nShield Security World

nShield 5c v13.4.5 User Guide

17 October 2024

Table of Contents

1. Introduction

1.1. Read this guide if …

Read this guide if you need to configure or manage:

- An Entrust Hardware Security Module (HSM).
- An associated *Security World*. nShield hardware security modules use the Security World paradigm to provide a secure environment for all your HSM and key management operations.

The nShield HSM is connected to a network by an Ethernet connection. Each HSM is configured to communicate with one or more client computers on the network. You can also configure clients to make use of any other network-connected HSMs on the network, as well as locally connected HSMs.

All nShield HSMs support standard cryptography frameworks and integrate with many standards based products.

This guide assumes that:

- You are familiar with the basic concepts of cryptography and Public Key Infrastructure (PKI)
- You have read the *Installation Guide*.
- You have installed your nShield HSM.

> 🛈 Throughout this guide, the term *Installation Guide* refers to the particular Installation Guide for your product.

1.2. Terminology

The nShield 5c is also referred to as the *hardware security module* or the *nShield HSM*.

This guide refers to other nShield HSMs by type:

nShield HSMs	nShield HSM type
Connect, Connect +, Connect XC, 5c	Network-attached HSMs
Solo, Solo +, Solo XC, 5s	PCIe HSMs
Edge	USB-attached HSMs

1.3. Model numbers

Model numbering conventions are used to distinguish different nShield hardware security devices.

Model number	Used for
NH2096-B	nShield 5c Base
NH2096-M	nShield 5c Medium
NH2096-H	nShield 5c High

1.4. Security World Software

The Security World for nShield is a collection of programs and utilities, including the hard-server, supplied by Entrust to install and maintain your nShield security system.

Entrust provides the firmware that runs on the nShield HSM, and software to run on each client computer. The nShield HSM is supplied with the latest version of the HSM firmware installed. For more information about:

- Upgrading the firmware, see Upgrading the image file and associated firmware.
- Installing and configuring the software on each client computer, see the Installation Guide and Client Software and module configuration.
- The supplied utilities, see Supplied utilities.
- Maintenance of your nShield hardware, see Maintenance of nShield Hardware.

1.4.1. Software architecture

The software, firmware, and utilities have version numbers and there is also a version number for the World which refers to the World data that is stored in encrypted form on the client computer, typically in the `opt/nfast/kmdata` (**Linux**) or `C:\ProgramData\nCipher\Key Management Data` (**Windows**) directory or on the RFS. This data includes information concerning the World itself and also concerning each key that was created within that World. The World version created is determined by the version numbers of the software and firmware used when it was first created, see Creating and managing a Security World.

The latest World version is version 3. You can query the version of the World loaded on your system by using the command `kmfile-dump`.

1.4.1.1. Hardserver

The *hardserver* software controls communication between the internal security module and applications on the network.

Separate instances of the hardserver run on the unit and each client that is configured to work with the unit. There is a secure channel, known as the *impath*, between the two software instances, which forms a single secure entity for transferring data between the unit and the clients. See also Compatibility issues.

The unit's hardserver is configured using the front panel on the unit, or by means of uploaded configuration data. Configuration data is stored on the unit and in files in a specially configured file system on each client computer. For more information about using:

- The front panel to configure the unit, see The front panel interface
- The specially configured file system to configure the unit and the client, see Client Software and module configuration.

1.4.1.2. Remote file system (RFS)

Each unit uses a *remote file system* (RFS). You can configure the RFS on any computer, but it is normally located on the first client that is configured. The RFS contains:

- The master configuration information for the unit
- The Security World files
- The key data.

Do not copy the master configuration to file systems on other clients. You can copy Security World files and key data to other clients to allow you to manage the unit from more than one client. To make it available to the unit, copy to the RFS the data for Security Worlds, cards or keys that you create on a client that does not contain the RFS.

1.4.1.3. Security World Software

The Security World Software is a collection of programs and utilities that you install on the client and use to maintain the nShield security system. The Security World Software includes the following:

- The appropriate installer for the client platform
- The client hardserver
- A set of utilities for configuring the nShield HSM
- A set of utilities and the KeySafe application for performing key management tasks on nShield HSMs.

The nShield HSM is referenced and used by a utility or application in the same way as a local (PCI-connected) nShield HSM.

1.4.2. Default directories

The default locations for Security World Software and program data directories on English-language systems are summarized in the following table:

Directory name	Default path (Linux)	Environment variable (Windows)	Default path (Windows)
nShield Installation	`/opt/nfast/`	`NFAST_HOME`	`C:\Program Files\nCipher\nfast`
Key Management Data	`/opt/nfast/kmdata/`	`NFAST_KMDATA`	`C:\ProgramData\nCipher\Key Management Data`
Dynamic Feature Certificates	`/opt/nfast/femcerts/`	`NFAST_CERTDIR`	`C:\ProgramData\nCipher\Feature Certificates`
Static Feature Certificates	`/opt/nfast/kmdata/hsm-ESN/features`		`%NFAST_KMDATA%\hsm-ESN\features`
Log Files	`/opt/nfast/log`	`NFAST_LOGDIR`	`C:\ProgramData\nCipher\Log Files`
User Log Files	`/home/<user>/nshieldlogs`	`NFAST_USER_LOGDIR`	`C:\Users\<user>\nshieldlogs`
Remote Static Feature Certificates	`opt/nfast/kmdata/hsm-ESN/features`		`%NFAST_KMDATA%\hsm-ESN\features`
Remote Dynamic Feature Certificates	`opt/nfast/kmdata/hsm-ESN/features`		`%NFAST_KMDATA%\hsm-ESN\features`

> 🛈 By default, the Windows `C:\ProgramData\` directory is a hidden directory. To see this directory and its contents, you must enable the display of hidden files and directories in the View settings of the Folder Options.

> 🛈 Dynamic feature certificates must be stored in the directory stated above. The directory shown for static feature certificates is an example location. You can store those certificates in any directory and provide the appropriate path when using the Feature Enable Tool. However, you must not store static feature certificates in the dynamic features certificates directory.

On Windows, the absolute paths to the Security World Software installation directory and

program data directories are stored in the indicated nShield environment variables at the time of installation. If you are unsure of the location of any of these directories, check the path set in the environment variable.

The instructions in this guide refer to the locations of the software installation and program data directories as follows:

- By name (for example, Key Management Data).
- **Linux**: By absolute path (for example, `/opt/nfast/kmdata`).
- **Windows**: By nShield environment variable names enclosed with percent signs (for example, `%NFAST_KMDATA%`).

If the software has been installed into a non-default location:

- **Linux**: Create a symbolic link from `/opt/nfast/` to the directory where the software is actually installed.
- **Windows**: Ensure that the associated nShield environment variables are re-set with the correct paths for your installation. For more information about creating symbolic links, see your operating system's documentation.

1.4.3. Utility help options

Unless noted, all the executable utilities provided in the `bin` subdirectory of your nShield installation have the following standard help options:

- `-h|--help` displays help for the utility
- `-v|--version` displays the version number of the utility
- `-u|--usage` displays a brief usage summary for the utility.

1.5. Setting the PATH for nShield utilities

It is recommended that the PATH environment variable be changed to include `opt/nfast/bin` (**Linux**) or `<%NFAST_HOME%\bin>`, which is usually `C:\Program Files\nCipher\nfast\bin` (**Windows**).

This is the directory in the nShield installation that contains the nShield command-line utilities and some DLLs.

This will allow all the nShield command-line utilities to be run without the need to type the full path, for example running `enquiry` instead of `opt/nfast/bin/enquiry>` (**Linux**) or `<%NFAST_HOME%\bin\enquiry>` (**Windows**).

`opt/nfast/bin` (**Linux**) or `<%NFAST_HOME%\bin>` (**Windows**) must be set in the PATH in order to use the OpenSSL module in the Python that is bundled with nShield.

The Python bundled with nShield is located in `opt/nfast/python3/bin` (**Linux**) or `%NFAST_HOME\python3\bin`, which is usually `C:\Program Files\nCipher\nfast\python3\bin` (**Windows**). If using the nShield Python, you may additionally want to add this directory to the PATH environment variable so that you can run the nShield python as just the python command. You may not want to do this if you are also using other Python installations on the same machine.

1.6. Further information

This guide forms one part of the information and support provided by Entrust.

If you have installed the Java Developer component, the Java Generic Stub classes, nCipherKM JCA/JCE provider classes, and Java Key Management classes are supplied with HTML documentation in standard `Javadoc` format, which is installed in the appropriate `nfast/java` directory when you install these classes.

1.7. Security advisories

If Entrust becomes aware of a security issue affecting nShield HSMs, Entrust will publish a security advisory to customers. The security advisory will describe the issue and provide recommended actions. In some circumstances the advisory may recommend you upgrade the nShield firmware and or image file. In this situation you will need to re-present a quorum of administrator smart cards to the HSM to reload a Security World. As such, deployment and maintenance of your HSMs should consider the procedures and actions required to upgrade devices in the field.

> 🛈 The Remote Administration feature supports remote firmware upgrade of nShield HSMs, and remote ACS card presentation.

We recommend that you monitor the Announcements & Security Notices section on Entrust nShield, https://nshieldsupport.entrust.com, where any announcement of nShield Security Advisories will be made.

1.8. Recycling and disposal information

For recycling and disposal guidance, see the nShield product's Warnings and Cautions documentation.

2. Security Worlds

This chapter describes the *Security World* infrastructure we have developed for the secure life-cycle management of cryptographic keys. The Security World infrastructure gives you control over the procedures and protocols you need to create, manage, distribute and, in the event of disaster, recover keys.

A Security World provides you with the following features:

- Security
- Application independence
- Platform independence
- Flexibility
- Scalability
- Robustness
- Audit logging

A Security World comprises:

- One or more Entrust nShield HSMs (such as the nShield HSM)
- An *Administrator Card Set* (ACS)
  A set of Administrator smart cards used to control access to the Security World config uration, as well as in recovery and replacement operations.

  > ⊗ Store the ACS in a secure location, separate from the HSMs, when you are not using them to provide authorization for administrative tasks.

- Optionally, one or more *Operator Card Sets* (OCSs)
  A set or sets of Operator smart cards used to control access to application keys.
- Some cryptographic key and certificate data that is encrypted using the Security World key and stored on a host computer or computers

You can add or remove cards, keys, and even hardware security modules at any time. These components are linked by the Security World key, which is unique to each world. To see how these components are related to one another, see the image below.

Distributing the keys used for different tasks within the Security World over different storage media means that the Security World can recover from the loss of any one component. It also increases the difficulties faced by an attacker, who needs to obtain all the components before gaining any information.
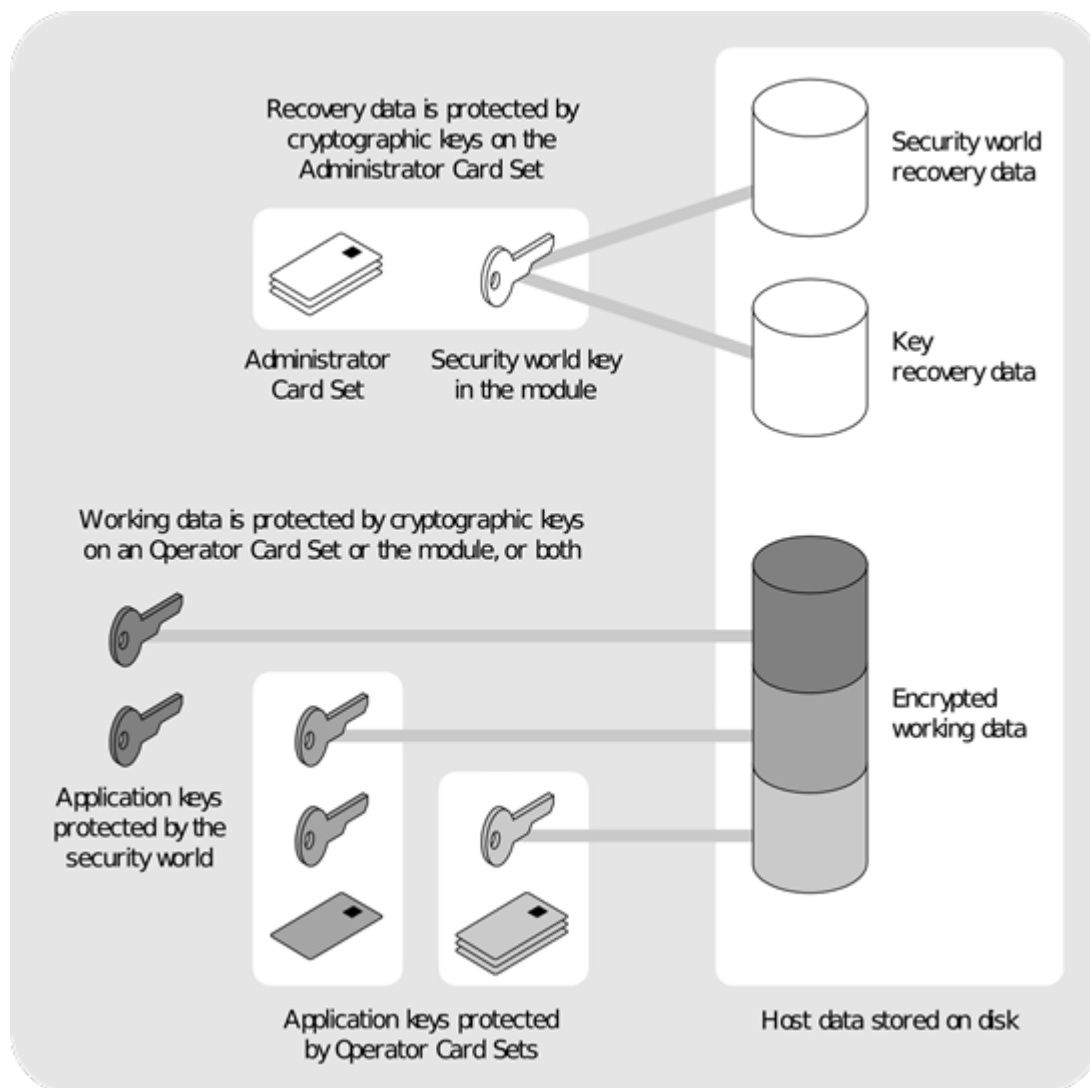
2.1. Security

We have designed the Security World technology to ensure that keys remain secure throughout their life cycle. Every key in the Security World is always protected by another key, even during recovery and replacement operations.

Because the Security World is built around nShield key-management modules, keys are only ever available in plain text on secure hardware.

> ⊗ All Security Worlds rely on you using the security features of your operating system to control the users who can access the Security World and, for example, write data to the host.

2.1.1. Smart cards

The Security World uses:

- An *Administrator Card Set* (ACS) to control access to recovery and replacement functionality

> ⛔ You must keep the ACS secure and separate from the HSMs when you are not using it, to minimize risk to the Security World.

- Zero or more *Operator Card Sets* (OCSs) to control access to application keys

> ℹ️ In FIPS 140 Level 3 Security Worlds, you require a card from either the ACS or an OCS to authorize most operations, including the creation of keys and OCSs.

Each card set consists of a number of smart cards, *N*, of which a smaller number, *K*, is required to authorize an action. The required number *K* is known as the *quorum*.

> ℹ️ The value for *K* should be less than *N*. We do not recommend creating card sets in which *K* is equal to *N* because an error on one card would render the whole card set unusable. If your ACS became unusable through such an error, you would have to replace the Security World and generate new keys. In Common Criteria CMTS Security Worlds the minimum value of *K* for the ACS is 2.

An ACS is used to authorize several different actions, each of which can require a different value for *K*. All the card sets are distinct: a smart card can only belong to the ACS or to one OCS.

Each user can access the keys protected by the Security World and the keys protected by their OCS. They cannot access keys that are protected by another OCS.

Operator Cards employ the Security World key to perform a challenge-response protocol with the hardware security module. This means that Operator Cards are only useable by an HSM that belongs to the same Security World.

2.1.2. Remote Operator

The Remote Operator feature is used to load a key protected by an OCS onto a machine to which you do not have physical access (for example, because it is in a secure area).

> ℹ️ The Remote Operator feature is not available in Common Criteria CMTS Security Worlds.

The Remote Operator feature enables the secure transmission of the contents of a smart card inserted into the slot of one module (the *attended module*) to another module (the

unattended module). To transmit to a remote module, you must ensure that:

- The smart card is from a persistent OCS
  See Using persistent Operator Card Sets for more about persistent cards.
- The attended and unattended modules are in the same Security World

To achieve secure communication channels between the attended and unattended modules, the hardserver uses an *impath* (an abbreviation of *intermodule path*), a secure protocol for communication over IP networks. The communication channels between the modules:

- Are secure against both eavesdroppers and active adversaries
- Can carry arbitrary user data as well as module-protected secrets, such as share data, that pass directly between modules.

See also Compatibility issues.

2.1.3. Remote Administration

Remote Administration is a collection of features that allow you to configure and operate an HSM or set of HSMs without being physically present at the HSM. This includes creating ACS when creating a Security World and presenting ACS to authorize loading of a Security World. It also includes creating OCS to protect application keys and presenting OCS to authorize the loading of application keys. The OCS may be persistent or non-persistent.

The ACS and/or OCS cards must be nShield Remote Administration smart cards. When presenting a card, a secure channel is formed directly between the Remote Administration smart card and the target HSM before any token shares are read from or written to the smart card. The secure channel is secure against both eavesdroppers and active adversaries.

For more information, refer to the nShield *Remote Administration User Guide*.

2.1.4. Security World and an nShield HSM

To enable the secure transmission of data between an nShield HSM and each client, the hardserver uses an impath (see Remote Operator). Any nShield HSM on the network can load a Security World securely over the network, and access its keys, wherever the HSM is installed.

Security World and key data is stored on the file system of the nShield HSM, where it is updated whenever card or key operations are performed on the HSM. The data is also auto-

matically transferred to the *remote file system* (RFS). If required, you can also share the data with client computers that use the Security World. For more information, see:

- Remote file system (RFS)
- Configuring the remote file system (RFS).

You can update the Security World on the host using:

- The nShield HSM front panel controls
- The command-line utilities
- The Cryptographic Service Provider wizard
- KeySafe.

You can also use these tools to create keys or cards. If you perform such tasks on a client other than the computer on which the RFS is installed, you must transfer the updated files to the RFS before they are available to the HSM.

2.1.5. NIST SP800-131A

When a new Security World is created it will be SP800-131A compliant.

2.1.5.1. Compatibility issues

In order to comply with the latest encryption standards, Entrust has adopted an enhanced NIST SP800-131A compliant encryption protocol between nShield HSM HSMs and their clients with Security World software installed. In some cases, this change may have an impact on the compatibility of network-attached HSMs in environments with mixed HSM deployments.

In most cases where versions of Security World software of v11.50 or later are deployed in conjunction with v11.40 software or lower, no action is required. However, there are two cases in which communication cannot be established between the HSM and clients or hosts:

- v11.50 or higher clients communicating with a v11.40 or lower nShield HSM, where the HSM client uses an nToken.
- v11.50 or higher nShield HSM communicating with a Remote File System (RFS) using v11.40 or lower.

Release version	Image versions nShield HSM	Security World Software[1] v11.40	Security World Software[1] v11.50 and later
Up to 11.40	Up to image version 0.3.5.	Supported.	For deployments with nTokens, please upgrade the nShield HSM netimage. As a less preferred option, you can downgrade the client-side software.
11.50 or later	Image 0.3.6 or later.	RFS and client software upgrade required.	Supported.

[1] Previously known as nCipher software, or nCSS.

2.1.6. FIPS 140 compliance

All Security Worlds are compliant with the Federal Information Processing Standards (FIPS) 140 specification. The default setting for Security Worlds complies with Level 2 of FIPS 140.

A Security World that complies with the roles and services section of FIPS 140 Level 2 does not require any authorization to create an OCS or an application key.

2.1.6.1. FIPS 140 Level 3 compliance

When you create a Security World, you can choose whether the Security World is compliant with the roles and services section of either:

- FIPS 140 at Level 2
- FIPS 140 at Level 3

The FIPS 140 Level 3 option is included for those customers who have a regulatory requirement for compliance with FIPS 140 at Level 3.

If you choose to create a Security World that complies with FIPS 140 Level 3, the nShield HSM initializes in that mode, conforming with the roles and services, key management, and self-test sections of the FIPS validation certificate.

Before you can create or erase an OCS in a Security World that complies with FIPS 140 Level 3, you must authorize the action with a card from the ACS or an OCS from that Security World.

For more details about FIPS 140, see http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf.

2.2. Platform independence

The Security World is completely platform independent. All key information is stored in a proprietary format that any computer supported by Security World Software can read, regardless of the native format used by that computer. This enables you to:

- Safely move a Security World between platforms with differing native formats. For example, you can move a Security World between Windows and Linux operating environments.
- Include hosts running different operating systems in the same Security World.

> When copying host data between computers using different operating systems or disk formats, use a mechanism that preserves the original data format and line endings (such as `.tar` file archives).

2.3. Application independence

A Security World can protect keys for any applications correctly integrated with the Security World Software. Each key belongs to a specific application and is only ever used by that application. Keys are stored along with any additional data that is required by the application.

You do not need to specify:

- Which applications you intend to use. You can add a key for any supported application at any time.
- How the key is used by an application. A Security World controls the protection for the key; the application determines how it is used.

Although keys belong to a specific application, OCSs do not. You can protect keys for different applications using the same OCS.

In the image above:

- **Card Set 1** protects multiple keys for use with **Application 1** and **Application 2**
- **Card Set 2** protects a single key for use with **Application 2**
- **Card Set 3** protects multiple keys for use with **Application 2** and **Application 3**
- The *Security World key* protects a single key for use with **Application 3**.

2.4. Flexibility

Within a Security World, you can choose the level of protection for each application key that you create.

When you create a Security World, a cryptographic key is generated that protects the application keys and the OCSs in the Security World.

2.4.1. Using the Security World key: module-protected keys

You can use the Security World key to protect an application key that you must make available to all your users at all times. This key is called a *module-protected key*. Module-protected keys:

- Have no passphrase
- Are usable by any instance of the application for which they were created, provided that this application is running on a server fitted with a hardware security module belonging to the correct Security World.

