor) can itself be split into N different components, of which M of them are needed to reconstitute that authentication secret. The decision to invoke M of N for any of the HSM's authentication secrets is no longer made via the command line. Instead, M of N is a PED function, a choice that you make when the secret is created (such as during HSM initialization or partition creation). M of N can therefore be applied to some secrets of an HSM and not to others, at your discretion, and as your organization's security policy dictates.

In usual practice, you select a number M which is the number of trusted people who must be present when HSM authentication is performed - each of them is issued a colored PED Key containing one share of that total M of N secret. The larger the number, the more operationally difficult it can be to get them all together when needed. Then you select a number N which should be a little larger than M, to allow for substitutions. This allows you to achieve M different secret shares in order to access your HSM, even though some of the total key holders might be absent due to illness, travel, etc. That N is the total number of shares into which the M of N secret will be split.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

To login with M of N in force, you are first prompted to supply a blue PED Key (or a black PED Key, as appropriate to the task), then you are prompted to supply each additional (different) key of that color until M splits have been presented - those can be any M of those keys, in any order, as long as all are different. That is, the secret is spread over N keys, but you need only M of them to recreate the complete secret when required (where M is usually less than N).

Tamper, Secure Transport, and Purple PED Keys

The HSM recognizes a number of tamper conditions (including over/under-temperature, physical interference, etc.), and allows you to choose how those are treated. The options range from simple reporting of an event in the HSM log, to temporarily (or even "permanently") disabling the HSM. In addition, the tamper function has been expanded to include Secure Transport Mode (STM) for ultimate security when shipping or storing your Luna HSMs. The advanced tamper features and ability to set STM are reserved for PED-authenticated Luna HSMs.

The use of purple PED Keys is optional unless your security policy dictates that tamper events must require a response from the HSM Admin.

The use of Secure Transport Mode (STM) is optional unless your security policy dictates that level of preparation before shipping or storage of the HSM.

If you wish to invoke Secure Transport Mode before shipping (or storing) a Luna SA HSM, you must enable the Secure Recovery Key (SRK). The SRK moves one of the two recovery splits (secure recovery vector or SRV, used to recover the Master Tamper Key in case it is destroyed by a tamper event or by STM) out of the HSM and imprints it onto a purple PED Key.

Those actions are described in detail elsewhere.

About the Purple SRK (secure recovery key)

Due to its nature, the purple PED Key (and its contained secret) behaves differently, in some respects, than all the other PED Keys.

- You choose to use this feature to enhance security during shipments or to enforce certain responses in case of physical tampering of the Luna SA (once again, it is optional - you can use all other features of the HSM without ever invoking a purple PED Key). You must put safeguards in place to ensure that the SRK does not go missing - without the purple PED Key, you cannot recover from STM or a tamper event, and must ship the HSM back to SafeNet for remanufacture.
- One of the safeguards that you can use is to make copies of the SRK at the time it is generated (*). If one of the copies is lost or destroyed, you can still recover the HSM.
- Another safeguard might be to extract the SRV onto multiple SRK splits (M of N greater than 1) rather than just one. If one of the N splits is lost or destroyed, you can still recover the HSM if you can locate quantity M of the remaining splits.
- As a safeguard against loss of the purple key in shipment, you do not need to ship the SRK to the site where the HSM is being installed. You can use Remote PED to perform the recovery from Secure Transport Mode.

Unlike all other PED Keys, the purple PED Key cannot be duplicated via Luna PED's stand-alone duplication facility in the PED's Admin menu. If you attempt to do so, the PED insists that the source key you have presented is blank, and does not continue. Therefore, if you expect to need more than one copy of the SRK, you must make those duplicates when the SRK is created - either at **hsm srk enable** or at **hsm srk keys resplit**.

 Copyright 2014 SafeNet, Inc. All rights reserved.

CHAPTER 5
General Security Guidance

This chapter provides information about handling/storing/using your Luna HSM in secure fashion, and about ensuring that your network connections to the HSM and HSM host are as secure as possible. It contains the following sections:

About Connection Security

The following is not critical if your Luna systems reside inside secure locations, behind strong firewalls, and are managed only within/between such secured locations (via VPN for example).