This level of protection is suitable for high-availability Web servers that you want to recover

immediately if the computer resets.

2.4.2. Using Operator Card Sets: OCS-protected keys

An OCS belongs to a specific Security World. Only a hardware security module within the Security World to which the OCS belongs can read or erase the OCS. There is no limit to the number of OCSs that you can create within a Security World.

An OCS stores a number of symmetric keys that are used to protect the application keys. These keys are of the same type as the Security World key.

Each card in an OCS stores only a fragment of the OCS keys. You can only re-create these keys if you have access to enough of their fragments. Because cards sometimes fail or are lost, the number of fragments required to re-create the key (K) are usually less than the total number of fragments (N).

To make your OCS more secure, we recommend that you make the value of K relatively large and the value of N less than twice that of K (for example, the values for K/N being 3/5 or 5/9). This practice ensures that if you have a set of K cards that you can use to recreate the key, then you can be certain that there is no other such card set in existence.

> **ℹ** | Some applications restrict K to 1.

2.4.2.1. Using Operator Card Sets to share keys securely

You can use OCSs to enable the same keys for use in a number of different HSMs at the same time.

If you have a non-persistent OCS, you must leave one of the cards in an appropriate card slot of each HSM. This should only be done if it is in accordance with the security policies of your organization.

To use OCS-protected keys across multiple HSMs, set:

- K to 1
- N at least equal to the number of the HSMs you want to use.

You can then insert single cards from the OCS into the appropriate card slot of each HSM to authorize the use of that key.

To issue the same OCS-protected key to a set of users, set:

- K to 1

- *N* equal to the number of users.

You can then give each user a single card from the OCS, enabling those users to authorize the use of that key.

> **i** If you have created an OCS for extra security (in which *K* is more than half of *N*), you can still share the keys it protects simultaneously amongst multiple modules as long you have enough unused cards to form a *K/N* quorum for the additional hardware security modules. For example, with a 3/5 OCS, you can load keys onto 3 hardware security modules because, after loading the key on the first device, you still have 4 cards left. After loading the key on a second device, you still have 3 cards left. After loading the key onto a third device, you have only 2 cards left, which is not enough to create the quorum required to load the key onto a fourth device.

If a card becomes damaged, you can replace the whole OCS if you have authorization from the ACS belonging to that Security World.

> **i** You can only replace OCSs that were created by Security Worlds that have the OCS/softcard replacement option enabled. For more information, see OCS and softcard replacement.

2.4.2.2. Using Operator Card Sets for high availability

If you cannot risk the failure of a smart card, but some keys must remain accessible at all times, you can create a 1/2 OCS.

Use the first card as the working card and store the second card in a completely secure environment. If the working card fails, retrieve the spare second card from storage, and use it until you re-create a new set of 2 cards (see Replacing an Operator Card Set or recovering keys to softcards).

> **i** You can only replace OCSs that were created by Security Worlds that have the OCS/softcard replacement option enabled. For more information, see OCS and softcard replacement.

2.4.2.3. Using persistent Operator Card Sets

If you create a standard (non-persistent) OCS, you can only use the keys protected by that OCS while the last required card of the quorum remains loaded in the card reader. The keys protected by this card are removed from the memory of the hardware security module as

soon as the card is removed from the card reader, which provides added security.

If you create a *persistent* OCS, the keys protected by a card from that OCS persist after the card is removed from the smart card reader.

This enables:

- The use of the same smart card in several hardware security modules at the same time
- Several users to load keys onto the same hardware security module at the same time.

The Security World Software maintains strict separation between the keys loaded by each user, and each user only has access to the keys protected by their OCS.

Keys protected by a persistent card are automatically removed from the hardware security module:

- When the application that loaded the OCS closes the connection to the hardware security module
- After a time limit that is specified when the card set is created
- When an application chooses to remove a key
- When the HSM is cleared. See Manually removing keys from an HSM for more information
- If there is a power loss to the module, for example, due to power outage.

> **ⓘ** Some applications automatically remove a key after each use, reloading it only when required. Such applications do not benefit from persistent OCSs. The only way of sharing keys between hardware security modules for such applications is by having multiple smart cards in an OCS.

Although the hardware security module stores the key, the key is only available to the application that loaded it. To use keys protected by this card in another application, you must re-insert the card, and enter its passphrase if it has one. Certain applications only permit one user at a time to log in, in which case any previously loaded persistent OCS used in that application is removed before the user is allowed to log in with a new OCS.

2.4.2.4. Manually removing keys from an HSM

You can manually remove all keys protected by persistent cards by clearing the hardware security module. For example, you could:

- Run the command `nopclearfail --clear --all`
- Turn off power to the hardware security module

Any of these processes removes all keys protected by OCSs from the hardware security module. In such cases, all users of any applications using the hardware security module must log in again.

Persistence is a permanent property of the OCS. You can choose whether or not to make an OCS persistent at the time of its creation, but you cannot change a persistent OCS into a non-persistent OCS, or a non-persistent OCS into a persistent OCS.

A Security World can contain a mix of persistent and non-persistent card sets.

2.4.3. Using passphrases for extra security

You can set individual passphrases for some or all the cards in an OCS.

You can change the passphrase for a card at any time provided that you have access to the card, the existing passphrase, and a hardware security module that belongs to the Security World to which the card belongs. For more information, see Changing card and softcard passphrase.

> **ℹ** Some applications do not support the use of passphrases.

2.4.3.1. Maximum passphrase length

> **ℹ** The maximum passphrase length limitation is not applicable to software versions before Security World Software v11.72.

passphrases are limited to a maximum length of 254 characters, when using the following commands:

- `new-world`
- `createocs`
- `cardpp`
- `ppmk`
- `racs`

Other commands are unaffected.

You can still use and edit existing passphrases that are longer than 254 characters.

Prior to Security World Software v11.72, we set no absolute limit on the length of passphrases, although individual applications may not accept passphrases longer than a specific number of characters. Likewise, the Security World does not impose restrictions on

which characters you can use in a passphrase, although some applications may not accept certain characters.

Entrust recommends that your password only contains 7-bit ASCII characters:

`A-Z, a-z, 0-9, ! @ # $ % ^ & * - _ + = [ ] { } | \ : ' , . ? / ` ~ " < > ( ) ;`

2.4.3.2. passphrase penalty timer

The HSM maintains a penalty time, measured in seconds and based on the number of failed PINs. Each failed attempt to enter a passphrase adds 4 seconds to the penalty time.

The penalty timer has a 14s penalty threshold, the first 3 failed passphrase verifications do not incur a penalty delay. Before verifying a passphrase, the HSM waits for the current penalty timer to be below 14s. The penalty time decays over time.

> **ⓘ** A HSM only has a small number of command processing threads, related to the kind of hardware in use (for example, 9 threads on an nShield Solo). Once all of these are waiting for a penalty to expire, any other submitted commands will be forced to wait. This can mean that even if penalty time isn't large, the total delay experienced by clients may be substantial.

2.4.4. Using softcard-protected keys

If you want to use passphrases to restrict key access but avoid using physical tokens (as required by smart-card protection), you can create a *softcard-protected key*.

A *softcard* is a file containing a logical token that you cannot load without a passphrase. You must load the logical token to authorize the loading of any key that is protected by the softcard. Softcard files:

- Are stored in `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**).
- Have names of the form `softcard_hash` (where `hash` is the hash of the logical token share).

Softcard-protected keys offer better security than module-protected keys and better availability than OCS-protected keys. However, because softcard-protected keys do not require physical tokens to authorize key-loading, OCS-protected keys offer better security than softcard-protected keys.

The passphrase of a softcard is set when you generate it, and you can use a single softcard to protect multiple keys. Softcards function as persistent 1/1 logical tokens, and after a soft

card is loaded, it remains valid for loading its keys until its `KeyID` is destroyed.

2.4.5. NVRAM key storage

Application keys protected by an nShield HSM are stored in an encrypted format and copied automatically to the remote file system. A hardware security module, such as an nShield HSM, and/or an OCS protects the keys, as described in the preceding sections. You can also store application keys within the nonvolatile memory of a suitable hardware security module.

NVRAM-stored keys are encrypted in exactly the same way as application keys that are pro tected by the unit. The encrypted application key on the unit is replaced by a file containing the name of the NVRAM file that contains the application key. This file allows applications to use NVRAM-stored keys in the same way as keys stored in the remote file system. You can protect an NVRAM-stored key with either the Security World or an OCS.

> **ℹ** NVRAM-stored keys differ from standard application keys only in their storage location. They still require protection by the unit or an OCS.

Use of an NVRAM-stored key is the same as for any other key protected by an nShield HSM Security World.

NVRAM key storage:

- Offers no additional security benefits above those offered by the standard Security World Software mechanisms
- Is available for only a limited number of keys
- Reduces a Security World's ability to offer load-balancing and recovery
- Requires backup and recovery procedures in addition to any that you implement for data stored on the client computer.

> **❗** Do not store keys in NVRAM unless you must do so to satisfy regulatory requirements.

NVRAM key storage was introduced only for those users who must store keys within the physical boundary of a hardware security module to comply with regulatory requirements. NVRAM-stored keys provide no additional security benefits and their use exposes your ACS to increased risk. Storing keys in nonvolatile memory also reduces load-balancing and recovery capabilities. Because of these factors, we recommend you always use standard Security World keys unless explicitly required to use NVRAM-stored keys.

2.5. Scalability

A Security World is scalable. You can add multiple hardware security modules to a server and share a Security World across multiple servers. You can also add OCSs and application keys at any time. You do not need to make any decisions about the size of the Security World when you create it.

To share a Security World across multiple clients:

- Ensure each client has at least one hardware security module configured
- Copy the Security World data to each client
- Load the Security World onto the hardware security modules for each client.

If you create cards or keys in a Security World from a client rather than on the hardware security module (using the command line, the Microsoft CSP wizard, or KeySafe), you must transfer the files from the client to the remote file system, unless the client is already on the same computer as a remote file system.

To provide access to the same keys on every server, you must ensure that all changes to the client data are propagated to the remaining servers. If your clients are part of a cluster, then the tools provided by the cluster should synchronize the data.

There is no risk of an attacker obtaining information by snooping on the network, as the data is only ever decrypted inside a hardware security module. Alternatively, you can maintain copies of the data on different clients.

> ℹ️ We provide the `rfs-sync` command-line utility to synchronize `opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**) between a cooperating client and the remote file system it is configured to access. Run `rfs-sync` whenever a cooperating client is initialized, to retrieve data from the remote file system, and also whenever a client needs to update its local copy of the data (or, if the client has write access, to commit changes to the data).

If you want to make cards or keys which are normally created from the client available from the front panel of the hardware security module, we recommend that you use client cooperation to automate the copying of files to the device.

2.5.1. Load-sharing

If you have more than one hardware security module configured with a client, your applications (that have been integrated with the Security World Software) can make use of the

load-sharing features in the Security World Software to share the cryptography between them. Two approaches are supported:

- API specific load-sharing modes
- HSM Pool mode: a more generic load-sharing approach for module protected keys introduced with module firmware version 2.65.2.

> ℹ | Some applications may not be able to make use of these features.

HSM Pool mode is supported on all major APIs except Java (i.e. nCipherKM JCA/JCE CSP). When HSM Pool mode is enabled for an API, the application sees the HSMs in the Security World as a single resource pool. A significant benefit is that when a failed HSM is restored to the Security World or a new HSM is added to the Security World, it is automatically added to the resource pool making it available for cryptographic operations without restarting the application (i.e. failback support). The pool of HSMs can be viewed as a single resource using the command `enquiry --pool`.

> ⊘ | For certain situations, `enquiry` and `nfkminfo` fields for modules in the pool have a particular meaning:
> * `Module #1: Not Present` indicates that there are no HSMs in the pool.
> * `hardware status: network error` on a module indicates that the Security World has changed, for example another HSM has been removed from the Security World. Run the same utility (`enquiry` or `nfkminfo`) directly on the module to obtain the hardware status for the module.

2.6. Robustness

Cryptography must work 24 hours a day, 7 days a week, in a production environment. If something does go wrong, you must be able to recover without compromising your security. A Security World offers all of these features.

2.6.1. Backup and recovery

The Security World data stored on the file system and remote file system of the hardware security module is encrypted using the Security World key.

You should regularly back up the data stored in the Key Management Data directory with your normal backup procedures. It would not matter if an attacker obtained this data because it is worthless without the Security World key, stored in your hardware security module, and the Administrator cards for that Security World.

When you create a Security World, it automatically creates recovery data for the Security World key. As with all host data, this is encrypted with the same type of key as the Security World key. The cryptographic keys that protect this data are stored in the ACS. The keys are split among the cards in the ACS using the same *K*/*N* mechanism as for an OCS. The ACS protects several keys that are used for different operations.

The cards in the ACS are only used for recovery and replacement operations and for adding extra hardware security modules to a Security World. At all other times, you must store these cards in a secure environment.

> In FIPS 140 Level 3 Security Worlds, the ACS or an OCS is needed to control many operations, including the creation of keys and OCSs.

2.6.2. Replacing a hardware security module

If you have a problem with a hardware security module, you can replace it with another hardware security module of the same type by loading the Security World data on the remote file system onto the replacement device. Alternatively, you may be able to erase the Security World from the device that has the problem, return the device to its default state and then reload the Security World on the same device.

If you have more than one hardware security module configured with a client and you use one of the load-sharing modes identified above, then your client application is resilient to the failure of individual hardware security modules. If you use HSM Pool mode, then an nShield HSM can be replaced and returned to the HSM pool without restarting the client application.

For information about replacing a hardware security module, see Adding or restoring an HSM to the Security World.

2.6.3. Replacing the Administrator Card Set

If you lose one of the smart cards from the ACS, or if the card fails, you must immediately create a replacement set using either:

- The front panel controls of the nShield HSM
- The KeySafe **Replace Administrator Card Set** option
- `racs` utility (see Replacing the Administrator Card Set).

> You should also use the front panel controls of the nShield HSM, `racs` or the KeySafe **Replace Administrator Card Set** option to migrate the ACS from standard nShield cards to nShield Remote Administration

Cards. Authorization needs to take place using the local slot of an HSM.

When using the `racs` utility, you cannot redefine the quantities in a K of N relationship for an ACS. The K of N relationship defined in the original ACS persists in the new ACS.

A hardware security module does not store recovery data for the ACS. Provided that *K* is less than *N* for the ACS, and you have at least *K* cards available, a hardware security module can re-create all the keys stored on the device even if the information from other cards is missing.

The loss or failure of one of the smart cards in the ACS means that you must replace the ACS. However, you cannot replace the ACS unless you have:

- The required number of current cards
- Access to their passphrases.

Although replacing the ACS deletes the copy of the recovery data on your host, you can still use the old ACS with the old host data, which you may have stored on backup tapes and other hosts. To eliminate any risk this may pose, we recommend erasing the old ACS as soon as you create a new ACS.

2.6.4. Replacing an Operator Card Set or recovering keys to softcards

If you lose an Operator Card, you lose all the keys that are protected by that card. To prevent this, you have the option to store a second copy of the working key that the recovery key protects in a Security World. Similarly, you can recover keys protected by one softcard to another softcard.

The ability to replace an OCS is an option that is enabled by default during Security World creation (see OCS and softcard replacement). You can only disable the OCS replacement option during the Security World creation process. You cannot restore the OCS replacement option, or disable this option, after the creation of the Security World.

You can only recover keys protected by an OCS to another OCS, and not to a softcard. Likewise, you can only recover softcard-protected keys to another softcard, and not to an OCS.

To create new copies of the keys protected by the recovery key on an OCS, you can use either:

- The front panel controls of the nShield HSM

- The `rocs` command-line utility.

> ℹ️ It is not possible to recover PKCS #11 keys using the front panel controls of the nShield HSM. You must use the `rocs` command-line utility.

2.6.4.1. The security of recovery and replacement data

Replacing OCSs and softcards requires authorization. To prevent the duplication of an OCS or a softcard without your knowledge, the recovery keys are protected by the ACS.

However, there is always some extra risk attached to the storage of any key-recovery or OCS and softcard replacement data. An attacker with the ACS and a copy of the recovery and replacement data could re-create your Security World. If you have some keys that are especially important to protect, you may decide:

- To issue a new key if you lose the OCS that protects the existing key

- Turn off the recovery and replacement functions for the Security World or the recovery feature for a specific key.

You can only generate recovery and replacement data when you create the Security World or key. If you choose not to create recovery and replacement data at this point, you cannot add this data later. Similarly, if you choose to create recovery and replacement data when you generate the Security World or key, you cannot remove it securely later.

If you have not allowed recovery and replacement functionality for the Security World, then you cannot recover any key in the Security World (regardless of whether the key itself was created as recoverable).

The recovery data for application keys is kept separate from the recovery data for the Security World key. The Security World always creates recovery data for the Security World key. It is only the recovery of application keys that is optional.

2.7. Audit Logging

Use of nShield HSMs in regulated environments where there is a requirement to provably log events in the HSMs can make use of the Audit Logging facility. This facility provides the following features:

- Tamper evident logging of relevant nCore command execution on the HSM

- Tied to Security World

- Traceability of cryptographic key lifetime
  - Authorization for key usage
  - Key loading onto HSM
  - Optional logging of key usage
  - Key destruction
- Compatibility with syslog and SIEM infrastructures
  - Logs produced in Common Event Format (CEF)
- Public key log verification without need for generating HSM.

For further information, see Audit Logging.

2.8. KeySafe and Security Worlds

KeySafe provides an intuitive and easy-to-use graphical interface for managing Security Worlds. KeySafe manages the Security World and the keys protected by it. For more information about using KeySafe, see Using KeySafe.

> ℹ️ Most applications store only their long-term keys in the Security World. Session keys are short term keys generated by the application which are not normally loaded into the Security World.

When you use KeySafe to create cards or keys, the data is written to the Key Management Data directory on the computer on which you run KeySafe. An nShield HSM can only use this data when it is transferred to the remote file system (if it is on a different computer), from where it is loaded automatically onto the unit. For this reason, you may find it most convenient to run KeySafe on the same computer as the remote file system.

Although you may use KeySafe to generate keys, it is your chosen application that actually uses them. You do not need KeySafe to make use of the keys that are protected by the Security World. For example, if you share a Security World across several hardware security modules, you do not need to install KeySafe on every computer. To manage the Security World from a single computer, you can install KeySafe on just that one computer even though you are using the Security World data on other computers.

KeySafe enables you to:

- Create OCSs
- List the OCSs in the current Security World
- Change the passphrase on an Operator Card
- Remove a lost OCS from a Security World

- Replace OCSs

- Erase an Operator Card

- Add a new key to a Security World

- Import a key into a Security World

- List the keys in the current Security World

- Delete a key from a Security World.

KeySafe does not provide tools to back up and restore the host data or update hardware security module firmware, nor does KeySafe provide tools to synchronize host data between servers. These functions can be performed with your standard system utilities.

In addition to KeySafe, we also supply command-line utilities to manage the Security World; for more information about the supplied utilities, see Supplied utilities. Current versions of these tools can be used interchangeably with the current version of KeySafe.

You can also perform all the tasks available from KeySafe using the front panel controls of the unit, except for adding, importing and deleting keys.

2.9. Applications and Security Worlds

A Security World can protect keys for a range of industry standard applications. For details of the applications that are currently supported, visit https://nshieldsupport.entrust.com.

We have produced Integration Guides for many supported applications. The Integration Guides describe how to install and configure an application so that it works with Entrust hardware security modules and Security Worlds.

For more information about the Entrust range of Integration Guides:

- Visit https://nshieldsupport.entrust.com.
- Contact Support.

2.10. The nShield PKCS #11 library and Security Worlds

> Do not use PKCS #11 to perform any task that requires an Administrator Card. Use the equivalent nShield utilities instead.

Many applications use a PKCS (Public Key Cryptography Standard) #11 library to generate and manage cryptographic keys. We have produced an nShield version of the PKCS #11 library that uses the Security World to protect keys.

Enabling a PKCS #11 based application to use nShield hardware key protection involves configuring the application to use the nShield PKCS #11 library.

The nShield PKCS #11 library treats a smart card from an OCS in the current Security World as a PKCS #11 token. The current PKCS #11 standard only supports tokens that are part of a 1-of-*N* card set, however the nShield PKCS #11 library has vendor specific extensions that support *K*/*N* card sets, see nShield PKCS #11 library with the preload utility.

A Security World does not make any distinction between different applications that use the nShield PKCS #11 library. Therefore, you can create a key in one PKCS #11 compliant application and make use of it in a different PKCS #11 compliant application.

2.11. Risks

Even the best-designed tools cannot offer security against every risk. Although a Security World can control which user has access to which keys, it cannot prevent a user from using a key fraudulently. For example, although a Security World can determine if a user is authorized to use a particular key, it cannot determine whether the message that is sent with that key is accurate.

A Security World can only manage keys that were created inside the Security World. Keys created outside a Security World, even if they are imported into the Security World, may remain exposed to a security risk.

Most failures of security systems are not the result of inherent flaws in the system, but result from user error. The following basic rules apply to any security system:

- Keep your smart cards safe.
- Always obtain smart cards from a trusted source: from Entrust or directly from the smart card manufacturer.

    > **ℹ** nShield Remote Administration Cards can only be supplied by Entrust.

- Never insert a smart card used with key management products into a smart card reader you do not trust.
- Never insert a smart card reader you do not trust into your hardware security module.
- Never tell anyone your passphrase.
- Never write down your passphrase.
- Never use a passphrase that is easy to guess.

    > **ℹ** If you have any doubts about the security of a key and/or Security World, replace that key and/or Security World with a newly generated

one.

3. The front panel interface

This chapter describes the nShield HSM user interface, including the front panel controls. You are also shown how to use a keyboard to control the unit.

3.1. Front panel controls

Figure 1. nShield HSM front panel controls

Label	Description
A	Power button
B	Warning LED (orange)
C	Display screen
D	Touch wheel
E	Status indicator LED (blue)
F	Display navigation button (left)
G	Display navigation button (right)
H	Select button
I	Slot for smart cards
J	Obsolete. Button retained for backward compatibility. No functionality on the nShield HSM.
K	USB connector

3.2. Display screen and controls

When the unit is powered, the display screen displays a menu or a dialog.

Each menu or dialog includes onscreen navigation labels that appear at the bottom of the display screen, on either side next to the display navigation buttons. Press the button next to the label to perform the action specified by the label.

To go back to the previous dialog or menu screen, use the navigation button to the left of the screen. To confirm a dialog value or select a menu option, either:

- Press the navigation button to the right of the screen.
- Touch the Select button.

Use the touch wheel to changes values or move the cursor on the display screen. To confirm a value, touch the Select button.

3.2.1. Menu screens

You can access menus from the display screen.

Menus are displayed as a list of selectable options. An onscreen arrow points to the currently selectable option. If the menu has more than four options, an arrow indicates the direction in which more options are available.

To select a menu option:

1. Move the indicator arrow up or down with the touch wheel.
2. When the indicator arrow points to the option you want to select, either:
   ◦ Press the navigation button to the right of the screen (labeled onscreen as **SELECT**).
   ◦ Touch the Select button.

At the top right of the display screen, a number sequence indicates the path to the current option. The last digit of the sequence shows the location of the menu you are currently viewing. The top level menu has no numbers, but when you select the System menu, the number **1** is shown.

The preceding digits in the sequence show the position of each option in turn that was selected in previous menu screens to reach the current menu. For example, the sequence **1-2** shows that the indicator is on the second option of the menu that was reached by selecting the first option on the top-level menu.

For a map of the menu screens, see the *Installation Guide*.

3.2.2. Dialogs

For some tasks, a dialog is displayed onscreen. When the dialog opens, the cursor is in the first field. To change and then enter values:

1. Use the touch wheel to change the displayed value of the fields.
2. Touch the Select button to enter the displayed value and move to the next field in the dialog.

Repeat the procedure to enter all necessary values in the dialog.

3.2.3. Information display

When you use a dialog to request information (for example, a log or details of a key), there is often too much information to display onscreen. In such cases, only the first part of the information is displayed.

To view the rest of the information:

• Use the touch wheel to scroll the displayed information in the direction indicated by the onscreen arrows.

• When an **Options** label is displayed, press the right-hand navigation button to see a menu of navigation options. You can normally choose to go to the top, to the bottom, or to a specified line in the display.

The numbers of the lines currently being displayed onscreen are shown at the left of the screen. They are followed in parentheses (**( )**) by the total number of lines available for display.

3.3. Using the front panel controls

You can use the front panel controls to configure the unit and to perform other tasks described in this guide. When the unit is working over the network with another computer (a client computer), you can program and control the unit as if it were part of the client computer.

> **ℹ** If the unit is powered down while you are logged in, you are logged out automatically.

3.3.1. Start-up information

When you turn on power to the unit and it has completed its initialization, the lower part of the display screen shows basic start-up information about the unit.

There is a series of start-up information topics available. By default, the first displayed topic is the current **System time**. Use the touch wheel to view the other start-up information topics.

3.3.2. Administrative control of the unit

You can view and control the status of the unit by using the front panel controls and menu options.

Tasks	Action
Understand and control the power status of the unit	Use the Power button to power up the unit.
	If the Power button is not illuminated, the unit is not powered. The Power button flashes intermittently as the unit powers up. It also flashes when the unit is in standby mode. For more information about the Power button, see the *Installation Guide*.

Tasks	Action
Control access to the unit	You can control access to the menus on the unit and the Power button on the front panel by using **System > System configuration > Login settings**. When **UI Lockout with OCS** has been enabled, you must log in with an authorized Operator Card before you can access the menus. You can still view information about the unit on the start-up screen. When you are logged in, you can log out and leave the unit locked. When **UI Lockout without OCS** has been enabled, you cannot access the menus, but you can still view information about the nShield HSM on the start-up screen. The only way to disable this setting (apart from returning the HSM to factory state) is to push an updated configuration file to the nShield HSM. See About user privileges and *ui_lockout* for more information. **Power button lockout** can be enabled and disabled independently when UI Lockout allows access to the menus.
Unlock the unit	When **UI Lockout with OCS** has been enabled and you have logged out, the display screen displays the label **Login** next to the right-hand navigation button. Press the right-hand navigation button, then insert an Operator Card that has been authorized for login, and follow the onscreen instructions.
Log out of the unit	Select **Logout**. This option is not available if UI Lockout with OCS has not enabled.
Put the unit in standby mode	Press the Power button or select **System > Shutdown/Reboot > Shutdown**.
Restore the unit to its original configuration	Select **System > System configuration > Default config**.
Restore the unit to its factory state	Select **System > Factory state**.
Clear the memory of the internal hardware security module	Run `nopclearfail --clear --all` or select **HSM > HSM reset**.
View information about the current state of the internal hardware security module	Select **HSM > HSM information**.
View information about the current state of the system	See the next section.
Set the Real-Time Clock on the unit	Select **Security World mgmt > Admin operations > Set secure RTC**. You need to reboot the nShield 5c after setting the RTC.

Tasks	Action
Change the mode of the unit	Select **HSM > Set HSM mode**. • Select **Operational** mode to run the unit normally. • Select **initialization** mode to configure the unit with software utilities rather than the front panel.

3.3.3. Viewing the current status of the unit

To view information about the current state of the system, from the main menu select **System > System information**. Select an option to view the associated information as follows:

Option	Description
View system log	Displays the system log.
View hardserver log	Displays the module hardserver log.
Display tasks	Displays the tasks that the system is currently performing.
Component versions	Displays the version numbers of the various system software components.
View h/w diagnostics	Displays the following environmental information about the module: • The current temperature at the left and right sensors • The minimum and maximum previous temperature at each sensor • The voltage on each power rail • The speed of each fan.
View tamper log	Displays the tamper log.
View unit id	Displays the ID of the unit.

3.3.4. Viewing the mode of the unit

See Identifying the current mode.

3.4. Using a keyboard to control the unit

You can connect a keyboard to the USB connector on the front panel. You can connect either a US or a UK keyboard. To configure the unit for your keyboard type, select **System > System configuration > Keyboard layout** and then choose the keyboard type you require.

When you have connected a keyboard and configured the unit for its use, you can enter numbers and characters directly into the display. You can also control the unit by using the

following keystrokes:

Keystroke	Use
F1	Same as pressing the left-hand navigation button on the front panel.
F2	Same as pressing the right-hand navigation button on the front panel.
F3	Same as touching the Select button.
Esc	Same as pressing the left-hand navigation button on the front panel.
Enter	Where the Select button is active, same as pressing Select: where Button B is active, same as pressing Button B.
Up arrow	Moves the indicator upwards in a menu.
Down arrow	Moves the indicator downwards in a menu.
Tab	Moves the cursor to the next field in a dialog.
Shift-Tab	Moves the cursor to the previous field in a dialog.
PgUp	Displays the previous screen.
PgDn	Displays the next screen.

4. Physical security of the HSM

This chapter provides a brief overview of the physical security measures that have been implemented to protect your nShield HSM. You are also shown how to check the physical security of your nShield HSM.

The tamper detection functionality on the nShield HSM provides additional physical security, over and above that provided by the holographic security seal, and alerts you to tampering in an operational environment. There is a removable lid on top of the nShield HSM, protected by the security seal and tamper switches. To prevent the insertion of objects into the nShield HSM, baffles are placed behind vents.

To optimize their effectiveness, use the physical security measures implemented on the nShield HSM in association with your security policies and procedures. For more information about creating and managing security policies, see the *Security Policy Guide* on the NIST CMVP website.

> **ⓘ** Currently, the FIPS 140 Level 3 boundary is at the internal module. Future software releases may move the FIPS boundary so that it includes the entire nShield HSM chassis.

> **ⓘ** For more information about FIPS 140, see http://csrc.nist.gov/publications/fips/fips140- 2/fips1402.pdf.

4.1. Tamper event

The nShield HSM offers several layers of tamper protection. The outer boundary of the box is tamper-responsive. When tampered, the unit ceases to provide cryptographic functionality, alerts the operator of the event, and ultimately forces the operator to reset the unit to factory defaults. Movements/vibrations, or replacing the fan tray module or a PSU, does not activate the tamper detection functionality.

If a tamper event does occur, you can use the Security World data stored on the RFS and the Administrator Card Set to recover the keys and cryptographic data.

4.1.1. nShield HSM lid is closed

If the nShield HSM is powered, a tamper event has occurred, and the lid is closed, the unit will automatically reset to a factory state.

Should this happen, examine your unit for physical signs of tampering (see Physical security

checks).

If you discover signs of tampering do *not* attempt to put the unit back into operation. The date and time of the tamper event are recorded in the log (see Logging, debugging, and diagnostics).

> The tamper-responsiveness circuitry has a Real Time Clock that is synchronised to the system time of the nShield HSM, however the times associated with events in the tamper log may still have slight offsets to times recorded in other log files.

If there are signs of tampering, and the tamper event occurred:

- During transit from Entrust, contact Support.
- After installation, refer to your security policies and procedures.

For more information about creating and managing security policies, see the *Security Policy Guide*.

You require a quorum of the Administrator Card Set (ACS) to restore the key data and reconnect the nShield HSM to the network.

4.1.2. nShield HSM lid is open

If the nShield HSM is powered, a tamper event has occurred, and the lid is open, the following message is displayed onscreen:

```
Unit lid is open
```

An open lid indicates that the physical security of the unit is compromised. You may want to examine your unit for other physical signs of tampering (see Physical security checks). Do *not* attempt to put the unit back into operation.

The date and time of the tamper event are recorded in the log files (see Logging, debugging, and diagnostics). If the tamper event occurred:

- During transit from Entrust, contact Support.
- After installation, refer to your security policies and procedures. For more information about creating and managing security policies, see the *Security Policy Guide* on the NIST CMVP website.

After closing the lid you must reboot the nShield HSM. The unit will then automatically reset to a factory state. If the lid remains open, the above message will remain on the screen and

all button presses are ignored.

4.2. Physical security checks

Check the physical security of your nShield HSM before installation and at regular intervals afterwards. For an alternative presentation of the physical security checks described here, see the *Physical Security Checklist*. For more information about tamper events, and what actions to take if you discover signs of tampering, see Tamper event.

To determine if the security of the nShield HSM is compromised:

1. Check that the physical security seal is authentic and intact. Look for the holographic foil bearing the nCipher logo. Look for cuts, tears and voiding of the seal. The seal is located on the top of the nShield HSM chassis.

   For information about the appearance of intact and damaged security seals, see the *Physical Security Checklist*.

2. Check that the metal lid remains flush with the nShield HSM chassis.

3. Check all surfaces — the top, bottom and sides of the nShield HSM — for signs of physical damage.

4. Check that there are no signs of physical damage to the vents, including attempts to insert objects into the vents.

4.3. Replacing the fan tray module and PSU

You can replace the fan tray module or a power supply unit (PSU) **without** activating a tamper event as both are outside the security boundary. You can access:

* The PSU(s) from the rear of the nShield HSM.
* The fan tray module through the removable front vent.

Should a problem occur with the fan tray module or a PSU, contact Support **before** taking further action. For more information about replacing the fan tray module or a PSU, see the *Fan Tray Module Installation Sheet* or the *Power Supply Unit Installation Sheet*.

> ℹ️ The fan tray module contains back-up batteries providing reserve capacity (a guaranteed minimum of 3 years) for tamper detection functionality even when the nShield HSM is in an unpowered state.

The tamper protection circuitry remains fully operational if the nShield HSM is placed on standby while a replacement operation is performed (whether you are replacing the fan tray module or one of the two PSUs, in the case of dual PSU units).

> ℹ️ Provided that the nShield HSM is connected to the mains power supply, it displays an onscreen error message when back-up battery power is low. The Status LED also displays a low power warning. For more information, see the *Installation Guide*.

4.3.1. Replacing the fan tray module

It is not necessary to remove mains power to replace a fan tray module (we recommend that you power down the unit into standby state using the front panel power button). However, if mains power is removed then a replacement fan tray module **must be installed within an hour** to ensure that a tamper event is not activated. If put in standby state the time required to change fan tray module is unlimited. For more information about replacing the fan tray module, see the *Fan Tray Module Installation Sheet*.

4.3.1.1. Fan tray module error messages

If you receive any of the following error messages on the nShield HSM display, accompanied by the orange warning LED, follow the related action in the table below:

Error message	Action
Single fan fail	Contact Support
Many fans fail	Replace fan tray
Battery power low	Consider replacing fan tray during the next scheduled service/maintenance period.
System Shutdown Both fans in a pair had failed	Replace fan tray

If the error message is **Single fan fail**, the nShield HSM can continue operating under the specified operating environment. Although you are advised to contact Support, the limited nature of such a failure means you can replace the fan tray module at your convenience.

If the error message is **Many fans fail**, you must replace the fan tray module immediately.

If the error message is **Battery Power low**, this indicates that one or both of the backup batteries located on the fan tray module (required only when the nShield HSM is removed from mains power) is running low.

The **Battery Power low** indication has no detrimental affect on the nShield HSM performance whilst the unit remains powered. Entrust recommend customers should consider replacing the fan tray module during the next service/maintenance.

If two fans fail from a redundant pair, the nShield HSM will display the error message **Many fans have failed** for a few seconds and it will then shutdown. On reboot, the nShield HSM will then display the error messages **System Shutdown** and **Both fans in a pair had failed**. In this situation the fan tray module must be replaced immediately.

4.3.2. Replacing the PSU

If you have a dual PSU nShield HSM, you do not have to remove power to the functioning PSU while replacing a faulty PSU. Tamper detection functionality will operate normally throughout the PSU replacement process. If you decide to remove power from both PSUs, tamper detection functionality will continue to operate normally for at least 3 years, as the fan tray module provides back-up capacity for this circuitry. For more information about replacing the PSU, see the *Power Supply Unit Installation Sheet*.

4.3.2.1. PSU error messages

If a PSU fails, an orange warning LED comes on and an error message is displayed on the nShield HSM display. Although you are advised to contact Support, the unit can continue to operate normally and you can replace the failed PSU at your convenience. There is no need to power down the unit when you replace the failed PSU.

In addition to the orange warning LED, an audible warning is given when a PSU fails on an nShield HSM. The audible warning is turned off when you navigate to the Critical errors screen.

4.3.3. Battery life when storing the nShield HSM

If a nShield HSM has been in storage for an extended period of time the fan tray module may need replacement.

Entrust guarantees a *minimum* battery life of three years, even if the nShield HSM remains in storage and is not connected to the mains power supply during this time.

5. Software installation

See the appropriate *Installation Guide* for your nShield module for more about installing the Security World software.

After you have installed the software, you must complete further Security World creation, configuration and setup tasks before you can use your nShield environment to protect and manage your keys.

5.1. After software installation

The *Installation Guide* provides brief explanations of how to perform all the post-installation tasks listed in this section. If this is the first time you are installing a unit and the Security World Software, or you are unfamiliar with the process, we recommend following the steps outlined in the *Installation Guide*.

After you have successfully installed the Security World Software, as described in the *Instal lation Guide*), complete the following steps to finish preparing your HSM for use:

1. Ensure that your public firewall is set up correctly. See the *Installation Guide* for your HSM for more information about firewall settings.
2. Perform the necessary basic HSM-client configuration tasks, as described in Basic HSM and remote file system (RFS) configuration.
3. Create and configure a Security World, as described in Creating a Security World.
4. Create an OCS, as described in Creating Operator Card Sets (OCSs).
5. Complete additional necessary HSM-client configuration tasks:

   a. To configure the unit so that it works with the client machine, see Configuring the nShield HSM to use the client.

   b. To configure client computers so that they work with the unit, see Configuring client computers to use the nShield HSM.

   For this release, you must generate a new client configuration file to take advantage of new functionality. To generate a new client configuration file, back up your existing configuration file and run the command `cfg-mkdefault`. This generates a template for the configuration file into which you can copy the set tings from your old configuration file.

   c. To enable the TCP sockets for Java applications (including KeySafe), run the com-

mand:

```
config-serverstartup -sp
```

For more information, see Client configuration utilities.

When all additional HSM configuration tasks are completed, you can:

1. Stop and then restart the hardserver, as described in Stopping and restarting the hardserver.
2. Test the installation and configuration. See the *Installation Guide* for your HSM for more information.

6. Client Software and module configuration

This chapter describes how to configure the internal security module of the nShield HSM and the client to communicate with each other, after you have installed the HSM and the Security World Software.

For more information about installing the HSM, see the *Installation Guide*. If you are configuring an HSM and client for the first time, or you want to complete a basic installation quickly, see the *Installation Guide*.

> **ℹ** The nShield HSM provides significant performance improvements, and can be deployed successfully with existing nShield products. Customers wishing to take advantage of these performance improvements **must** update their client machines with the latest Security World Software.

6.1. About user privileges

Cryptographic security does not depend on controlling user privileges or access but maintaining the integrity of your system from both deliberate or accidental acts can be enhanced by appropriate use of (OS) user privileges.

Linux

There are three levels of user:

- Superuser
- nfast group user
- normal

Typically, normal users can carry out operations involving Security Worlds, cardsets and keys, but not create Security Worlds, keys and cardsets. nfast group users have enhanced access, enabling them to create Security Worlds, cardsets and keys. For example, encrypted copies of keys are held in `kmdata` (`/opt/nfast/kmdata`). Normal users only have read access to the files, whereas nfast group users have read and write access, enabling them to create and use keys. nfast group users can also change the mode of an HSM remotely.

Superuser access (for example, `root`) is required for such tasks as software installation, starting and stopping the hardserver and SNMP.

Windows

There are two levels of user in Windows:

- Administrator access

- Normal

Administrator access (an Administrator on Windows) is required for such tasks as:

- Software installation, starting and stopping the hardserver and SNMP.

- Typically, any operation that requires write access, such as the creation of Security Worlds, card sets and keys.

Typically, normal users can only carry out read only operations involving Security Worlds, card sets and keys. For example, encrypted copies of keys are held in **Key Management Data** (`C:\ProgramData\nCipher\Key Management Data`). The default permissions allow all users to read these files, enabling them to use keys but not create them. File permissions can be altered to restrict access to specific keys to specific users/groups.

> 🛈 You should use the `nt_privpipe_users` option to define the name of the user who can carry out privileged operations, for example, using the `nopclearfail` utility. See *nt_privpipe_users* for more information.

6.2. About client configuration

> 🛈 You can add more HSMs to a client and more clients to an HSM at any time.

The HSM and a client communicate by means of the hardserver, which handles secure transactions between the HSM and applications that run on the client. You must configure:

- Each client hardserver to communicate with the HSM that the client needs to use

- The HSM to communicate with clients that are allowed to use the HSM.

Information about the current configuration of the HSM or a client is stored in configuration files that are stored in specified file systems on the clients and on the HSM. For more information about the contents of configuration files, see HSM and client configuration files.

For information about configuring the HSM by importing an edited configuration file, see About user privileges.

6.2.1. Remote file system (RFS)

Each HSM must have a remote file system (RFS) configured. The RFS contains master copies of all the files that the HSM needs:

- The HSM configuration file

- Feature-enabling certificates

- The encrypted Security World and key data for Security Worlds created on the HSM.

The RFS normally resides on a client computer, but it can be located on any computer that is accessible on the network.

For more information about setting up the remote file system, see Configuring the Remote File System (RFS).

6.2.2. HSM configuration

The data that defines the configuration of the HSM hardserver is stored in a file on the HSM. This file is automatically:

- Updated when the HSM is configured from the front panel

- Exported to the remote file system (RFS) directory.

Each HSM has separate configuration files on the RFS, stored in the directories with names of the form `/opt/nfast/kmdata/hsm-ESN/config` (**Linux**) or `%NFAST_KMDATA%\hsm-ESN\config` (**Windows**) where `ESN` represents the electronic serial number of the HSM from which the files were exported. These directories can contain the following files:

Option	Description
`config`	The master configuration file. This contains the current configuration for the HSM. It is always present in the directory.
`config-name`	An alternative configuration file saved by the system.
`config.new`	A hand-edited configuration file that can be read by the HSM.

You normally configure the HSM using the front panel controls. However, in some cases (for example, if you need to configure an HSM remotely, or if you are importing a number of clients), you may prefer to edit the exported configuration file and then re-import the file into the HSM. For more information, see:

- About user privileges

- HSM and client configuration files.

6.2.3. Client configuration

The data that defines the configuration of the client hardserver is stored in a file on the

client's file system.

You must load the configuration file for the configuration to take effect. For information about loading a client configuration remotely, see Remote configuration of additional clients.

You can configure a client to use multiple HSMs. All the HSMs configured for use by a client can fail over if the application that uses them is set up appropriately.

For more information about the contents of the client configuration file, see HSM and client configuration files.

> **ℹ** You can also configure the client's hardserver by setting environment variables, as described in Setting environmental variables. Environment variable settings override settings in the client configuration files.

6.3. Basic HSM and remote file system (RFS) configuration

After installing the HSM hardware and software, there are several HSM and RFS configuration tasks you must perform. You perform these RFS tasks before:

- Creating the Security World and an Operator Card Set (OCS)
- Completing the process of configuring the HSM and client to work together.

6.3.1. Configuring the Ethernet interfaces

The HSM communicates with one or more clients. Each client is an Ethernet connected computer that has the Security World Software installed and configured. You must supply Internet Protocol (IP) addresses for the HSM and the client. Contact your system administrator for this information if necessary.

To configure the Ethernet interfaces (IPv4 and IPv6), see the *Installation Guide* for your HSM.

6.3.2. Optionally configure hardserver interfaces

By default, the hardserver listens on all interfaces. However, you can alter the hardserver settings. Altering the hardserver settings would prove necessary, for example, if you wanted to connect one of the Ethernet interfaces to external hosts.

Ensure that you have configured the Ethernet interfaces on the HSM before attempting to configure the hardserver. See the Installation Guide for more information about configuring

the Ethernet interfaces.

You can configure the following options to specify network interfaces on which the hard-server listens:

Option	Description
All interfaces	This option (which is the default) specifies that the hardserver listens on all interfaces.
IP address of interface #1	This option specifies that the hardserver listens only on interface 1. This option only appears if interface 1 has been configured.
IP address of interface #2	This option specifies that the hardserver listens only on interface 2. This option only appears if interface 2 has been configured.
Will not listen	This option specifies that the hardserver does not listen on any interfaces.