However, if your application places the Luna appliance or HSM host in the "DMZ", please consider the following:

- Attackers are known to be making concerted efforts to compromise server administrator account passwords. Given research published over the past few years showing the capabilities of popular game console hardware, for example, to act as extremely fast brute force password generators, it is very likely that these recent attacks are making use of automated systems. For this reason, it is strongly recommended that particular attention be given to creating strong passwords for the HSM host system's accounts. If possible, pass-phrases of 15 characters or more should be considered. One established technique for generating pass-phrases is to select a phrase at random from a book, remove spaces and punctuation, and insert numeric and special characters to replace some of the letters. If you use this sort of technique, it is good to avoid some of the more common replacements such as capitalizing the first or last character in a word, replacing “e” with “3” or “s” with “$”, etc. since they would be the first ones tried by an attacker or password generator system.
- Given the sheer numbers of computer-using people in the world, any 'rule of thumb' that you can devise for streamlining your password-making has undoubtedly been thought of by someone else. If it's a rule, it can be automated, so assume that it has been automated by password-cracking programs everywhere. For example, look to the emerging "language" of text-message abbreviations for examples of substitutions that are already widely practiced and would therefore be easily cracked.
- Longer and more complicated passwords are progressively harder to crack, but they are also far more difficult to remember. Therefore long complicated passwords are more likely to need writing down, which drastically increases risk of exposure by means of 'social engineering' or simple detective work on the part of attackers.
- Currently he most secure text password seems to be a string of several **unrelated** words in your language of choice, preferably with no double characters, and totaling more than twenty characters. So, don't choose an unmodified phrase from a book. Your password should be nonsense, but nonsense that you can remember.
- For example *battery_trick$rapid6pink* - to get around the rule of "no doubles", you could insert a dash, or space or other character *bat-tery_trick$rapid6pink* but don't use that exact example - once this Help becomes public, that combination will be in a dictionary.
- Change the SSH port number from the well-known number 22 to something in the range of 1025 to 65535.
- Use the lunash:>sysconf ssh port command to change the SSH port number.

Release 5.4.1 007-011329-006Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Consider Using Certificate-based Authentication

You can choose to use certificate-based authentication for your "admin", "operator", and "monitor" users (or named users with those roles) to connect, instead of password authentication.

See the sysconf ssh publickey commands. When creating your certificate on a client/admin computer, select a key size of 1024 bits or greater to generate the certificate.

However, because the certificate resides on a computer, it is ultimately only as secure as access to that computer, which is likely protected by password (see above).

DRAFT SP 800-118 Guide to Enterprise Password Management

NIST announced that Draft Special Publication (SP) 800-118, Guide to Enterprise Password Management, has been released for public comment. SP 800-118 is intended to help organizations understand and mitigate common threats against their character-based passwords. The guide focuses on topics such as defining password policy requirements and selecting centralized and local password management solutions.

http://csrc.nist.gov/publications/PubsDrafts.html#800-118

Security and Handling Considerations - HSM Appliance

This section discusses general security and handling issues related to the Luna SA HSM appliance.

Physical Security of the Appliance

The HSM appliance is a commercial-grade secure appliance. This means that:

- It is provided with anti-tamper external features that make physical intrusion into the unit difficult - tamper-resistant screws must be drilled out, in order to open the case, and tamper-evident stickers secure the seams. These measures do not deter a determined attacker, they merely deter casual intrusion and leave visible evidence of attempts (successful or otherwise) to open the unit.
- Vents and other paths into the unit are baffled to prevent probing from the outside.
- The HSM Keycard, inside the appliance, that houses the actual HSM components, is encased in an aluminum shell, filled with hardened epoxy. Attempts to gain access to the circuit board itself would result in physical evidence of the attempted access and likely physical destruction of the circuitry and components, thus ensuring that your keys and sensitive objects are safe from an attacker.