To define the interface and port on which the hardserver listens:

1. From the main menu, select **System > System configuration > Hardserver config**. The following screen appears:

   ```
   Hardserver config
     Select network I/F
     hardserver listens on:
     All interfaces
     Select TCP port: 9004
     CANCEL        FINISH
   ```

2. Select the network interfaces on which the hardserver is to listen.

   > For security reasons, do not allow the hardserver to listen on any interface that is to connect to the public Internet.

3. Press the Select button to move to the TCP port field, and set the port on which the hardserver is to listen. The default is 9004.

   > Make sure that your firewall settings are consistent with your port settings. See the Installation Guide for more about firewall settings.

4. When the network interface and port are correct, press the right-hand navigation button.

5. Press the right-hand navigation button again to continue.

6. You are asked if you wish to reboot the system now or later. Press the right-hand navigation button to reboot now.

6.3.3. Configuring the Remote File System (RFS)

The RFS contains the master copy of the Security World data for backup purposes. The RFS can be located on either a client or another network-accessible computer where the Security World Software is installed. If the RFS is on a client, the same file structure also contains the configuration files for that client.

> ℹ️ We recommend that you regularly back up the entire contents of the RFS. The `kmdata` (**Linux**) or `%NFAST_KMDATA%` (**Windows**) directory is required to restore the nShield HSM, or a replacement, to its current state, in case of failure.

You can specify a new remote file system, and modify or delete an existing remote file system configuration. To create or modify a remote file system configuration, specify the IP address of the computer on which the file system resides.

> ℹ️ You must have created an RFS on the client computer before you specify the IP address of the client.

For more information about the RFS and its contents, see:

- Remote file system (RFS)
- Location of Security World files.

The nShield HSM must be able to connect to TCP port 9004 of the RFS. If necessary, modify the firewall configuration to allow this connection on either the RFS itself or on a router between the RFS and the nShield HSM.

You can configure the connection to use secure authentication using software-based authentication or with an nToken (or local HSM) installed in the RFS. When enabled the nShield HSM not only examines the RFS's IP address, but also requires the RFS to identify itself using a signing key.

> ℹ️ If an nToken is installed in the RFS, it can be used to both generate and protect a key that is used for the impath communication between the nShield HSM and the RFS. Thus a strongly protected key is used at both ends of the impath. A local nShield PCIe or USB-attached HSM can also be used to perform the role of the nToken.

Obtain the following information about the nShield HSM before you set up an RFS for the

first time:

- The IP address

- The electronic serial number (ESN)

- The hash of the K_{NETI} key (HK_{NETI}). The K_{NETI} key authenticates the nShield HSM to clients. It is generated when the nShield HSM is first initialized from factory state.

If your network is secure and you know the IP address of the nShield HSM, you can use the `anonkneti` utility to obtain the ESN and hash of the K_{NETI} key by giving the following command on the client computer. For guidance on network security, see the *nShield Security Manual*.

```
anonkneti <Unit IP>
```

In this command, *<Unit IP>* is the IP address of the nShield HSM, which could be one of the following:

- an IPv4 address

- an IPv6 address, including a link-local IPv6 address

- a hostname

The command returns output in the following form:

```
A285-4F5A-7500 2418ec85c86027eb2d5959fef35edc5e1b3b698f
```

In this example output, `A285-4F5A-7500` is the ESN and `2418ec85c86027eb2d5959fef35edc5e1b3b698f` is the hash of the K_{NETI} key.

Alternatively, you can find this information on the nShield HSM startup screen. Use the touch wheel to scroll to the appropriate information.

When you have the necessary information, set up an RFS as follows:
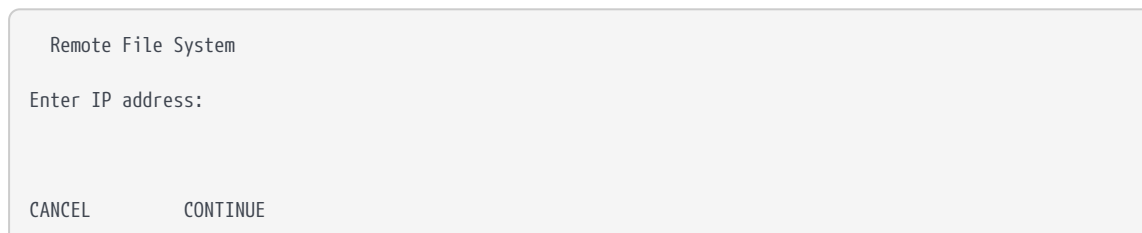
1. Prepare the RFS on the client computer (or another appropriate computer) by running the following command on that computer:

   ```
   rfs-setup <Unit IP> <EEEE-SSSS-NNNN> <keyhash>
   ```

   In this command:

   ◦ *<Unit IP>* is the IP address of the nShield HSM.

   ◦ *<EEEE-SSSS-NNNN>* is the ESN of the nShield HSM.

- ◦ *<keyhash>* is the hash of the K_{NETI} key.

2. On the nShield HSM display screen, use the right-hand navigation button to select **Sys tem > System configuration > Remote file system**, and enter the IP address of the client computer on which you set up the RFS:
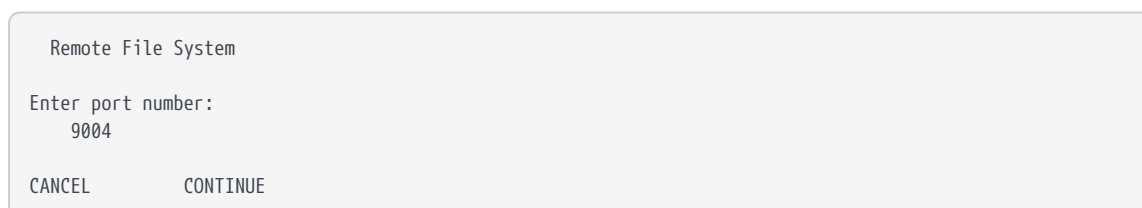
```
   Remote File System

Enter IP address:



CANCEL        CONTINUE
```

3. The next screen asks for the port number on which the RFS is listening. Enter the port number and press the right-hand navigation button to continue:
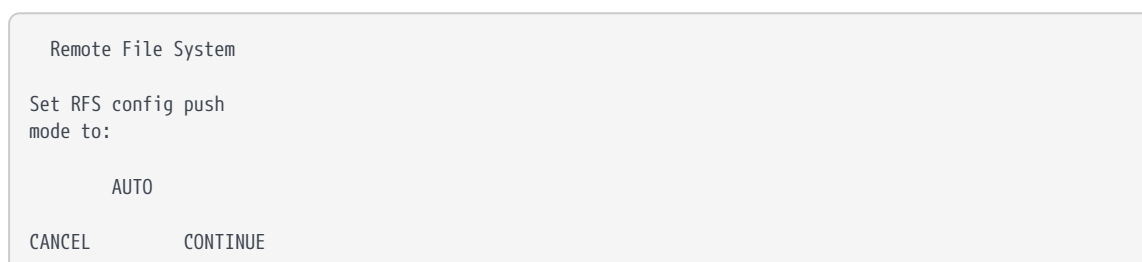
```
   Remote File System

Enter port number:
    9004

CANCEL        CONTINUE
```

> 🛈 | Leave the port number at the default setting of 9004.

4. Select the config push mode and press the right-hand navigation button to continue:
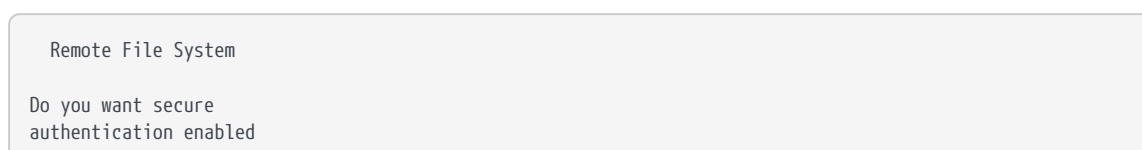
```
   Remote File System

Set RFS config push
mode to:

      AUTO

CANCEL        CONTINUE
```

Three options are available:

- ◦ AUTO: The RFS is only allowed to push configuration files to the nShield HSM if secure authentication is enabled. This is the default value.
- ◦ ON: The RFS is allowed to push configuration files to the nShield HSM.
- ◦ OFF: The RFS is not allowed to push configuration files to the nShield HSM.

5. You must confirm whether to enable or disable secure authentication when setting up the RFS. The following screen is displayed:

```
   Remote File System

Do you want secure
authentication enabled
```

```
on the RFS?

        YES
CANCEL      CONTINUE
```

    a. Select No and press the right-hand navigation button to configure the RFS without secure authentication. The authentication of the RFS will be based on the IP address only.

    b. Select Yes and press the right-hand navigation button to configure the RFS with secure authentication.

6. Skip this step if you have not selected secure authentication.

If an nToken is installed in the RFS, you will be asked to choose which authentication key to use. Select the desired option and press the right-hand navigation button:

```
>0DA8-A5AE-BA0D
 Software Key



BACK      SELECT
```

    a. The ESN of the nToken installed in the RFS.

    b. "Software Key" for software-based authentication.

If no nToken is installed in the RFS, then software-based authentication is automatically selected.

7. Skip this step if you have not selected secure authentication.

At the next screen, verify that the key hash displayed by the nShield HSM matches the RFS key hash:

```
Remote 0DA8-A5AE-BA0D
reported the key hash:
 9e0020264af732933574
 0cfe10446d33cea33f4a
Is this EXACTLY right?

CANCEL      CONFIRM
```

The RFS key hash is obtained by running the commands described below. Take a copy of the returned key hash and compare it to the value reported on the nShield HSM display.

With software-based authentication

    Run the following command on the RFS:

```
enquiry -m0
```

This command returns the software key hash, tagged as `kneti hash`, as part of its output, for example:

```
Server:
  enquiry reply flags   none
  enquiry reply level   Six
  ...
  kneti hash            d4c3d757a67416cb9ba31f33febd6ead688629e5
  ...
```

With nToken authentication

Run the following command on the RFS:

```
ntokenenroll -H
```

This command produces output of the form:

```
nToken module #1
nToken ESN:      0DA8-A5AE-BA0D
nToken key hash: 9e0020264af732933574
                 0cfe10446d33cea33f4a
```

Check that the ESN also matches the one reported on the nShield HSM display.

If the RFS key hash matches the one reported on the nShield HSM display, press the right-hand navigation button to continue the RFS configuration. Otherwise press the left-hand navigation button to cancel the operation.

8. The nShield HSM displays "Transferring files..." followed by a message reporting that the RFS has been set up. Press the right-hand navigation button again to exit.

After you have defined the RFS, the nShield HSM configuration files are exported automatically. See the *User Guide* for more about configuration files.

To modify the RFS at a later date, select **System > System configuration > Remote file system**, and then select the required action.

> You can allow other clients to access the remote file system and share Security World and key data that is stored in the `kmdata` (**Linux**) or `%NFAST_KMDATA%` (**Windows**) directory in the same way as the HSM. Clients that access data in this way are described as *cooperating clients*. To configure client cooperation, you need to know the details of each client including IP address and optionally their key hash and nTo-

> ken ESN.

6.3.4. Configuring log file storage

You can choose to store log files on both the HSM and RFS or on the HSM only.

To configure log file storage, use the right-hand navigation button to select **System > System configuration > Log config**. Then select one of:

1. **Log** to store log files on the HSM only.
2. **Append** to store log files on both the HSM and remote file system.

We recommend selecting **Append** because if you select **Log** you can only view the log file from the nShield HSM front panel. Moreover, the log file stored on the HSM is cleared every time it is powered down.

You may also additionally configure the logs to be sent to a remote syslog server, see Configuring Remote Syslog.

6.3.5. Setting the time and date

If you do not intend to use NTP time synchronization, set the time and date as described in this section. If you configure the nShield HSM to use NTP time synchronization, then the time and date will be maintained by NTP.

To set the time and date on the HSM as UTC:

1. Use the right-hand navigation button to select and display the **System** menu:

```
1-1
System configuration
System information
Login settings
Upgrade system
Factory state
Shutdown/Reboot
BACK              SELECT
```

2. Select **System configuration** to display the **System configuration** menu:

```
1-1-1
Network config
Hardserver config
Remote file system
Client config
Resilience config
Config file options
```

```
┌─────────────────────────────────────────────┐
│  BACK              SELECT                     │
└─────────────────────────────────────────────┘
```

3. Use the touch wheel to move the arrow to **Date/time setting**, and press the right-hand navigation button to select it. The **Set system date** screen is displayed:

```
Set system date
Please enter the
current UTC date as
DD/MM/YYYY:
 27/ 5/2013
CANCEL           NEXT
```

4. For each date field, use the touch wheel to set the value and move the cursor to the next field.

   When you have completed all the fields, press the right-hand navigation button to confirm the date. The **Set system time** screen is displayed:

```
Set system time
Please enter the
current UTC time as
hour/mins/seconds:
 18:08:19
CANCEL           FINISH
```

> ℹ️ Setting the time and date of the HSM as UTC does not reset the value of the Real Time Clock (RTC) on the HSM. The UTC date and time settings are used only in log messages.

6.3.6. Keyboard layout

You can connect a keyboard to the USB connector on the nShield HSM front panel. This enables you to control the nShield HSM using a special set of keystrokes instead of the standard front panel controls.

You can connect either a US or a UK keyboard. To configure the nShield HSM for your keyboard type, select **System > System configuration > Keyboard layout** and then choose the keyboard type you require.

6.4. Configuring the nShield HSM to use the client

You must inform the HSM hardserver of the location of the client computer.

You can configure the connection to use secure authentication using software-based authentication or with an nToken (or local PCIe HSM) installed in the client. When enabled

the nShield HSM not only examines the client IP address, but also requires the client to iden
tify itself using a signing key.

> **ℹ** If an nToken is installed in a client, it can be used to both generate and
> protect a key that is used for the impath communication between the
> nShield HSM and the client. Thus a strongly protected key is used at
> both ends of the impath. A local PCIe HSM can also be used to perform
> the role of the nToken.

> **ℹ** Software-based authentication is only supported from version 12.60.
> Previously enrolled clients using software-based authentication will
> need to be re-enrolled if an earlier version of Security World software is
> installed.

The client configuration process varies slightly depending on whether you are enrolling the
client with or without secure client authentication:

1. On the nShield HSM front panel, use the right-hand navigation button to select **Sys-
   tem > System configuration > Client config > New client**.

   The following screen is displayed:

   ```
   Client configuration

   Please enter your
   client IP address:



   CANCEL          NEXT
   ```

   Enter the IP address of the client, and press the right-hand navigation button.

2. Use the touch wheel to confirm whether you want to save the IP or not, and press the
   right-hand navigation button.

   ```
   Client configuration

   Do you want to save
   the IP in the config?
   (No for dynamic client
   IPs)
            No
   BACK            NEXT
   ```

3. Use the touch wheel to select the connection type between the nShield HSM and the
   client.

```
Client configuration

Please choose the
client permissions

     Unprivileged

BACK            NEXT
```

The following options are available:

Option	Description
Unprivileged	Privileged connections are never allowed.
Priv. on low ports	Privileged connections are allowed only from ports numbered less than 1024. These ports are reserved for use by root on Linux.
Priv. on any ports	Privileged connections are allowed on all ports.

> **ℹ** A privileged connection is required to administer the nShield HSM (for example, to initialize a Security World). If privileged connections are allowed, the client can issue commands that can interfere with the normal operation of the nShield HSM, such as clearing it. Entrust recommends that you allow only unprivileged connections unless you are performing administrative tasks.

4. When you have selected a connection option, press the right-hand navigation button.

   The following screen is displayed:

```
Client configuration

Do you want secure
authentication enabled
on this client?

        Yes
BACK            NEXT
```

   a. Select No and press the right-hand navigation button to configure the client without secure authentication. The authentication of the client will be based on the IP address only.

   b. Select Yes and press the right-hand navigation button to configure the client with secure authentication.

5. On the nShield HSM, enter the number of the port on which the client is listening (the default is 9004), and press the right-hand navigation button. The following screen is

displayed:

```
Client configuration

On what port is the
client listening?

       9004

CANCEL          NEXT
```

6. Skip this step if you have not selected secure authentication.

If an nToken is installed in the client, you will be asked to choose which authentication key to use. Select the desired option and press the right-hand navigation button:

```
>3138-147F-2D64
 Software Key



BACK      SELECT
```

    a. The ESN of the nToken installed in the client.

    b. "Software Key" for software-based authentication.

If no nToken is installed in the client, then software-based authentication is automatically selected.

> ℹ️ Software-based authentication is only supported from version 12.60.

7. Skip this step if you have not selected secure authentication.

The next screen asks you to verify that the key hash displayed by the nShield HSM matches the client key hash:

```
Remote 3138-147F-2D64
reported the key hash:
 691be427bb125f387686
 38a18bfd2eab75623320
Is this EXACTLY right?

CANCEL          CONFIRM
```

The client key hash is obtained by running the commands described below. Take a copy of the returned key hash and compare it to the value reported on the nShield HSM display.

With software-based authentication

Run the following command on the client:

```
enquiry -m0
```

This command returns the software key hash, tagged as `kneti hash`, as part of its output, for example:

```
Server:
  enquiry reply flags   none
  enquiry reply level   Six
  ...
  kneti hash            f8222fc007be38b78ebf442697e244dabded38a8
  ...
```

With nToken authentication

Run the following command on the client:

```
ntokenenroll -H
```

This command produces output of the form:

```
nToken module #1
nToken ESN:      3138-147F-2D64
nToken key hash: 691be427bb125f387686
                 38a18bfd2eab75623320
```

Check that the ESN also matches the one reported on the nShield HSM display.

If the client key hash matches the one reported on the nShield HSM display, press the right-hand navigation button to continue the RFS configuration. Otherwise press the left-hand navigation button to cancel the operation.

8. The nShield HSM displays a message reporting that the client has been configured. Press the right-hand navigation button again.

To modify or delete an existing client, select **System > System configuration > Client config** and perform the appropriate procedure.

If you want to use multiple clients with the nShield HSM, you must enable additional client licenses (see Enabling optional features). When you have additional client licenses enabled, to configure more clients, repeat the appropriate steps of the procedure described in this section for each client.

6.4.1. Remote configuration of additional clients

Additional clients can be added remotely provided that config push is enabled. This can be done from the RFS or a client computer (refer to the nShield *Remote Administration User Guide* for more information.).

> ℹ️ Before you can use multiple clients with the nShield HSM, you must enable the additional clients as described in Enabling optional features.

To add clients remotely:

1. Copy the HSM configuration file, `NFAST_KMDATA/hsm-ESN/config/config` (**Linux**) or `NFAST_KMDATA\hsm-ESN\config\config` (**Windows**), to `config.new` in the same directory.

   The path to the new file will be `NFAST_KMDATA/hsm-ESN/config/config.new` (**Linux**) or `NFAST_KMDATA\hsm-ESN\config\config.new` (**Windows**).

2. Add a new entry to `config.new` in the `hs_clients` section to contain the details of the client to be added.

   The following two entries must exist in the configuration file:

   ```
   addr=<client_IP>
   clientperm=permission_type
   ```

   Where:

   `<client_IP>` can be either the IP address of the client or `0.0.0.0`, `::`, or blank if the HSM is to accept clients identified by their key hash instead of their IP address.

   `0.0.0.0` or `::`, and blank result in the same behavior. You can only use them in the configuration file, you cannot enter these values in the front-panel user interface.

   The default is blank.

   If you set both the `<client_IP>` field (the client's IP address) and the key hash, the HSM must identify clients from both of these fields.

   `permission_type` defines the type of commands the client can issue (`unpriv`, `priv` or `priv_lowport`).

   a. If the client is using an nToken, two additional entries will need to be added to the configuration file:

      ```
      esn=nToken_ESN
      keyhash=nToken_keyhash
      ```

Where `nToken_ESN` is the ESN of the client's nToken and `nToken_keyhash` is the hash of the key that the client's nToken should authenticate itself with.

b. If the client is using software-based authentication, one additional entry will need to be added to the configuration file:

```
keyhash=software_keyhash
```

Where `software_keyhash` is the hash of the software generated key that the client should authenticate itself with. The `ESN` entry must be blank or omitted for software-based authentication.

> ℹ️ Each client entry after the first must be introduced by a line consisting of one or more hyphens.

3. Load the updated configuration file on to the nShield HSM. To do this, run the following command:

```
cfg-pushnethsm --address=<module_IP_address> <new_config_file>
```

In this command, `<module_IP_address>` is the IP address of the nShield HSM and `<new_config_file>` is the location of the updated configuration file (`config.new`).

Alternatively, you can load the configuration file using the nShield HSM front panel without enabling config push. The configuration file (`config.new`) must be in the `/opt/nfast/kmdata/hsm-ESN/config` (**Linux**) or `%NFAST_KMDATA%\hsm-ESN\config` (**Windows**) directory on the remote file system. To do this, select **System > System configuration > Config file options > Fetch configuration**.

> ℹ️ An SEE machine cannot be installed or configured using the fetch configuration option from the front panel, the config push feature must be used for this. See Remotely loading and updating SEE machines for more information.

6.5. Changing the nShield HSM configuration from the Front Panel to use a client

From the Front Panel, you can change the settings that your nShield HSM stores about its client.

> ❗ The Front Panel does not display all current properties for the client. Entrust recommends that you retrieve the current client settings before

> you start the update. See HSM and client configuration files. During the configuration update, check the current configuration file against the values displayed on the Front Panel. Check the values at each configuration step before proceeding to the next step and finally confirming the changes.

6.6. Configuring client computers to use the nShield HSM

Each client computer must be configured to use the internal security module of your nShield HSM. There are two methods for achieving this:

- Enrolling the client with the configuration file.
- Enrolling the client with command-line utilities.

6.6.1. Enrolling the client with the client configuration file

The client configuration files are in the `/opt/nfast/kmdata/config` (**Linux**) or `%NFAST_KM-DATA%\config` (**Windows**) directory on the client computer's file system.

> **ℹ** For this release, you must generate a new client configuration file to take advantage of the new functionality. To generate a new client configuration file, back up your existing configuration file and run the command `cfg-mkdefault`. This generates a template for the configuration file into which you can copy the settings from your old configuration file.

The `nethsm_imports` section defines the network HSMs that the client imports (See nethsm_imports). It can also be set up by the `nethsmenroll` utility.

- Edit the following mandatory fields: `local_module`, `remote_ip`, `remote_esn`, `remote_key-hash` and `privileged`. The default value for `remote_keyhash` (40 zeros) specifies that no authentication should occur.
- The `ESN` and `hash` of the HSM to import can be retrieved by running the command

```
anonkneti remote_ip
```

- If the client is to be enrolled with an nToken, open a command line window, and run the command: `ntokenenroll --H`. This command produces output of the form:

```
nToken module #1
nToken ESN: 3138-147F-2D64
```

```
nToken key hash: 691be427bb125f3876838a18bfd2eab75623320
```

- Enter the nToken's `ESN` in the field `ntoken_esn` in the config file.
- Each HSM entry after the first must be introduced by a line consisting of one or more hyphens, i.e. `---`.
- At the command line run the command `cfg-reread`, to reload the hardserver's configuration.
- Verify that the client can use the nShield HSM by running `enquiry`, which reports the HSM's status.

> **ⓘ** If the client is to be enrolled with either software-based authentication or no authentication, the `ntoken_esn` field must be left **empty**.

For information about configuration file contents, see HSM and client configuration files.

6.6.2. Enrolling the client from the command line

The `nethsmenroll` command-line utility edits the client hardserver's configuration file to add the specified nShield HSM. For more information about the options available to use with `nethsmenroll`, read the following section Client configuration utilities, or run the command:

```
nethsmenroll --help
```

To retrieve the nShield HSM's `ESN` and `HKNETI`, run the command

```
anonkneti <Unit IP>
```

This command produces output of the form:

```
3138-147F-2D64 691be427bb125f38768638a18bfd2eab75623320
```

If the nShield HSM's `ESN` and `HKNETI` are not specified, `nethsmenroll` attempts to contact the HSM to determine what they are, and requests confirmation.

1. If you are enrolling the client *with* an nToken, run the command:

   ```
   nethsmenroll --ntoken-esn <nToken ESN> [Options] --privileged <Unit IP> <Unit ESN> <Unit KNETI HASH>
   ```

2. If you are enrolling the client *without* an nToken, i.e. software-based authentication or no authentication, run the command:

```
nethsmenroll [Options] --privileged <Unit IP> <Unit ESN> <Unit KNETI HASH>
```

These commands produce output of the form:

```
OK configuring hardserver's nethsm imports.
```

6.6.3. Client configuration utilities

We provide the following utilities for client configuration:

Utility	Description
nethsmenroll	This utility is used to configure the client to communicate with the nShield HSM.
config-serverstartup	This utility is used to configure the client's hardserver to enable TCP sockets.

6.6.3.1. nethsmenroll

The `nethsmenroll` command-line utility edits the client hardserver's configuration file to add the specified nShield HSM. If the nShield HSM's `ESN` and `HKNETI` are not specified, `nethsmenroll` attempts to contact the nShield HSM to determine what they are, and requests confirmation.

Usage:

```
nethsmenroll [Options] --privileged <hsm-ip> <hsm-esn> <hsm-kneti-hash>
```

Options	Description
-m|--module=MODULE	Specifies the local HSM number to use (default is 0 for dynamic configuration by hardserver).
-p|--privileged	Causes the hardserver to request a privileged connection to the nShield HSM (default `unprivileged`).
-<hsm-ip>	The IP address of the nShield HSM, which could be one of the following: • an IPv4 address • an IPv6 address, including a link-local IPv6 address • a hostname
-r|--remove	Deconfigures the specified nShield HSM.
-f|--force	Forces reconfiguration of an nShield HSM already known.

Options	Description
`--no-hkneti-confirmation`	Does not request confirmation when automatically determining the nShield HSM's `ESN` and `HKNETI`
	ⓘ This option is potentially insecure and should only be used on secure networks where there is no possibility of a man-in-the-middle attack. For guidance on network security, see the *nShield Security Manual*.
`-V|--verify-nethsm-details`	When the `ESN` and `HKNETI` have been provided on the command line, verifies that the selected HSM is alive, reachable and matches those details.
`-P|--port=PORT`	Specifies the port to use when connecting to the given nShield HSM (default `9004`).
`-n|--ntoken-esn=ESN`	Specifies the `ESN` of the nToken to be used to authenticate this client. If the option is omitted, then software authentication will be used instead.

6.6.3.2. config-serverstartup

The `config-serverstartup` command-line utility automatically edits the `[server_startup]` section in the local hardserver configuration file in order to enable TCP ports for Java and KeySafe. Any fields for which values are not specified remain unchanged. After making any changes you are prompted to restart the hardserver.

Run `config-serverstartup` using commands of the form:

```
config-serverstartup [OPTIONS]
```

For more information about the options available to use with `config-serverstartup`, run the command:

```
config-serverstartup --help
```

6.7. Configuring NTP in the nShield HSM

The nShield HSM has a standard NTP client that can be configured to support synchronization of system time on the HSM with one or more NTP servers. Network Time Protocol (NTP) is intended to synchronize all participating computers to within a few milliseconds of Coordinated Universal Time (UTC). System time on the nShield HSM is independent of the Real Time Clock (RTC) in the HSM and is used for log messages and front panel display.

> ℹ️ Entrust recommends that the NTP Server(s) are trusted servers within your local network, not internet time servers.

After configuring NTP the HSM has to be re-started for the configuration to take effect. When starting up after being configured, the NTP client can make a step change to the system time to bring it into line with that of the NTP server(s). At all other times, the NTP client will only slew (gradually change) the system time. When using NTP there should be no need to set the system time by setting time and date from the front panel of the nShield HSM.

Before configuring NTP you must ensure the following:

- **config push** is enabled for the remote computer used to configuring NTP. Refer to the nShield *Remote Administration User Guide* for more information.
- The client computer enabled for **auto push** is configured for Privileged connections, see Configuring the nShield HSM to use the client, so that the nShield HSM can be rebooted from the client computer.

6.7.1. Using the NTP configuration tool

NTP is configured using the `cfg-pushntp` utility on a client computer.

Utility	Description
`cfg-pushntp`	This tool enables or disables NTP time synchronization on the specified nShield HSM. When enabling NTP synchronization, the IP addresses of up to 3 NTP servers may be specified. ℹ️ The new NTP settings will take effect the next time the target nShield HSM is restarted.

Usage:

```
cfg-pushntp -a ADDR [-p PORT -k -m MODULE] -1 ADDR [-2 ADDR -3 ADDR] enable
```

```
cfg-pushntp -a ADDR [-p PORT -k -m MODULE] disable
```

Options:

Field	Description
`enable	disable`	Enable or disable NTP service on the nShield HSM.
`-a|--address=ADDR`	IP address of nShield HSM to configure NTP on. This can be an IPv4 or an IPv6 address, or it can be a hostname that resolves to the HSM's IP address.

Field	Description
`-p|--port=PORT`	Set port to use to connect to the nShield HSM (default=9004).
`-k|--use-kneti`	Use KNETI to authenticate ourselves.
`-m|--module=MODULE`	Set the HSM to use for KNETI authentication. The default is HSM 1. This option can only be used with the `--use-kneti` option.
`-1|--ntp1=ADDR`	IP address of NTP server. This can be an IPv4 or an IPv6 address.
`-2|--ntp2=ADDR`	IP address of NTP server. This can be an IPv4 or an IPv6 address.
`-3|--ntp2=ADDR`	IP address of NTP server. This can be an IPv4 or an IPv6 address.

For example, running the command:

Linux

```
cfg-pushntp --address=192.30.100.150 --ntp1=192.23.24.256 enable
```

Windows

```
cfg-pushntp.exe --address=192.30.100.150 --ntp1=192.23.24.256 enable
```

Returns:

```
The requested NTP configuration changes have been uploaded and will take effect when the target nShield HSM is
restarted.
```

6.7.2. Restarting the nShield HSM

After configuring NTP, restart the nShield HSM, for example using `nethsmadmin --mod-ule=<MODULE> --reboot` (see Utilities for network-attached HSMs). Once the HSM has rebooted and the syslog output is available, check that there are no NTP failures reported in the syslog output.

6.7.3. Enable NTP for IPv6

You can use HSM images that were built with buildroot from version 13 onwards with NTP over IPv6.

To run ntpd as a non-root user, start ntpd as root and use `switch -u` on the command line to change to the "ntp:ntp" non-root user.

To set NTP up for IPv6:

1. Reset the HSM to factory state using the UI (menu: 1-5). See

2. Assign an IPv6 address to the HSM on interface #1 (eth0).

   For example: `fd01:11::260:e0ff:fe81:4ac1`.

3. Assign an IPv6 address on the same subnet to the RFS host machine that will act as the NTP server.

   For example: `fd01:11::ae9f:98bc:e202:9941`.

4. Enroll the HSM on the RFS using the following commands:

   Ensure that the displayed kneti hash is not just a string of zeroes. If it is, reset the HSM to factory state.

   ```
   opt/nfast/bin/anonketi <HSM IP>

   opt/nfast/bin anonketi <HSM IPv6 address>

   sudo /opt/nfast/bin/rfs-setup --force <HSM IPv6 address> `/opt/nfast/bin/anonketi <HSM IPv6 address>`

   sudo /opt/nfast/bin/nethsmenroll -fp <HSM IPv6 address>
   ```

5. Add the RFS IPv6 address to the HSM (menu: 1-1-3)
   - RFS IP address: `<RFS IPv6 address>`
   - Set RFS config push mode to: `AUTO`
   - Do you want secure authentication enabled on RFS? `YES`

6. Add the RFS IPv6 address to the HSM as a client (menu: 1-1-4)
   - Client IP address: `<RFS IPv6 address>`
   - Do you want to save IP in the config? `Yes`
   - Client permissions: `Priv on any port`
   - Do you want secure authentication enabled on this client? `No`

7. Check that the RFS can detect the HSM. If any errors appear, such as `AccessDenied`, wait for one minute and then run the command again.

   ```
   /opt/nfast/bin/enquiry
   ```

8. Ensure that auto-push is enabled on the HSM. If auto-push is not enabled, `cfg-pushntp` reports an `UnknownHostVolume` error.

   a. Toggle auto-push on:

   System > System configuration > Config file options > Setup auto push > Yes

b. Add the RFS IPv6 address:

System > System configuration > Config file options

9. Using the HSM front panel, set the date and time to be different from those on the RFS host machine. This will help you verify that HSM is updating from the RFS when completing this task.

   You can view the HSM's status, including the date and time, by scrolling through the front (top) menu on the front panel.

10. Install NTP on the RFS.

    For an Ubuntu 18.04 host:

    a. Run the following command:

    ```
    sudo apt install ntp
    ```

    b. Add the following lines to the `/etc/ntp.conf` file.

    ```
    0.au.pool.ntp.org
    1.au.pool.ntp.org
    2.au.pool.ntp.org
    3.au.pool.ntp.org
    ```

    c. Restart the NTP service:

    ```
    sudo service ntp restart
    ```

    d. Verify that the service successfully restarted and is running:

    ```
    sudo service ntp status
    ```

    e. Check that the firewall is not blocking UDP and TCP ports 123.

11. Push the NTP configuration to the HSM from the RFS host machine:

    ```
    /opt/nfast/bin/cfg-pushntp --address=<HSM IPv6 address> --ntp1=<RFS IPv6 address> enable
    ```

12. Reboot the HSM from the RFS host machine. The HSM updates its system time from the NTP server by running the `/sbin/setostime` script when it starts up from the `/etc/inittab` file.

    ```
    sudo opt/nfast/bin/nethsmadmin --reboot --module 1
    ```

13. When the HSM has rebooted, check that the date and time in the front (top) menu
    match the date and time on the RFS host machine.

6.8. Configuring Remote Syslog

The nShield HSM can be configured to send logs directly to a remote syslog server, listen-
ing on a User Datagram Protocol (UDP) port, by editing the `remote_sys_log` section of the
config file: `remote_syslog_ip` can be an IPv4 or IPv6 address.

This behavior can be configured in addition to storing the log files on the RFS (i.e. you can
configure the logs to be sent to a remote syslog server regardless of whether the nShield
HSM logs are configured to be stored on the HSM or the RFS). For further information see
Configuring log file storage.

There is no additional formatting of log messages (the logs sent are the same log messages
that will appear on the unit or the RFS). It is the responsibility of the remote syslog server /
SIEM application to format, sort and aggregate the incoming log messages.

To configure an HSM to send logs to a remote syslog server:

1. Configure the remote host to receive logs from the HSM.

   For instructions, see the documentation for the operating system of the remote host.

2. Make sure that the HSM can communicate with an RFS host so that you can push the
   new configuration file to the HSM. See *Basic HSM, RFS and client configuration* in the
   Installation Guide for your HSM.

3. If your HSM configuration file predates the functionality to send logs to a remote sys-
   log server, manually create the `remote_sys_log` section in the config file for the remote
   module.

   See remote_sys_log.

4. In the `remote_sys_log` section, add the following settings:

   ```
   remote_sys_log_ip=REMOTE_SYSLOG_SERVER_IP
   remote_sys_log_port=REMOTE_SYSLOG_SERVER_PORT
   ```

5. Run the following command to push the new config file to the HSM:

   ```
   cfg-pushnethsm
   ```

   > ℹ️ If you are using an older version of the Security World software

with a Connect image that supports remote syslog, you will see an error message: **unrecognized section name: 'remote_sys_log'**. Use the following command to push the updated configuration file:

```
cfg-pushnethsh --force
```

> If you are using a version of the Security World software that supports remote syslog with an image that does not support remote syslog, the configuration file will be pushed to the HSM but the HSM will reject it. You will see that the upgraded configuration file on the RFS is unchanged.

To turn off sending logs to a remote syslog server, remove the entries from the `remote_sys_log` section of the configuration file and push the updated configuration file.

6.9. Setting up client cooperation

If you do not need to allow multiple clients access to your remote file system (RFS), you only need to follow the instructions provided in Configuring the Remote File System (RFS) to initialize your RFS. If you need to allow other clients to access your RFS (that is, able to access the RFS to share key data), complete the following steps:

1. Configure the RFS to accept access by cooperating clients:

   ◦ For every authenticated client (with write access and K_{NETI} authorization) that needs to be a client of this remote file system, run the command:

   ```
   rfs-setup --gang-client <client_IP_address> <EEEE-SSSS-NNNN> <keyhash>
   ```

   In this command:

   ▪ `<client_IP_address>` is the IP address of the client. This can be an IPv4 or IPv6 address.

   ▪ `<EEEE-SSSS-NNNN>` is the ESN of the nToken used by the client when using a nToken K_{NETI} key to authenticate itself. When using software-based authentication, it must be empty (that, is "") or can be omitted altogether.

   ▪ `<keyhash>` is the hash of the software or module K_{NETI} key used by the client.

   ◦ For every unauthenticated client (with write access but without K_{NETI} authorization), run the command:

Linux

```
rfs-setup --gang-client --write-noauth <client_IP_address>
```

Windows

```
rfs-setup.exe --gang-client --write-noauth <client_IP_address>
```

> **ℹ** The `--write-noauth` option should be used only if you believe your network is secure. This option allows the client you are configuring to access the RFS without K_{NETI} authorization.

To limit a gang-client to read-only, use the `--readonly` flag.

2. On each client that is to be a cooperating client, you must run the `rfs-sync` command-line utility with appropriate options:

   ° for clients using a software K_{NETI} key to authenticate themselves to the RFS, run the command with the default options:

   ```
   rfs-sync --setup <RFS_IP_ADDRESS>
   ```

   ° for clients using a module K_{NETI} key to authenticate themselves to the RFS, run the command:

   ```
   rfs-sync --setup --authenticate --module=<MODULE> <RFS_IP_ADDRESS>
   ```

   In this command:

   - `<RFS_IP_ADDRESS>` is the IP address of the RFS.
   - `<MODULE>` is the local module to use for authentication.

   ° for clients to authenticate the RFS using software-based authentication, use the `--rfs-hkneti=HKNETI` option to specify the hash of the software K_{NETI} key of the RFS.

   ° for clients to authenticate the RFS using nToken authentication, use the `--rfs-esn=ESN` and `--rfs-hkneti=HKNETI` options to specify the `ESN` and hash of the K_{NETI} key of the nToken installed in the RFS.

The `rfs-sync` utility uses lock files to ensure that updates are made in a consistent fashion. If an `rfs-sync --commit` operation (the operation that writes data to the remote file system) fails due to a crash or other problem, it is possible for a lock file to be left behind. This would cause all subsequent operations to fail with a lock time-out error.

The `rfs-sync` utility has options for querying the current state of the lock file, and for delet-

ing the lock file; however, we recommend that you do not use these options unless they are necessary to resolve this problem. Clients without write access cannot delete the lock file.

For more information about the `rfs-sync` utility, see rfs-sync.

To remove a cooperating client so the RFS no longer recognizes it, you must:

- Know the IP address of the cooperating client that you want to remove
- Manually update the `remote_file_system` section of the hardserver configuration file by removing the following entries for that particular client:

  Linux

  ```
  remote_ip=<client_IP_address>
  remote_esn=keyhash=00000000000000000000000000000000000000000
  native_path=/opt/nfast/kmdata/local
  volume=kmdata-local
  allow_read=yes
  allow_write=yes
  allow_list=yes
  is_directory=yes
  is_text=no
  ```

  and

  ```
  remote_ip=<client_IP_address>
  remote_esn=keyhash=00000000000000000000000000000000000000000
  native_path=/opt/nfast/kmdata/local/sync-store/
  volume=kmdata-backup
  allow_read=yes
  allow_write=yes
  allow_list=yes
  is_directory=yes
  is_text=no
  ```

  Windows

  ```
  remote_ip=client_IP_address
  remote_esn=keyhash=00000000000000000000000000000000000000000
  native_path=%NFAST_KMDATA%\local
  volume=kmdata-local
  allow_read=yes
  allow_write=yes
  allow_list=yes
  is_directory=yes
  is_text=no
  ```

  and

  ```
  remote_ip=client_IP_address
  remote_esn=keyhash=00000000000000000000000000000000000000000
  native_path=%NFAST_KMDATA%\local\sync-store
  volume=kmdata-backup
  allow_read=yes
  ```

```
allow_write=yes
allow_list=yes
is_directory=yes
is_text=no
```

6.9.1. Useful utilities

6.9.1.1. anonkneti

To find out the ESN and the hash of the K_{NETI} key for a given IP address, use the `anonkneti` command-line utility. A manual double-check is recommended for security.

The IP address could be one of the following:

- an IPv4 address
- an IPv6 address, including a link-local IPv6 address
- a hostname

6.9.1.2. rfs-sync

This utility synchronises the `opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**) folder between a cooperating client and the remote file system it is configured to access. It should be run when a cooperating client is initialised in order to retrieve data from the remote file system and also whenever a client needs to update its local copy of the data or, if the client has write access, to commit changes to the data.

6.9.1.2.1. Usage

```
rfs-sync [-U|--update] [-c|--commit] [-s|--show] [--remove] [--setup [setup_options] ip_address]
```

6.9.1.2.2. Options

-U|--update

These options update local key-management data from the remote file system.

> If a cooperating client has keys in its `kmdata/local` directory that are also on the remote file system, if these keys are deleted from the remote file system and then `rfs-sync --update` is run on the client, these keys remain on the client until manually removed.

-c|--commit

These options commit local key-management data changes to the remote file system, and update the client from the remote file system.

-s|--show

These options display the current synchronisation configuration.

--setup

This option sets up a new synchronisation configuration. Specifics of the configuration can be altered using `setup_options` as follows:

-a|--authenticate

This set-up option specifies the use of a module KNETI key to authenticate this client to the RFS. By default the software KNETI key is used instead.

-m|--module=module

This option selects the local module to use for authentication. The default is 1. This option can only be used with the `--authenticate` option.

-p|--port=port

These options specify the port on which to connect to the remote file system. The default is 9004.

--rfs-hkneti=HNETI

This option specifies the hash of the K_{NETI} key to use for nToken or software-based authentication of the RFS.

--rfs-esn=ESN

This option specifes the `ESN` of the nToken to use for authentication of the RFS.

ip_address

This option specifies the IP address of the remote file system, which could be one of the following:

- an IPv4 address
- an IPv6 address, including a link-local IPv6 address
- a hostname

--remove

This option removes the synchronisation configuration.

A client can use `rfs-sync --show` to display the current configuration, or `rfs-sync --remove` to revert to a standalone configuration. Reverting to a standalone configuration leaves the

current contents of the Key Management Data directory in place.

The `rfs-sync` command also has additional administrative options for examining and removing lock files that have been left behind by failed `rfs-sync --commit` operations. Using the `--who-has-lock` option displays the task ID of the lock owner. As a last resort, you can use the `rfs-sync` command-line utility to remove lock files. In such a case, the `--kill-lock` option forcibly removes the lock file.

> ⓘ   The lock file can also be removed via menu item 3-3-2, **Remove RFS Lock**: this executes the `rfs-sync --kill-lock` command.

6.9.2. Setting environmental variables

This section describes how to set Security World Software-specific environment variables. You can find detailed information about the environment variables used by Security World Software in Environment variables.

Linux

You can set Security World Software-specific environment variables in the file `/etc/nfast.conf`. This file is not created by the installation process: you must create it yourself. `/etc/nfast.conf` is executed by the start-up scripts of nShield HSM services as the root user. This file should only contain shell commands that can safely be run in this context. `/etc/nfast.conf` should be created with access permissions that allow only the root user to write to the file.

> ⓘ   Ensure that all variables are exported as well as set.

Windows

You can set Security World Software-specific environment variables as follows:

1. Open the **System** dialog by clicking **System** in the control panel menu.
2. Select the **Advanced** tab and click the **Environment Variables** button.
3. To add a variable, click **New**. Alternatively, to edit an existing variable select an entry in the **System Variables** list and click **Edit**.
4. In the **Variable Name** field, type or edit the name of the environment variable (for example, **NFAST_HOME**).
5. In the **Variable Value** field, type or edit the value to use.
6. Click the **OK** button to set the value, and then click the **OK** button to close the dialog.
7. Open the **Administrative Tools** dialog by clicking the **Administrative Tools** icon in the Control Panel

8. Open the **Services** console by clicking the **Services** icon.

9. From the displayed list of services, select the **nFast Server** icon, and select **Restart the service**.

6.9.3. Logging and debugging

The nShield HSM and applications that use it generate log files. You can view the logs using the unit front panel. Application log messages are handled on the client.