If an attacker with unlimited resources were to simply steal the appliance, and apply the resources of a well-equipped engineering lab, it might be possible to breach the physical security. However, without the Password (password authenticated HSMs) or the PED Keys (PED-authenticated HSMs), such an attacker would be unable to decipher any signal or data that they managed to extract.

With that said, it is your responsibility to ensure the physical security of the unit to prevent such theft, and it is your responsibility to enforce procedural security to prevent an attacker ever having possession of (or unsupervised access to) both the HSM and its authentication secrets.

Physical Environment Issues

The data sheets provided by SafeNet show the environmental limits that the device is designed to withstand. It is your responsibility to ensure that the unit is protected throughout its working lifetime from extremes of temperature, humidity, dust, vibration/shock that exceed the stated limits.

 Copyright 2014 SafeNet, Inc. All rights reserved.

It is also your responsibility to ensure that the HSM appliance is installed in a secure location, safe from vandalism, theft, and other attacks. In summary, this usually means a clean, temperature-, humidity-, and access-controlled facility. We also strongly recommend power conditioning and surge suppression to prevent electrical damage, much as you would do for any important electronic equipment.

Communication

Communications with the unit are either local and, therefore, subject to direct oversight and control (you decide who is allowed to connect to the serial port or the PED port) or via secure remote links. All remote communications are as secure as SSH and TLS with tunneling protocol can make them.

Authentication Data Security

It is your responsibility to protect passwords and/or PED Keys from disclosure or theft and to ensure that personnel who might need to input passwords do not allow themselves to be watched while doing so, and that they do not use a computer or terminal with keystroke logging software installed.

HSM Audit Data Monitoring

The HSM Keycard of the Luna HSM appliance stores a record of past operations that is suitable for security audit review. The easiest way in which to retrieve this record is to use the "`hsm supportinfo`" command and extract the dual port data provided within the `supportinfo.txt` file that is returned by the command. Because of the limited storage capacity of the HSM card, it has a limited size window in which to write these records and it must periodically re-start from the beginning of the window and overwrite existing records. For this reason, it is important that the audit data be retrieved often enough to ensure no data loss. Under typical load conditions, retrieving the file once every eight hours should be sufficient. However, for very heavy loads or operations containing large input data payloads, it might be necessary to retrieve the file as often as once per hour.

Audit Logging

Beginning with Luna HSM 5.2, the secure Audit Logging feature provides an Audit role (white PED Key) separate from all other HSM roles, to manage a secure audit logging function. Audit logging sends HSM log event records to a secure database on the local file system, with cryptographic safeguards ensuring verifiability, continuity, and reliability of HSM event log files.

Intended Installation Environment

The following assumptions are made about the environment in which the Luna cryptographic modules will be located and installed:

- Those responsible for the Luna HSM must ensure that the authentication data for each Luna HSM account is held securely and not disclosed to persons not authorized to use that account.
- Those responsible for the Luna HSM must ensure that it is installed, managed, and operated in a manner that is consistent with the local security policy.
- The host IT environment must be configured and checked to ensure that any applications installed in the host environment, that require access to the HSM are legitimate, are valid and have been vetted for authenticity and integrity (i.e., have not been modified for malicious purposes).
- Those responsible for the Luna HSM must ensure that it is installed and operated in an environment that is protected from unauthorized physical access.

 Copyright 2014 SafeNet, Inc. All rights reserved.

- Those responsible for the Luna HSM must ensure that there are procedures in place such that, after a system failure or other discontinuity, recovery of the Luna HSM and the host IT environment is possible without compromise of IT security.
- Those responsible for the Luna HSM must ensure that those using the Luna HSM (including Security Officers and Token/Partition Users), have a level of competence sufficient to ensure its correct management and operation. This competence may be established through a combination of training and the accompanying Installation Guide and Configuration, Administration, and Reference documentation.
- Procedural and physical measures must prevent the disclosure of cryptography-related IT assets to unauthorized individuals or users via the electromagnetic emanations of the Luna HSM .
- Those responsible for the host IT environment must ensure that no connections are provided to outside systems or users that would undermine IT security.
- Those responsible for the host IT environment must ensure that the power supplied to the Luna HSM is adequately protected against unexpected interruptions and the effects of surges and voltage fluctuations outside the normal operating range of the device.
- Those responsible for the host IT environment must ensure that the Luna HSM is operated in an environment in which there is provided adequate protection against disasters such as fire and flood.
- Those responsible for the host IT environment must ensure that the Luna HSM is located in an environment that is adequate to protect security-relevant and cryptographic key data and the Luna HSM firmware from interference or inadvertent modification by strong electromagnetic radiation from other sources.