The Security World Software generates logging information that is configured through a set of four environment variables:

- NFLOG_FILE
- NFLOG_SEVERITY
- NFLOG_DETAIL
- NFLOG_CATEGORIES

> **ⓘ** If none of these logging environment variables are set, the default behavior is to log nothing, unless this is overridden by any individual library. If any of the four logging variables are set, all unset variables are given default values.

Detailed information about controlling logging information by means of these environment variables is supplied in Logging, debugging, and diagnostics.

Some components of the Security World Software generate separate debugging information which you can manage differently. Debugging information for applications is handled on the client. If you are setting up the unit to develop software that uses it, you should configure debugging before commencing software development.

6.9.4. Configuring Java support for KeySafe

To use KeySafe, follow the instructions in Using KeySafe.

6.10. Routing

If you have configured only one network interface, you do not need to configure a static route for the unit, although you can do so if you wish. If you have configured a second network interface, you can choose to configure a static route.

If the unit is to connect to a remote host or network that is unreachable through the default gateway, you must set up extra static routes in the system routing table.

To set up the Ethernet interfaces (IPv4 and IPv6), see the *Installation Guide* for your HSM.

After you have defined static routes, you should test them as described in Testing routes.

> **ⓘ** If you define, edit, or delete a route, you must reboot the unit before the route can be used and the routing table is updated.

6.10.1. Testing routes

When you have set up or edited a route, you should test the route.

6.10.1.1. Testing a route between the unit and the client

When you have installed the unit in its final location, you should test the connection between the unit and the client. You can do this from the client, as described in this section, or by using the **Ping remote host** option on the unit. To do this from the unit, select **System > System configuration > Network config > Ping remote host.**

You can also use the method described in this section to test the route between the client and a remote computer.

To test the route from the client to the unit, issue a `ping` command from the client for the IP address that you specified for the unit. (The format of the command and results may vary according to the platform that you are using.)

```
ping <xxx.xxx.xxx.xxx>
```

If the `ping` operation is successful, a message similar to the following is displayed:

```
Pinging xxx.xxx.xxx.xxx with 32 bytes of data
Reply from xxx.xxx.xxx.xxx: bytes=32 time=10ms TTL=125
Reply from xxx.xxx.xxx.xxx: bytes=32 time=10ms TTL=125
Reply from xxx.xxx.xxx.xxx10ms TTL=125
Reply from xxx.xxx.xxx.xxx10ms TTL=125
```

6.10.1.2. Testing a route between the unit and a remote host

When you have defined or edited a route from the unit to a remote computer, you should test it. To do this you can issue a `ping` command from the unit to the IP address of the host.

You can also use this method to test the connection between the unit and the client.

To test a route from the unit to a host:

1. Select **System > System configuration > Network config > Ping remote host**. The following screen appears:

```
Ping remote host
Select IP address:
 0. 0. 0. 0
RESET FINISH
```

2. Enter the IP address of the remote host.

3. Press **FINISH** to issue the ping. The following message appears:

```
Please wait, running ping
```

4. After a short wait, a display similar to the following should appear, showing that the unit has managed to communicate with the host:

```
Ping xxx.xxx.xxx.xxx:
#0: rtt=0.0503 ms
#1: rtt=0.0503 ms
#2: rtt=0.0503 ms
#3: rtt=0.0503 ms
4 of 4 packets back.
 min avg max SD
 0.29 0.36 0.50 0.09 ms
```

If not all of the information is visible onscreen, use the touch wheel to scroll up and down the page.

5. Press the left-hand navigation button to return to the **Network config** menu.

6.10.2. Tracing network routes

You can trace network routes from the unit and from clients.

6.10.2.1. Tracing the route from the unit

You can trace the route taken from the unit to a remote computer. You can also use this method to trace the route from the unit to the client.

1. Select **System > System configuration > Network config > Trace route to host**. The following screen appears:

```
Trace route
Select IP address:
```

```
    0. 0. 0. 0
    CANCEL FINISH
```

2. Press the right-hand navigation button to issue the `traceroute`. The following message appears:

```
    Please wait, running traceroute.
```

3. After a short wait, a display similar to the following should appear, showing the IP addresses encountered along the route:

```
    1: xxx.xx.xxx.x
     0.40 ms
    2: *
    3: xx.xxx.xx.xxx
     2.1 ms
    4: xxx.xx.xxx.xxx
     2.4 ms
    BACK
```

If not all of the information is visible onscreen, use the touch wheel to scroll up and down the page.

4. Press the left-hand navigation button to return to the **Network config** menu.

6.10.2.2. Tracing the route from the client

You can trace the route from the client to the unit or (if the client is connected to the public Internet) to a remote computer.

To trace the route from the client to the unit, issue a `traceroute` command from the client for the IP address of the unit. (The format of the command, and results, may vary depending upon the platform that you are using.)

```
    tracert <xxx.xxx.xxx.xxx>
```

After a short wait, a display similar to the following should appear, showing the IP addresses encountered along the route:

```
    Tracing route to modulename (xxx.xxx.xxx.xxx)/ over a maximum of 30 hops:
     1 xxx.xxx.xxx.xxx (xxx.xxx.xxx.xxx) 1.457 ms 0.513 ms 0.311 ms
     2 xxx.xxx.xxx.xxx (xxx.xxx.xxx.xxx) 0.773 ms 0.523 ms 1.615 ms
```

6.10.2.3. Displaying the routing table

You can view details of all the IP addresses for which the internal security module has a route stored. The routing table includes entries for static routes (which are stored permanently) and local hosts to which the module has set up temporary routing entries.

To view the routing table:

1. Select **System > System configuration > Network config > Show routing table**. A display similar to the following appears:

```
Dest Gateway Flg
1 xxx.xxx.xxx.xxx
  xxx.xxx.xxx.xxx UG
2 xxx.xx.xx.x
  xxx.x.x.xxx UG
BACK
```

   If not all of the information is visible on the unit display screen, use the touch wheel to scroll up and down the page.

2. Press the left-hand navigation button to return to the **Network config** menu.

6.11. Configuring an nShield HSM using the Serial Console

On supported hardware variants (see *Model numbers* in the nShield Security World: nShield 5c v13.4.5 User Guide) there is an RJ45 serial port connector at the rear of the nShield HSM marked **Console**. The serial port provides access to a serial console command line interface that enables remote configuration of the unit whilst also facilitating status monitoring. Tasks which would typically require a physical presence in front of the HSM, including setting IP addresses, can be done remotely using the serial console.

Rear Panel

This functionality can help provide separation of duty between the data center technician installing the nShield HSM and the administrator configuring and using the HSM. The only required local activity is to connect the nShield HSM to power and to serial and Ethernet ports. Everything that can be configured using the front panel can then be configured remotely either over the serial interface, by using the `nethsmadmin` utility (see nethsmadmin) or by pushing an updated configuration file to the nShield HSM (refer to the nShield *Remote Administration User Guide* for more information.).

The Serial Console supports IPv4 and IPv6 addresses.

6.11.1. Serial port configuration

The RJ45 connector can be directly connected to your client machine or connected to a serial port aggregator for remote access.

To access the serial console command line interface, first determine the device name of the serial connection once it is connected to your client machine. Then configure the serial port connection in your serial port communication software with the following parameters:

- Line Speed (baud): 115200

This is the default baud rate. If you have manually changed the baud rate to 9600 (see here), enter this value instead.

- Data bits: 8
- Parity: None
- Stop bits: 1.

Once the connection is established, press **Return** until a login prompt is displayed. The login prompt should look like:

```
nethsm login:
```

6.11.2. Change the baud rate

To change the baud rate using the front panel, navigate to **System** > **System configuration** > **Remote config options** > **Serial Console** > **Serial baud rate** and select the required baud rate.

to change the baud rate remotely:

1. Copy the config file to `config.new`.
2. If the `config.new` file does not already contain the following section, add it to the file.

```
[cli]
# Start of the cli section
# The serial CLI baud rate configuration. Restart your CLI connection after
# changing.
# Each entry has the following fields:
#
# Set to "115200" or "9600" to select the relevant baud rate for the serial
# CLI connection. Note active CLI serial connections are broken upon the
# setting of a new baud rate.
baud_map=BAUDRATE
```

3. Edit the `baud_map` value.
4. Push `config.new` to the HSM using `cfg_pushnethsm`.

> **i** Changing the baud rate using the front panel creates the `[cli]` section in the config file if it does not already exist.

6.11.2.1. Troubleshooting

Error	Action
Nothing on the screen	Press **Enter** a few times. Make sure that the RJ45 connector is properly connected and that remote configuration is enabled on the nShield HSM, see Enabling and disabling the serial console.
Miscellaneous characters displayed on the screen	Make sure the serial port connection is configured correctly, see Serial port configuration.

6.11.3. Creating a serial console session

The username for accessing the serial console is `cli` and the default password is `admin`.

On first login you will be prompted to change the password for the `cli` user. The minimum length of the new password is 5 characters. You must enter a new password. You cannot reuse `admin` the first time you change the password. For guidance on selecting a password, see the *nShield Security Manual*.

Once you are successfully logged in to a serial console session you will see the welcome message:

```
Welcome to the nShield 5c Serial Console. Type help or ? to list commands.

(cli)
```

The serial console session will automatically logout after 180 seconds of inactivity. To manually end a session, use the **logout** command.

6.11.4. Enabling and disabling the serial console

The serial console interface can be enabled and disabled using the nShield HSM front panel.

* To enable the serial console interface, navigate to **System** > **System configuration** > **Remote config options** > **Serial Console** and set to **On**.
* To disable the serial console interface, set **Serial Console** to **Off**.

The serial interface is enabled by default and will turn back on with the default password if the unit is reset to its factory state. This means if you do not want the serial console enabled you will need to disable it after each factory state.

If you do not see the menu option **System** > **System configuration** > **Remote config options** > **Serial Console** on the front panel then this means that your nShield HSM does

not support the serial console feature (the hardware does not support serial access).

The availability of the serial console feature can also be checked remotely from an enrolled client by running the enquiry utility.

Feature availability	Enquiry output
Serial console available	... level six flags SerialConsoleAvailable ...
Serial console not available	... level six flags none ...

6.11.5. Serial console commands

The serial console command line interface provides the following commands:

Command	Description
`cs5`	Configures CodeSafe functionality on the nShield 5c
`date`	Get/set the HSM system time
`enquiry`	Prints enquiry data from the HSM
`esn`	Show the Electronic Serial Number (ESN) of the HSM
`expose`	Get/set the expose NIC `nethsm_settings` configuration
`factorystate`	Reset unit to its original (factory) state ℹ️ | This will reset the IP and serial console settings
`gateway`	Get/set the default IPv4 gateway address
`gateway6`	Get/set the default IPv6 gateway address
`getrtc`	Get the real time clock (RTC) of the nShield 5c Only available on the nShield 5c
`help` or `?`	List available commands with `help` or detailed help with `help cmd`
`hsmdiagnose`	Runs hsmdiagnose on the embedded HSM
`kneti`	Show the (hash of) Kneti (used for enrolling the HSM with clients)

Command	Description
ks5agent	Manage the KeySafe 5 agent Configure and enable a KeySafe 5 agent on this Connect to report to a central KeySafe 5 management platform
logs	Print nShield 5c diagnostic logs Only available on the nShield 5c
logout	Log out of the nShield HSM Serial Console
maintmode	Enable/disable maintenance mode Only available on the nShield 5c
netcfg	Get/set the IPv4 network interface configuration
netcfg6	Get/set the IPv6 network interface configuration
netdiagnose	Print network interface statistics
netenable	Enable IPv6
netlink	Get/set the network interface link, get the network interface MAC address
bondcfg	Get/set the HSM bond interface configuration
bondlink	Get/set the bond interface link
passwd	Set the serial console password
ping	Ping a remote host
push	Get/set the config push setting
reboot	Reboot the HSM
rfsaddr	Get/set the RFS IP address, port, optional secure authentication and push mode
route	Get/set IPv4 network routes
route6	Get/set IPv6 network routes
routing	Show the IPv4 routing table
routing6	Show the IPv6 routing table
setrtc	Set the RTC on the nShield 5c Only available on the nShield 5c Before you can use this command, you must put the nShield 5c into maintenance mode with maintmode

Command	Description
`stattree`	Run the `stattree` command
`sworldcheck`	Check for any Security World data on the HSM
`tamperlog`	Show the nShield HSM tamper log
`uptime`	Show how long the nShield HSM has been running (since last boot)
`version`	Show nShield HSM Serial Console version information
`watchdog`	Get/set the `nethsm_watchdog` configuration

For additional help on a command, run `help` command to see additional guidance on command usage, syntax and parameter documentation.

6.11.6. Using multiple modules

The hardserver can communicate with multiple modules connected to the host. By default, the server accepts requests from applications and submits each request to the first available module. The server can share load across buses, which includes the ability to share load across more than one module.

If your application is multi-threaded, you can add additional modules and expect performance to increase proportionally until you reach the point where cryptography no longer forms a bottleneck in the system.

6.11.6.1. Identifying modules

Modules are identified in two ways:

- By serial number
- By `ModuleID`.

You can obtain the `ModuleID`'s and serial numbers of all your modules by running the `enquiry` command-line utility.

6.11.6.2. Electronic Serial Number (ESN)

The serial number is a unique 12-digit number that is permanently encoded into each module. Quote this number in any correspondence with Support.

6.11.6.2.1. ModuleID

The `ModuleID` is an integer assigned to the module by the server when it starts. The first module it finds is given a `ModuleID` of 1, the next is given a `ModuleID` of 2, and this pattern of assigning `ModuleID` numbers continues for additional modules.

The order in which buses are searched and the order of modules on a bus depends on the exact configuration of the host. If you add or remove a module, this can change the allocation of ModuleIDs to all the modules on your system.

You can use the `enquiry` command-line utility to identify the PCI bus and slot number associated with a module.

All commands sent to nShield modules require a `ModuleID`. Many Security World Software commands, including all acceleration-only commands, can be called with a `ModuleID` of 0. Such a call causes the hardserver to send the command to the first available module. If you purchased a developer kit, you can refer to the developer documentation for information about the commands that are available on nShield modules.

In general, the hardserver determines which modules can perform a given command. If no module contains all the objects that are referred to in a given command, the server returns an error status.

However, some key-management operations must be performed together on the same module. In such cases, your application must specify the `ModuleID`.

To be able to share OCSs and keys between modules, the modules must be in the same Security World.

6.11.6.3. Adding a module

If you have a module installed, you can add further modules without reinstalling the server software.

However, we recommend that you always upgrade to the latest server software and upgrade the firmware in existing modules to the latest firmware.

1. Install the module hardware. Refer to the *Installation Guide* for information on installing nShield hardware.
2. (**Linux**) Run the script `/opt/nfast/sbin/install`.
3. Add the module to the Security World. Refer to Adding or restoring an HSM to the Security World.

6.11.6.4. Module fail-over

The Security World Software supports fail-over: if a module fails, its processing can be transferred automatically to another module provided the necessary keys have been loaded. Depending on the mode of failure, however, the underlying bus and operating system may not be able to recover and continue operating with the remaining devices.

To maximize uptime, we recommend that you fit any additional nShield modules for failover on a bus that is physically separate from that of the primary modules.

6.12. Stopping and restarting the hardserver

If necessary, you can stop the hardserver on the client, and where applicable the Remote Administration Service, by running the following command. On Windows, this must be a command window with administrative privileges.

Linux

```
/opt/nfast/sbin/init.d-ncipher stop
```

Windows

```
net stop "nfast server"
```

If the Remote Administration Service is running, you will be warned and given the option of continuing or not.

Similarly, you can start the hardserver on the client, and where applicable the Remote Administration Service, by running the following command. On Windows, this must be a command window with administrative privileges.

Linux

```
/opt/nfast/sbin/init.d-ncipher start
```

You can also restart the hardserver on the client, and where applicable the Remote Administration Service, by running the following command:

```
/opt/nfast/sbin/init.d-ncipher restart
```

Windows

```
net start "nfast server"
net start "nfast Remote Administration Service"
```

On Windows, you can also stop, start, or restart the hardserver, and where applicable the

Remote Administration Service, from the Windows Control Panel:

1. From the Windows Start menu, open the Windows Control Panel.

2. Double-click Administrative Tools.

3. Double-click Services.

4. Locate **nFast Server** or **nFast Remote Administration Service** in the list of services, and from the **Action** menu, select **Stop**, **Start**, or **Restart** as required.

> ℹ️ The **nFast Remote Administration Service**, where applicable, is dependent on the **nFast Server** so should be started or restarted after the **nFast Server**.

6.13. Resetting and testing the nShield HSM

6.13.1. Default configuration

To reset the unit to its default configuration, select **System > System configuration > Default config** and confirm that you want to set the default configuration.

This removes the configuration of the module but does not change its K_{NETI}.

6.13.2. Factory state

To reset the unit to its original (factory) state, select **Factory state** from the main menu and confirm that you want to return the unit to its factory state.

This gives a new K_{NETI} to the unit, which means that you must update the keyhash field of the unit's entry in the `nethsm_imports` section of the configuration file of all the clients that use it.

After a factory reset, ensure you re-enable any dynamic features. See Remotely enabling dynamic feature certificates including nShield HSM client licenses.

For more information about:

- The contents of the configuration files, see Module and client configuration files
- Configuring a new remote file system for the unit, see Configuring the Remote File System (RFS).

6.13.3. Testing the installation

To test the installation and configuration, follow these steps:

1. sign in to the client computer as a regular user, and open a command window.

2. Run the command:

   Linux

   ```
   opt/nfast/bin/enquiry
   ```

   Windows

   ```
   enquiry
   ```

   A successful `enquiry` command returns an output of the following form:

   ```
   Server:
    enquiry reply flags  none
    enquiry reply level  Six
    serial number        ####-####-####
    mode                 operational
    version              #-#-#
    speed index          ######
    rec. queue           ####..####
    ---
    version serial       #
    remote port (PV4)    ####

   Module ##:
    enquiry reply flags  none
    enquiry reply level  Six
    serial number        ####-####-####
    mode                 operational
    version              #-#-#
    speed index          #####
    rec. queue           ##..###
    ---
    rec. LongJobs queue  ##
    SEE machine type     PowerPCELF
    supported KML types  DSAp1024s160 DSAp3072s256
    hardware status      OK
   ```

   If the mode is `operational`, the unit is installed correctly.

   If the `enquiry` command says that the unit is not found:

   a. Restart the client computer.

   b. Re-run the `enquiry` command.

6.14. Configure the hardserver to export the module for guest VM usage

ⓘ | This feature is not available on nToken models, but works with all other

> local HSM models. If you previously configured this feature using rserverperm, you may wish to update to using these instructions using the config file to specify the permissions for guest VMs to access the HSMs in a persistent manner.

1. Configure the host hardserver to permit guest VM hardservers to share access to the module:

   a. Edit the host hardserver config file `NFAST_KMDATA/config/config` (**Linux**) or `NFAST_KMDATA\config\config` (**Windows**).

   b. Add a new entry in the `hs_clients` section to contain the details of the client to be added.

   > **ⓘ** If your config file does not already contain a `hs_clients` section you may add it yourself with a line containing only `[hs_-clients]`.

   The `addr` and `clientperm` fields are required for each client, and `keyhash` is recommended for authentication: :

   ```
   [hs_clients]
   addr=<client_IP>
   clientperm=permission_type
   keyhash=software_keyhash
   ```

   Where:

   `<client_IP>` can be either the IP address of the guest VM or any of `0.0.0.0`, `::`, or blank if the host hardserver is to accept clients identified by their key hash instead of their IP address.

   If you set both the `<client_IP>` field (the guest VM's IP address) and the key hash, client connections will be restricted based on both values.

   `permission_type` defines the type of commands the client can issue (`unpriv` for unprivileged only, `priv` for privileged or `priv_lowport` for privileged connections restricted to low port numbers).

   `software_keyhash` is the hash of the software-generated authentication key that the client should authenticate itself with.

   If there is more than one client being configured, the fields for each client must be separated by line consisting of one or more hyphens (e.g. `----`).

   > **ⓘ** It is recommended that the firewall on the host be configured

> so that only connections from intended network interfaces can
> be made to the host hardserver on its Impath port (port 9004
> by default).

   c. Load the updated configuration file in the host hardserver. To do this, run the following command:

```
hsc_nethsmexports
```

> ℹ️ This command only needs to be run when the config is added
> or modified. The permissions for guest VMs will be re-applied
> automatically when the host hardserver is restarted.

2. Configure the hardserver in the guest VM to enroll to the host hardserver with an IP address using the virtual switch. Enter the following command for each guest hardserver that should have unprivileged access:

```
nethsmenroll <host-hardserver-ip>
```

Run the following command if the guest hardserver should have privileged access for mode change and administration:

> ℹ️ Not all administration operations will be permitted from a privileged
> guest VM, such as firmware updates, which must be carried out
> from the host.

```
nethsmenroll -p <host-hardserver-ip>
```

You will be asked to confirm your entries. You should then see the following message:

```
OK configuring hardserver's nethsm imports
```

3. Confirm the connection from the guest VMs by running `enquiry`.

7. Enabling optional features

nShield HSMs support a range of optional features. Optional features must be enabled with a certificate that is supplied by Entrust. You can order features when you purchase a unit, or you can obtain them at a later date (from your Entrust account manager). Feature certificates are supplied as a file made available for download or requested as a smart (Activator) card, to be delivered by post. Features are enabled using the front panel of the nShield 5c or by using either the `nethsmadmin` or the Feature Enable Tool (FET) utility from a privileged client.

Features provide additional functionality that must be enabled using the certificate file before the HSM can perform certain actions and use particular mechanisms. Features are either static or dynamic. Static features are persistent and remain enabled even if the HSM is factory stated or upgraded, most features are static. Conversely, dynamic features (i.e. client licenses, for adding further clients IP addresses to the nShield 5c above the default three and SEE restricted), are non persistent and, if the HSM is factory stated, must be enabled again using the features file or activator card.

> **ℹ** Dynamic features are identified in Available optional features. If a feature is not identified as dynamic it is a static feature.

For more information about:

- Ordering optional features, see Ordering additional features
- Feature-enabling procedures, see Enabling features.

The HSM checks to confirm whether any features that it attempts to use are enabled. It normally does this when it authorizes the commands or command options that relate to a specific feature.

> **ℹ** If you are enabling the Remote Operator feature, you must enable it on the HSM that is to be used as the unattended HSM.

For information about Remote Operator, see Remote Operator.

7.1. Available optional features

This section lists the features that can be added to the HSM. For details of all available features, contact Sales.

7.1.1. Elliptic Curve

Cryptography based on elliptic curves relies on the mathematics of random elliptic curve elements. It offers better performance for an equivalent key length than either RSA or Diffie-Hellman public key systems. Using RSA or Diffie-Hellman to protect 128-bit AES keys requires a key of at least 3072 bits. The equivalent key size for elliptic curves is only 256 bits. Using a smaller key reduces storage and transmission requirements.

Elliptic curve cryptography is endorsed by the US National Security Agency and NIST (the National Institute of Standards and Technology), and by standardization bodies including ANSI, IEEE and ISO.

nShield modules incorporate hardware that supports elliptic curve operations for ECDH (Elliptic curve Diffie-Hellman) and ECDSA (Elliptic Curve Digital Signature Algorithm) keys.

7.1.2. Elliptic Curve activation

All nShield HSMs require specific activation to utilize the elliptic curve features. HSMs use an activator smart card to enable this feature. Refer to Enabling features with a smart card for instructions on how to enable the EC feature. Additionally it is possible to activate the elliptic curve feature without a physical smart card. In this case the certificate details can be provided by email and entered locally. Refer to Enabling features without a smart card

Contact Sales if you require an EC activation.

nShield modules with elliptic curve activation support *MQV* (*Menezes-Qu-Vanstone*) modes.

7.1.3. Elliptic Curve support on the nShield product line

The following table details the range of nShield HSMs and the level of elliptic curve support that they offer.

HSM module type	Elliptic Curve support		Elliptic Curve offload acceleration[3]	
	Named curves[2]	Custom curves[1, 5]	Named curves[2]	Custom curves[1, 5]
nShield Edge (Windows only)	Yes	Yes	No	No
nShield Solo 500 and 6000 nShield 500, 1500, and 6000	Yes	Yes	No	No
nShield Solo 500+, 6000+ nShield 6000+	Yes	Yes	Yes, Prime curves and twisted Brain pool curves are accelerated[4].	Yes

HSM module type	Elliptic Curve support		Elliptic Curve offload acceleration[3]	
nShield Solo XC	Yes	Yes	Yes, Prime curves and both twisted and non-twisted Brainpool curves are accelerated[4].	Yes
nShield 5s	Yes	Yes	Yes, Prime curves and both twisted and non-twisted Brainpool curves are accelerated.	Yes

[1]Accessed via nCore, PKCS#11 and JCE APIs.

[2]Both Prime and Binary named curves are supported. Refer to Named Curves, below, which lists the most commonly supported elliptic curves.

[3]Offload acceleration refers to offloading the elliptic curve operation from the main CPU for dedicated EC hardware acceleration.

[4]Binary curves are supported, but are not hardware offload accelerated.

[5]Brainpool curves are supported as named curves via nCore, PKCS#11 and JCE only.

7.1.4. nShield software / API support required to use elliptic curve functions

	Security World Software for nShield	CodeSafe
Elliptic curve supported / API	Microsoft CNG, PKCS#11, Java Cryptographic Engine (JCE)[1].	Microsoft CNG, PKCS#11, Java Cryptographic Engine (JCE)[1].

[1]Java elliptic curve functionality is fully supported by the nShield security provider, nCipherKM. There is also the option to use the Sun/IBM PKCS #11 Provider with nCipherKM configured to use the nShield PKCS#11 library.

7.1.5. Named Curves

This table lists the supported named curves that are pre-coded in nShield module firmware.

Supported named curves			
ANSIB163v1	BrainpoolP160r1	NISTP192	SECP160r1
ANSIB191v1	BrainpoolP160t1	NISTP224	SECP256k1
	BrainpoolP192r1	NISTP256	

Supported named curves			
	BrainpoolP192t1	NISTP384	
	BrainpoolP224r1	NISTP521	
	BrainpoolP224t1	NISTB163	
	BrainpoolP256r1	NISTB233	
	BrainpoolP256t1	NISTB283	
	BrainpoolP320r1	NISTB409	
	BrainpoolP320t1	NISTB571	
	BrainpoolP384r1	NISTK163	
	BrainpoolP384t1	NISTK233	
	BrainpoolP512r1	NISTK283	
	BrainpoolP512t1	NISTK409	
		NISTK571	

7.1.6. Custom curves

nShield modules also allow the entry of custom elliptic curves which are not pre-coded in firmware. If the curve is Prime, it may benefit from hardware acceleration if supported by the nShield HSM (see nShield software / API support required to use elliptic curve functions, above).

Custom curves are supported by nCore and PKCS #11 APIs.

7.1.7. Further information on using elliptic curves

For more information on how to use elliptic curves, see the following sections:

- PKCS #11:
  - Mechanisms supported by PKCS #11: Mechanisms
- CNG (**Windows**):
  - Supported algorithms for CNG: Supported algorithms for CNG
  - Key exchange for CNG: Key exchange
- Symmetric and asymmetric algorithms: Cryptographic algorithms
- Using `generatekey` options and parameters to generate ECDH and ECDSA keys: Key generation options and parameters

> Java elliptic curve functionality is fully supported by the nShield security provider, nCipherKM. There is also the option to use the Sun/IBM PKCS #11 Provider with nCipherKM configured to use the PKCS #11 library.

7.1.8. Secure Execution Engine (SEE)

The SEE is a unique secure execution environment. The SEE features available to you are:

SEE Activation, CodeSafe 5	This feature enables the ability to run signed SEE applications within your HSM. To develop your own SEE applications, you must also purchase the CodeSafe SDK and obtain a CodeSafe developer id certificate from Entrust.
	For more information about how to develop SEE applications, see the *CodeSafe 5 Developer Guide*.

7.1.9. Remote Operator support

Many Entrust customers keep critical servers in a physically secure and remote location. The Security World infrastructure, however, often requires the physical presence of an operator to perform tasks such as inserting cards. Remote Operator enables these customers to remotely manage servers running Security World Software using a secure nShield communications protocol over IP networks.

The Remote Operator feature must be enabled on the module installed in the remote server. Remote Operator cannot be enabled remotely on an unattended module.

For more information about using Remote Operator, see Remote Operator.

For v12 and later, Entrust recommends that you use Remote Administration, which is more flexible than the Remote Operator functionality.

7.1.10. ISO smart card Support (ISS)

ISS, also called Foreign Token Open (FTO) allows data to be read to and written from ISO 7816 compliant smart cards in a manner prescribed by ISO7816-4. ISS allows you to develop and deploy a security system that can make full use of ISO 7816 compliant smart cards from any manufacturer.

7.1.11. Korean algorithms

This feature enables the following mechanisms:

- Korean Certificate-based Digital Signature Algorithm (KCDSA), which is a signature mechanism.

  KCDSA is used extensively in Korea as part of compliance with local regulations specified by the Korean government. For more information about the KCDSA, see the *nCore API Documentation*.

- SEED, which is a block cipher.
- ARIA, which is a block cipher.
- HAS160, which is a hash function.

7.1.12. Fast RNG for ECDSA

Utilise a faster alternative for Random Number Generation (RNG) for Elliptic Curve Digital Signature Algorithm (ECDSA). This feature is applicable only for nShield Solo XC, nShield Connect XC, nShield 5s, and nShield 5c.

The faster performance, comparable with v12.40 performance, is achieved by the RNG part of ECDSA being done on the NXP C291 Crypto Coprocessor.

This implementation of ECDSA uses an RNG that is not within scope for the nShield HSM certifications and for this reason it will not be used when the HSM is in a fips-140-level-3 or common-criteria-cmts Security World (regardless of the feature bit setting).

7.1.12.1. Client licenses

You can purchase additional client licenses that allow you to run multiple clients for the unit. Three clients are always enabled on each unit.

7.2. Ordering additional features

When you have decided that you require a new feature, you can order it from Sales. Before you call Sales, collect information about your HSM as follows:

- If possible, make a note of the unit serial number. This can be found on the base of the unit.
- Make a note of the Electronic Serial Number of the unit. You can find this from the front panel menu, by going to **HSM > HSM information > Display details**.

You must provide the ESN number to order a new feature.

If you prefer, you can include this information in an e-mail to Sales.

When your order has been processed, you receive a Feature Enabling Certificate in one of the following ways:

- Entrust e-mails you the Feature Enabling Certificate.
- Entrust sends you a smart card that contains the Feature Enabling Certificate.

The Feature Enabling Certificate contains the information that you need to enable the features you have ordered.

For more information, including pricing of features, telephone or email your nearest Sales representative using the contact details from this guide, or contact Entrust nShield Support, https://nshieldsupport.entrust.com.

7.3. Enabling features

Feature enabling differs for static and dynamic features.

- You can enable static features from the front panel of the unit or from the client.
- Entrust recommends that you enable dynamic features from the client. If the dynamic feature applies directly to nShield HSM, for example client licenses, you can use a `nethsmadmin` option to apply them. See Remotely enabling dynamic feature certificates including nShield HSM client licenses

> ℹ️ When enabling static feature(s) from the front panel, either using a card or a file, the module is cleared without warning. This will cause the HSM to drop or restart any SEE machine, and lose all the application keys that were loaded. In some cases, applications may need to be restarted.

7.3.1. Viewing enabled features

To see which (if any) features have already been enabled on the nShield HSM, from the main menu select **HSM > HSM feature enable > View current state**.

To print this list to a file on the unit, select **HSM > HSM feature enable > Write state to file**. The resulting file is transferred when the unit configuration is pushed to the remote file system. You can find it in `/opt/nfast/kmdata/hsm-ESN/features/fet.txt` (**Linux**) or `%NFAST_KM-DATA%\hsm-ESN\features\fet.txt`.

7.3.2. Enabling features with a smart card

To enable a new feature with a Feature Enabling smart card from Entrust:

1. Insert the Feature Enabling smart card into the unit's slot.
2. From the front panel, select **HSM > HSM feature enable > Read FEM from card**.

A message is displayed if the features are enabled successfully. If you do not see this message confirming a successful upgrade, see Enabling features without a smart card.

7.3.3. Enabling features without a smart card

You can also provide the Feature Enabling Certificate information supplied by Entrust from a file.

To enable a feature without a smart card:

1. Put the file that contains the feature enabling certificate in `/opt/nfast/kmdata/hsm-ESN/features` (**Linux**) or `%NFAST_KMDATA%\hsm-ESN\features` (**Windows**) on the remote file system. In this path, `ESN` is the ESN of the module.

   You can give the file any name that you wish. You must enter the file name on the unit's front panel, so you may prefer to use a short name.

2. From the front panel, select **HSM > HSM feature enable > Read from a file**.
3. Specify the name of the file that contains the certificate.

If the feature is enabled successfully, a message confirms this.

7.4. Remotely enabling dynamic feature certificates including nShield HSM client licenses

Feature certificates contained on the remote file system (RFS) can be applied to the nShield HSM. The main use case for applying feature certificates is for enabling the client licenses dynamic feature which have been purchased after the initial nShield HSM purchase, although both static and other dynamic feature certificates can be applied.

> 🛈 If you have performed a factory reset of your HSM, ensure you re-enable any dynamic features.

To apply a dynamic feature certificate, such as an nShield HSM client license, do the following:

> ℹ️ Feature certificates must be present on the RFS in the folder `$NFAST_KM DATA/hsm-ESN/features`.

1. Use the `nethsmadmin` utility to list the nShield HSM feature files on the RFS. Run the command:

```
nethsmadmin --module=<MODULE> --rfs=<RFS_IP> --list-features
```

In this command:

- `<MODULE>` specifies the HSM to use, by its ModuleID (default = 1).
- `<RFS_IP>` specifies the IP address of the RFS.
- Additionally the `--rfs-hkneti=<RFS_HKNETI>` and `--rfs-esn=<RFS_ESN>` options can be set to enable secure authentication of the RFS. There are three possible cases:

  - **Without secure authentication**: The authentication of the RFS will be based on the IP address only if the `--rfs-hkneti` and `--rfs-esn` options are not specified.

  - **Software-based authentication**: The `--rfs-hkneti` option specifies the software KNETI hash of the RFS. The `--rfs-esn` option shall not be specified.

    `<RFS_HKNETI>` can be obtained by running `anonkneti -m0 localhost` on the RFS.

  - **nToken authentication**: Only if an nToken (or local HSM) is installed in the RFS. The `--rfs-hkneti` and `--rfs-esn` options specify the KNETI hash and ESN of the nToken.

    `<RFS_HKNETI>` and `<RFS_ESN>` can be obtained by running `ntokenenroll -H` on the RFS.

2. Use the `nethsmadmin` utility to make the nShield HSM use a specific feature file from the RFS. Run the command:

```
nethsmadmin --module=<MODULE> --rfs=<RFS_IP> --apply-feature=<feature_file>
```

In this command:

- `<feature_file>` must be the path to the feature file that is displayed when you run the `nethsmadmin` command with the `--list-features` option. Errors are reported if you use either just the feature name, or the full path. The file must be alphanumeric, and no longer than 150 characters.

  The following is an example of the output expected when applying a dynamic fea-

ture:

```
Applying feature <DYNAMIC_FEATURE> to module <MODULE_NO> ...
Feature <DYNAMIC_FEATURE> application process started on module <MODULE_NO>
*DYNAMIC_FEATURE DETECTED*
Please restart you clientside hardserver and check the enquiry output to ensure the dynamic feature
has been applied correctly!
For the client licences feature check the 'max exported modules' section in enquiry to see if the new
client number has been applied correctly.
```

The following is an example of the output expected when applying a static feature:

```
Applying feature <STATIC_FEATURE> to module <MODULE_NO> ...
Feature <STATIC_FEATURE> application process started on module <MODULE_NO>
*STATIC_FEATURE DETECTED*
To be able to use the static feature please clear module MODULE_NO.
Use the fet utility to verify the feature was applied correctly.
```

In the output examples:

- `<DYNAMIC_FEATURE>` specifies the name of the dynamic feature file applied.
- `<STATIC_FEATURE>` specifies the name of the static feature file applied.
- `<MODULE_NO>` specifies the HSM that the feature was applied to.

8. Creating and managing a Security World

This chapter describes how to create and manage a Security World. You must create a Security World before using the HSM to manage keys.

You normally create a Security World after installing and configuring the module and its software. For more information, see:

- The Installation Guide for more about installing the module and software.
- Client Software and module configuration

You create a Security World with a single HSM. If you have more than one module, select one module with which to create the Security World, then add additional modules to the Security World after its creation. For more information, see Adding or restoring an HSM to the Security World. If you create a Security World with the audit logging feature enabled, all additional HSMs added to this Security World will also have audit logging enabled.

> 🛈 To use the module to protect a different set of keys, you can replace an existing Security World with a new Security World.

For more information about the type of user that is required for different operations, see About user privileges.

> ❗ All Security Worlds rely on you using the security features of your operating system to control the users who can access the Security World and, for example, write data to the host.

8.1. Creating a Security World

You can use the following to create Security Worlds:

- The unit front panel controls
  See Creating a Security World using the nShield HSM front panel.
- The `new-world` command line utility
  See Creating a Security World using new-world.

8.1.1. The creation process

When you create a Security World:

- The HSM is erased

- A new HSM key for this Security World is generated

- A new ACS to protect this HSM key is created

- The Security World information is stored within the file system of the nShield HSM operating system and on the RFS

  - The information is encrypted using the secrets stored on the ACS

- The HSM and Security World are configured for Audit Logging if selected

> ❗ If you want to re-use the physical cards created in a previous Security World, you must erase all Operator Cards, except for nShield Remote Administration Cards, while the previous Security World still exists. See Erasing cards and softcards.

> ❗ We recommend that you regularly back up the entire contents of the RFS. Either the `%NFAST_KMDATA%` directory on Windows, or the `kmdata` directory on Linux, is required to restore an nShield HSM or its replacement, to the current state in case of failure.

> ℹ️ Due to the additional primality checking required by SP800-131A, Security World generation will take longer when using the new default Ciphersuite (from v12.40 onwards) - on nShield USB-attached HSMs, this could be up to 45 minutes.

8.1.2. Security World Files

The Security World infrastructure stores encrypted key material and related data in files on the remote file system on the client. For multiple clients to use the same Security World, the system administrator must ensure that these files are copied to all the clients and updated when required.

For more information about the remote file system, see:

- Remote file system (RFS)
- Configuring the remote file system (RFS)

Other nShield HSMs can also use a Security World created on an nShield HSM using client cooperation. For more information, see Setting up client cooperation.

8.1.2.1. Location of Security World files

The logic for finding the security world data directory is:

1. If `NFAST_KMLOCAL` is set, use that.

2. Otherwise, if `NFAST_KMDATA` is set, use `${NFAST_KMDATA}/local` on Linux, `%NFAST_KM-DATA%\local` on Windows.

3. Otherwise, if `NFAST_HOME` is set, use `${NFAST_HOME}/kmdata/local` on Linux, `%NFAST_HOME%\kmdata\local` on Windows.

4. Otherwise, use `/opt/nfast/kmdata/local` on Linux, `C:\nfast\kmdata\local` on Windows.

If you want to make cards or keys which are normally created from the client available from the module's front panel, we recommend that you use client co-operation to automate the copying of files to the module. For information about configuring client co-operation, see Setting up client cooperation.

If you do not use client cooperation, you must manually copy the appropriate card and key files from the client or host on which the card set or key was created to the remote file system of `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**). These files must then be updated on the module by selecting **Security World mgmt** > **RFS operations** > **Update World files** from the main menu.

To be able to create Operator Cards or keys, the user on the client must have write permission for this directory. All other valid users must have read permission.

> 🛈 By default, the Key Management Data directory, and sub-directories, inherit permissions from the user that creates them. Installation of the Security World Software must be performed by a user with Administrator rights that allow read and write operations, and the starting and stopping of applications.

Security World operations create or modify Security World files as follows:

Operation	creates/modifies	file(s)
Create a Security World	creates	`world` (for each module in the Security World) `module_ESN`
Load a Security World	creates or modifies	(for each module in the Security World) `module_ESN`
Replace an ACS	modifies	`world`

Create an OCS	creates	`card_HASH`
		
		`cards_HASH_NUMBER`

Create a softcard	creates	`softcard_HASH`

Operation	creates/modifies	file(s)
Generate a key	creates	`key_APPNAME__IDENT`
Recover a key	modifies	`key_APPNAME` (for each key that has been recovered)

- `<ESN>` - Electronic serial number of the module on which the Security World is created.
- `<IDENT>` - Identifier given to the card set or key when it is created.
- `<NUMBER>` - Number of the card in the card set.
- `<APPNAME>` - Name of the application by which the key was created. It's a 40-character string that represents the hash of the card set's logical token. It's either user supplied or a hash of the key's logical token, depending on the application that created the key.

8.1.2.2. Required files

The following files must be present and up to date in the `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**) directory, or the directory specified by the `NFAST_KMLOCAL` environment variable, for a client to use a Security World:

- `world`
- A `module_ESN` file for each module that this host uses
- A `cards_<IDENT>` file for each card set that is to be loaded from this host
- A `card_<IDENT>_NUMBER` file for each card in each card set that is to be loaded from this host
- A `key_<APPNAME>_<IDENT>` file for each key that is to be loaded from this host.

These files are not updated automatically. You must ensure that they are synchronized whenever the Security World is updated on the module.

8.1.3. Security World options

Decide what kind of Security World you need before you create it. Depending on the kind of Security World you need, you can choose different options at the time of creation. For convenience, Security World options can be divided into the following groups:

- Basic options, which must be configured for all Security Worlds
  - Optionally enable Audit Logging for the Security World

- Recovery and replacement options, which must be configured if the Security World, keys, or passphrases are to be recoverable or replaceable
- SEE options, which only need be configured if you are using CodeSafe
- Options relating to the replacement of an existing Security World with a new Security World.

Security World options are highly configurable at the time of creation but, so that they will remain secure, not afterwards. For this reason, we recommend that you familiarize yourself with Security World options, especially those required by your particular situation, before you begin to create a Security World.

8.1.3.1. Security World basic options

When you create a Security World, you must always configure the basic options described in this section.

8.1.3.1.1. Cipher suite

Only one Cipher suite is supported and this is SP800-131 compliant.

8.1.3.1.2. ACS quorum

You must decide the total number of cards (N) in a Security World's ACS and must have that many blank cards available before you start to create the Security World. You must also decide how many cards from the ACS must be present (K) when performing administrative functions on the Security World.

> ℹ️ We recommend that you do not create ACSs for which K is equal to N, because you cannot replace such an ACS if even 1 card is lost or damaged.

> ℹ️ In Common Criteria CMTS Security Worlds the minimum value of K for the ACS is 2.

In many cases, it is desirable to make K greater than half the value of N (for example, if N is 7, to make K 4 or more). Such a policy makes it harder for a potential attacker to obtain enough cards to access the Security World. Choose values of K and N that are appropriate to your situation.

The total number of cards used in the ACS must be a value in the range $1 - 64$.

8.1.3.1.3. FIPS 140 Level 3 compliance

By default, Security Worlds are created to comply with the roles and services, key management, and self-test sections of the FIPS 140 standard at Level 2. However, you can choose to enable compliance with the FIPS 140 standard at Level 3.

> ℹ️ This option provides compliance with the roles and services of the FIPS 140- 2 Level 3 standard. It is included for those customers who have a regulatory requirement for compliance.

If you enable compliance with FIPS 140 Level 3 roles and services, authorization is required for the following actions:

- Generating a new OCS
- Generating or importing a key, including session keys
- Erasing or formatting smart cards (although you can obtain authorization from a card you are about to erase).