Security and Handling Issues - Luna HSM

This section chapter discusses general security and handling issues related to the Luna HSM and its host computer.

Physical Security of the Cryptographic Module

The Luna cryptographic module is a multi-chip standalone module as defined by FIPS PUB 140–2 section 4.5. The module is enclosed in a strong enclosure that provides tamper-evidence. Any tampering that might compromise the module's security is detectable by visual inspection of the physical integrity of the module. In addition, any attempts to physically tamper with the token would likely result in the destruction of its circuitry and components, thus ensuring that your keys and sensitive objects are safe from an attacker.

The module's physical design also resists visual inspection of the device design, physical probing of the device and attempts to access sensitive data on individual components of the device.

If an attacker with unlimited resources were to simply steal a Luna HSM, and apply the resources of a well-equipped engineering lab, it might be possible to breach its physical security. However, without the Password (password authenticated module) or the PED Keys (PED-authenticated module), such an attacker would be unable to decipher any signal or data that they managed to extract.

It is your responsibility to ensure the physical security of the module to prevent such theft, and it is your responsibility to enforce procedural security to prevent an attacker ever having possession of (or unsupervised access to) both the cryptographic module and its authentication secrets.

It is your responsibility to ensure the physical security (access) of passwords or PED Keys, and to ensure that personnel who might need to input passwords do not allow themselves to be watched while doing so, and that they do not use a computer or terminal with keystroke logging software installed.

Release 5.4.1 007-011329-006 Rev C July 2014 Copyright 2014 SafeNet, Inc. All rights reserved.

Intended Installation Environment

The following assumptions are made about the environment in which the Luna cryptographic modules will be located and installed:

- Those responsible for the Luna HSM must ensure that the authentication data for each Luna HSM account is held securely and not disclosed to persons not authorized to use that account.
- Those responsible for the Luna HSM must ensure that it is installed, managed, and operated in a manner that is consistent with the local security policy.
- The host IT environment must be configured and checked to ensure that any applications installed in the host environment, that require access to the HSM are legitimate, are valid and have been vetted for authenticity and integrity (i.e., have not been modified for malicious purposes).
- Those responsible for the Luna HSM must ensure that it is installed and operated in an environment that is protected from unauthorized physical access.
- Those responsible for the Luna HSM must ensure that there are procedures in place such that, after a system failure or other discontinuity, recovery of the Luna HSM and the host IT environment is possible without compromise of IT security.
- Those responsible for the Luna HSM must ensure that those using the Luna HSM (including Security Officers and Token/Partition Users), have a level of competence sufficient to ensure its correct management and operation. This competence may be established through a combination of training and the accompanying Installation Guide and Configuration, Administration, and Reference documentation.
- Procedural and physical measures must prevent the disclosure of cryptography-related IT assets to unauthorized individuals or users via the electromagnetic emanations of the Luna HSM .
- Those responsible for the host IT environment must ensure that no connections are provided to outside systems or users that would undermine IT security.
- Those responsible for the host IT environment must ensure that the power supplied to the Luna HSM is adequately protected against unexpected interruptions and the effects of surges and voltage fluctuations outside the normal operating range of the device.
- Those responsible for the host IT environment must ensure that the Luna HSM is operated in an environment in which there is provided adequate protection against disasters such as fire and flood.
- Those responsible for the host IT environment must ensure that the Luna HSM is located in an environment that is adequate to protect security-relevant and cryptographic key data and the Luna HSM firmware from interference or inadvertent modification by strong electromagnetic radiation from other sources.

 Copyright 2014 SafeNet, Inc. All rights reserved.