In addition, you cannot import or export private or symmetric keys in plain text.

8.1.3.1.4. UseStrongPrimes Security World setting

From firmware version 12.70, the nShield HSM always targets FIPS 186-4 compliance when generating RSA keys of 1024 bits or more. It typically does this using a "strong primes" strategy, however Entrust only guarantees this strategy if the `UseStrongPrimes` setting is enabled.

If your firmware is version 12.70 or higher, you do not need this setting enabled for FIPS 186-4 compliance.

If you are using an older version of firmware, meaning it has a version number *lower than* 12.70, then you need the `UseStrongPrimes` setting enabled to grant FIPS 186-2 compliance.

If your Security World is FIPS 140 Level 3, then this setting is on by default. If your Security World is not FIPS 140 Level 3, then you can disable the `UseStrongPrimes` setting for faster RSA key generation, however this removes FIPS 186-2 compliance.

8.1.3.1.5. Remote Operator

To use a module without needing physical access to present Operator Cards, you must enable the Remote Operator feature on the module. For more information, see Enabling optional features.

By default, modules are initialized into Security Worlds with remote card set reading

enabled. If you add a module for which remote card reading is disabled to a Security World for which remote card reading is enabled, the module remains disabled.

8.1.3.2. OCS and softcard replacement

By default, Security Worlds are created with the ability to replace one OCS or softcard with another. This feature enables you to transfer keys from the protection of the old OCS of softcard to a new OCS or softcard.

> You can replace an OCS with another OCS, or a softcard with another softcard, but you cannot replace an OCS with a softcard or a softcard with an OCS. Likewise, you can transfer keys from an OCS to another OCS, or from a softcard to another softcard, but you cannot transfer keys from an OCS to a softcard or from a softcard to an OCS.

You can choose to disable OCS and softcard replacement for a Security World when you create it. However, in a Security World without this feature, you can never replace lost or damaged OCSs; therefore, you could never recover the keys protected by lost or damaged OCSs, even if the keys themselves were generated as recoverable (which is the default for key generation).

> OCS and softcard replacement cannot be enabled after Security World creation without reinitializing the Security World and discarding all the existing keys within it.

For an overview of Security World robustness and OCS or softcard replacement, see Replacing an Operator Card Set or recovering keys to softcards. For details about performing OCS and softcard replacement operations, see Replacing Operator Card Sets and Replacing the Administrator Card Set.

8.1.3.3. passphrase replacement

By default, Security Worlds are created so that you cannot replace the passphrase of a card or softcard without knowing the existing passphrase.

However, you can choose to enable passphrase replacement at the time you create a Security World. This option makes it possible to replace the passphrase of a card or softcard even if you do not know the existing passphrase. Performing such an operation requires authorization from the Security World's ACS.

For details about performing passphrase replacement operations, see Changing unknown or lost passphrase.

8.1.3.4. Nonvolatile memory (NVRAM) options

Enabling nonvolatile memory (NVRAM) options allows keys to be stored in the module's NVRAM instead of in the Key Management Data directory of the host computer. Files stored in the module's non-volatile memory have Access Control Lists (ACLs) that control who can access the file and what changes can be made to the file. NVRAM options are relevant only if your module's firmware supports them, and you can store keys in your module's NVRAM only if there is sufficient space.

> ℹ️ When the amount of information to be stored in the NVRAM exceeds the available capacity, you can instead store this data in a blob encrypted with a much smaller key that is itself then stored in the NVRAM. This functionality allows the amount of secure storage to be limited only by the capacity of the host computer.

8.1.3.5. Security World SEE options

You must configure `SEE options` if you are using the nShield Secure Execution Engine (SEE). If you do not have SEE installed, the SEE options are irrelevant.

8.1.3.5.1. SEE debugging

SEE debugging is disabled by default, but you can choose whether to enable it for all users or whether to make it available only through use of an ACS. In many circumstances, it is useful to enable SEE debugging for all users in a development Security World but to disable SEE debugging in a production Security World. Choose the SEE debugging options that best suit your situation.

8.1.3.5.2. Real-time clock (RTC) options

Real-time clock (RTC) options are relevant only if you have purchased and installed the CodeSafe Developer kit. If so, by default, Security Worlds are created with access to RTC operations enabled. However, you can choose to control access to RTC operations by means of an ACS.

8.1.3.6. Security World replacement options

Options relating to Security World replacement are relevant only if you are replacing a Security World.

If you replace an existing Security World, its `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KM`

`DATA%\local` (**Windows**) directories are not overwritten but renamed to
`/opt/nfast/kmdata/local_N` (**Linux**) or `%NFAST_KMDATA%\local_N` (**Windows**), (where *N* is an
integer assigned depending on how many Security Worlds have been previously saved during overwrites). A new Key Management Data directory is created for the new Security
World. If you do not wish to retain the `/opt/nfast/kmdata/local_N` (**Linux**) or `%NFAST_KM-DATA%\local_N` (**Windows**) directory from the old Security World, you must delete it manually.

8.1.4. Creating a Security World using the nShield HSM front panel

When initiated from the nShield HSM front panel, while a Security World
is being created the nShield HSM disconnects itself from the network to
ensure that the operation is not interrupted. This means that the
Remote Administration feature cannot be used to present cards from a
remote location when creating a Security World from the front panel.

8.1.4.1. Before you start

Before you start to create a Security World:

- The `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**) directory
  must exist on the remote file system and be empty.

- Before configuring the Security World, you should know:

  - The security policy for the HSM

  - The number and quorum of Administrator Cards and Operator Cards to be used.

  To help you decide on the Security World you require, see Security World options.

- You must have enough smart cards to form the Security World's card sets.

To create a Security World from the nShield HSM Front Panel:

1. From the main menu, select **Security World mgmt** > **Module initialization** > **New Security World**.

2. Specify the Security World mode:

   a. **FIPS 140 Level 3** creates a Security World compliant with FIPS 140 requirements
      for roles and services at Level 3.

   b. **Common Criteria CMTS** creates a Security World supporting Common Criteria
      Protection Profile EN 419 221-5.

   c. **Unrestricted** creates a Security World which doesn't impose any particular confor

mance. With appropriate environmental constraints, an unrestricted Security World can be compliant with FIPS 140 Level 2.

3. Select the Cipher suite for the Security World. Currently only one option is available for the Security World key, AES (SP800-131AR1).

4. Enter the default quorum for the ACS. This consists of:

   a. The maximum number of cards from the ACS required by default for an operation. This number must be less than or equal to the total number of cards in the set.

   b. The total number of cards to be used in the ACS. This must be a value in the range 1 − 64 except for the Common Criteria CMTS Security World mode, for which the range is 2 − 64.

   > We recommend that you do not create an ACS for which the required number of cards is equal to the total number of cards because you cannot replace such an ACS if even a single card is lost or damaged.

5. If you answer the question **Specify all quorums?** by selecting:

   a. **no** - all operations and features (with the exception of passphrase recovery) will be enabled and require the maximum number of cards

   b. **yes** - you can specify which operations and features you want to enable (including passphrase recovery) and what the required number of cards for each of these will be.

6. If you chose to disable individual features or require a lower number of cards required for an operation, specify these parameters now. You can select a different number of Administrator Cards (K) to be required for each operation. You can also disable recovery and replacement operations and choose to use KNSO to authorize SEE (Secure Execution Engine) operations. The options for which you can specify a separate value of K are as follows:

Operation	Action allowed on HSM
Module reprogramming	Initializing an HSM into a Security World. You must specify a value of K for this operation.
passphrase replacement	Replacement of passphrases from backup files when recovering an OCS. You can disable this operation, see passphrase replacement. This operation is disabled in Common Criteria CMTS mode and cannot be enabled.
OCS/softcard replacement	Recovery of keys from backup files when replacing an OCS. You can disable this operation if you are using the nShield HSM, see OCS and softcard replacement.

Operation	Action allowed on HSM
NVRAM access	Reading from and writing to the NVRAM. You can choose to authorize this operation with K_{NSO}, see Nonvolatile memory (NVRAM) options.
RTC access	Updating the real time clock. This is not applicable for nShield 5c. You can choose to authorize this operation with K_{NSO}, see Real-time clock (RTC) options.
SEE debugging	Viewing full SEE debug information. You can specify a value of K for this operation, all it for all users or authorize it with K_{NSO}, see SEE debugging. This operation is disabled in Common Criteria CMTS mode.
FTO	Use of an Foreign Token Open (FTO) Delegate Key (ISO Smart Card Support). You can specify a value of K for this operation or authorize it with K_{NSO}. This operation is disabled in Common Criteria CMTS mode.

7. Specify if audit logging should be enabled.

> 🛈 In Common Criteria CMTS mode, audit logging is automatically enabled and cannot be disabled.

8. Specify whether the HSM is a valid target for remote shares (that is, whether it can import slots), see Remote Operator. This option is disabled for Common Criteria CMTS mode.

9. For Common Criteria CMTS mode only, choose whether to specify the maximum number of times an Assigned key can be used since it was authorized. A use limit compatible with the specified maximum will be imposed at key creation time and can be verified for Assigned keys. If you choose to specify a maximum key usage limit:

   a. Enter the key usages allowed, up to a maximum of 9999.

10. For Common Criteria CMTS mode only, choose whether to specify a maximum timeout for Assigned keys since key authorization. A time limit compatible with the specified maximum will be imposed when the key is created, and can be verified for Assigned keys. If you choose to specify a key timeout:

    a. Select the units from **Seconds**, **Minutes**, **Hours**, or **Days**.

    b. Enter a value up to a maximum of 9999 in your selected unit.

11. Format a card for the ACS as follows:

    a. Insert a card for the ACS and confirm that you want to use it.

    b. If the card is not blank, choose whether to overwrite it or to use a different card.

    c. Choose whether to specify a passphrase for the card. If you choose to specify a

passphrase:

    i. Enter the passphrase.

    ii. Enter the passphrase again to confirm it. If the two passphrases do not match, you must enter the correct passphrase twice.

  d. When prompted, remove the card.

12. Repeat the previous step to format additional cards for the ACS, setting their passphrases as described, until the ACS is complete. Each prompt screen shows how many cards are required and how many have been used.

13. At completion, a message confirms that the Security World has been created.

8.1.5. Creating a Security World using new-world

8.1.5.1. Before you start

Before you start to create a Security World:

- The HSM must be in pre-initialization mode. See Checking and changing the mode on the HSM for more about changing the mode.

- You must be logged in to the computer that is running the RFS. The RFS should be a privileged client that has the client tools installed. For more information, see server_settings.

- (**Windows**) You must have set the `NFAST_HOME` environment variable.

  > 🛈 | This variable is set by default during product software installation.

- (**Linux**) If you have installed the Security World Software in a directory other than `/opt/nfast/`, you must have created a symbolic link from `/opt/nfast/` to the directory in which the software is actually installed.

- Before configuring the Security World, you should know:

  - The security policy for the HSM

  - The number and quorum of Administrator Cards and Operator Cards to be used

  To help you decide on the Security World you require, see Security World options.

- You must have enough smart cards to form the Security World's card set.

When you have finished creating a Security World, you must change the mode to "Operational" using `nopclearfail -I m1` or `nopclearfail -O -m1`.

8.1.5.2. Using nethsmadmin to copy a Security World to an nShield HSM and check the current version

If a Security World is created using new-world, the `nethsmadmin` command-line utility enables you to copy the resultant files to a nShield HSM. Run the command:

```
nethsmadmin --module=<MODULE> --update-world
```

`nethsmadmin` can also be used to check if the Security World files have been copied to the nShield HSM. Run the command:

```
nethsmadmin --module=<MODULE> --check-world
```

In these commands:

`--module=<MODULE>`

Specifies the HSM to use, by its `ModuleID` (default = 1).

Follow the directions in this section to create a Security World from the command line with the `new-world` utility.

8.1.5.3. Running the new-world command-line utility

Open a command prompt window and type the command `new-world` using the options in the table.

The example below will create a Security World supporting FIPS140 Level 3 with a ACS quorum of 3/5 and with audit logging enabled.

```
new-world --mode=fips-140-level-3 --acs-quorum=3/5 --audit-logging
```

In this command:

Option	Description
`--initialize`	This option creates a new Security World, replacing any existing `/opt/nfast/kmdata/local/` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**) directory.
	ℹ️ Replacing an existing Security World in this way does not delete the Security World's host data and recovery and replacement data, but renames the existing `/opt/nfast/kmdata/local/` (**Linux**) or `%NFAST_KM-DATA%\local` (**Windows**) directory in which these reside as `%NFAST_KMDATA%\localN` (**Linux**) or `/opt/nfast/kmdata/localN` (**Windows**) (where *N* is an integer assigned depending on how many Security Worlds have been previously saved during overwrites).
`--factory`	This option erases an HSM, restoring it to factory state.
`--no-remoteshare-cert`	This option prevents making the HSM from becoming a target for remote shares.
`--no-strict-rsa-keygen`	If you have not specified a mode parameter you can use the `-no-strict-rsa-keygen` flag to disable the UseStrongPrimes setting. Otherwise it will be enabled by default. See UseStrongPrimes Security World setting.
`--mode=MODE`	`FIPS-140-level-3` creates a Security World compliant with FIPS 140 Level 3. `common-criteria-cmts` creates a Security World supporting *Common Criteria PP 419 221-5*. Omitting this option will create a default Security World compliant with FIPS 140 Level 2.

Option	Description
`--no-recovery`	This option disables the ability to recovery or replace OCSs and softcard (which is otherwise enabled by default). This is equivalent to setting `!r`, where the `!` operator instructs the system to turn off the specified feature (`r`).
	By default, `new-world` creates key recovery and replacement data that is protected by the cryptographic keys on the ACS. This option does not give Entrust or any other third party access to your keys. Keys can only be recovered if authorization from the ACS is available. We recommend that you leave OCS and softcard recovery and replacement functionality enabled.
	❗ We recommend that you do not disable the recovery and replacement option.
	❗ If you set the `--no-recovery` option, you can never replace lost or damaged OCSs generated for that Security World. Therefore, you could never recover any keys protected by lost or damaged OCSs, even if the keys themselves were generated as recoverable (which is the default for key generation).
	❗ OCS and softcard replacement cannot be enabled after Security World creation without reinitializing the Security World and discarding all the existing keys within it.
`--cipher-suite=<CIPHER -SUITE>`	This option specifies the Cipher suite and type of key that is used to protect the new Security World. `<CIPHER-SUITE>` should be set to `DLf3072s256mAESc-SP800131Ar1`.
`--nso-timeout=<TIMEOUT>`	This option allows you to specify the time-out (`<TIMEOUT>`) for new Security Worlds. By default, an integer given for *TIMEOUT* is interpreted in seconds, but you can supply values for *TIMEOUT* in the form *N* `s`, *N* `h`, or *N* `d` where *N* is an integer and `s` specifies second, `h` specifies hours, and `d` specifies days.
`--module=<MODULE>`	This option specifies the module to use (by its `ModuleID`). If you have multiple modules, `new-world` initializes them all together.

Option	Description
`--acs-quorum=<K>/<N>`	In this option, `<K>` specifies the minimum number of smart cards needed from the ACS to authorize a feature. You can specify lower *K* values for a particular feature. All the *K* values must be less than or equal to the total number of cards in the set. If a value for *K* is not specified, `new-world` creates an ACS that requires a single card for authorization.
	When the Security World is created in `common-criteria-cmts` mode, `new-world` requires a minimum K of 2.
	Some applications do not have mechanisms for requesting that cards be inserted. Therefore any OCSs that you create for use with these applications must have *K*=1.
	`<N>` specifies the total number of smart cards to be used in the ACS. This must be a value in the range 1 – 64. If a value for this option is not specified, `new-world` creates an ACS that contains a single card.
	We recommend that you do not create an ACS for which the required number of cards is equal to the total number of cards because you will not be able to replace the ACS if even a single card is lost or damaged.
	This option only takes effect if you are creating a new Security World.
`--reduced-features`	This option instructs `new-world` to use a reduced default feature set when creating a Security World. A Security World created with the `--reduced-features` option has no passphrase recovery; no NVRAM, RTC, or FTO; and no NSO dele gate keys. However, such a reduced-features Security World can perform many operations faster than more fully featured Security Worlds.
`--disablepkcs1pad`	This option disables the use of PKCS#1 v1.5 padding. All attempts to use PKCS#1 v1.5 padding for encryption or decryption operations will be rejected. PKCS#1 v1.5 signature operations are not affected. PSS and OAEP are not affected.

Option	Description
`--pp-min=LENGTH`	This option enables a minimum passphrase length check for the Administrator Card Set (ACS) the Operator Card Set (OCS) and any associated softcards when you create a Security World. The minimum passphrase length check is then applied after the Security World is created. When enabled and you attempt to create a card passphrase with fewer characters than the specified minimum length, the following warning message displays: ⚠️ | Warning: short passphrase. However, the passphrase can still be used. Example: ``` new-world --initialize --acs-quorum=K/N --pp-min=14 ``` If `--pp-min=<length>` is not used, the minimum passphrase length is set to the default value (0).
`--pp-strength`	This option enables passphrases to have at least one uppercase, lowercase, number, and symbol. If the `--pp-strength` argument is omitted, the complexity requirements are not enforced.
`--audit-logging`	This option configures the Security World and the HSM on which it is being created for audit logging, creating a log signing key for each HSM. ℹ️ | The log destination must have already been set in the hardserver configuration file. See Audit Logging. Audit logging is automatically enabled when the Security World is created in `common-criteria-cmts` mode.
`--max-keyusage`	This option allows the administrator to specify a maximum reauthorization condition in terms of number of key usages since authorization for Assigned keys in common-criteria-cmts mode. A use limit compatible with the specified maximum will be applied at key creation time and can be verified for Assigned keys. If this is not set then no `--max-keyusage` limit is applied to Assigned keys on creation.
`--max-keytimeout`	This option allows the administrator to specify a maximum reauthorization condition in terms of a TIMEOUT since authorization for Assigned keys in common-criteria-cmts mode. By default, an integer given for TIMEOUT is interpreted in seconds, but you can supply values for TIMEOUT in the form *Ns*, *Nh*, or *Nd* where *N* is an integer and *s* specifies second, *h* specifies hours, and *d* specifies days. A use limit compatible with the specified maximum will be applied at key creation time and can be verified for Assigned keys. If this is not set then no limit is applied to Assigned keys on creation.

> The `--max-keyusage` and `--max-keytimeout` options are only available in common-criteria-cmts mode. They provide support for the Protection Profile requirement that reauthorization conditions are set by an administrator on creating an Assigned Key.

8.1.5.4. new-world command-line utility features

Features for the Security World can be specified using the command line.

Security world features are selected by *feature expressions*. A feature expression is a comma-separated list of *feature terms*. Each term consists of a feature name, optionally preceded by either a double dash `--`, an exclamation point `!`, or `no-` to turn off the feature, and optionally followed by an equals sign `=` and the quorum of cards from the ACS required to use the feature. The default quorum is taken from the `K` argument of the `--acs-quorum` option.

> **Linux-only**
>
> The `!` character is interpreted by some shells as history expansion and must be escaped with a backslash, `\!`. The dash may be interpreted as being the start of an command-line option unless you have used the `-f` option or specified an HSM without including the `-m` flag.

> If you set the `!fto` flag, that is, turn off FTO, you will not be able to use smart cards to import keys.

> To use extended debugging for the HSM, you must set the `dseeall` flag.

The following feature names are available:

Feature name	Description
m	This feature makes it possible to add new HSMs into the Security World. This feature cannot be disabled.
r	This feature enables OCS and softcard replacement; see Replacing Operator Card Sets.
p	This feature enables passphrase replacement; see passphrase replacement and Changing card and softcard passphrase.
nv	This feature specifies that ACS authorization is needed to enable nonvolatile memory (NVRAM) allocation.

Feature name	Description
rtc	This feature specifies that ACS authorization is needed to set the real-time clock (RTC); (see Real-time clock (RTC) options). This feature is not applicable for nShield 5c.
dsee	This feature specifies that that ACS authorization is needed to enable SEE World debugging.
dseeall	This feature enables SEE World debugging for all users.
fto	This feature specifies that ACS authorization is needed to enable foreign token operations (FTO).

The following features remain available for use on presentation of the standard ACS quorum, even if turned off using the ! operator:

- nvram
- rtc
- fto

Setting the quorum of one these features to 0 has the same effect as turning it off using the ! operator.

The passphrase replacement (p) and dseeall features are turned off by default; the other options are turned on by default.

> **ℹ** The nonvolatile memory and SEE world debugging options are relevant only if you are using the Secure Execution Engine. If you have bought the CodeSafe Developer Kit, refer to the *CodeSafe Developer Guide* for more information.

> **ℹ** To use extended debugging for the HSM, you must set the dseeall flag.

> **!** The dseeall option is designed for testing purposes only. Do not enable this feature on production Security Worlds as it may enable SEE applications to leak security information.

For example, the following features:

```
m=1, r, !p, nv=2, rtc=1
```

Create a Security World for which:

- A single card from the ACS is required to add a new HSM

- The default number is required to replace an OCS

- passphrase replacement is not enabled

- Two cards are required to allocate nonvolatile memory

- One card is required to set the real-time clock (applies to SEE only).

8.1.5.5. new-world command-line utility output

If `new-world` cannot interpret the command line, it displays its usage message and exits.

If you attempt to set a quorum for a feature that you have disabled or if you attempt to set a quorum too high, `new-world` displays an error and exits.

If the HSM is not in the pre-initialization mode, `new-world` advises you that you must put the HSM in this mode and waits until you have changed the HSM mode before continuing.

The HSM must be in pre-initialization mode. See Checking and changing the mode on the HSM for more about changing the mode.

> **ℹ** If the HSM is in the pre-initialization mode, `new-world` prompts you for smart cards and passphrases as required.

8.1.6. After you have created a Security World

Store the ACS in a safe place.

> **❗** If you lose more than *N* minus *K* of these Administrator Cards you cannot restore the Security World or lost Operator Cards. For example, if you have a 2/3 ACS and you lose more than one card, you cannot restore the Security World. If you have created an Administrator card set where *K* = *N*, then the loss of one card stops you from being able to restore the Security World.

To prevent this situation from occurring, replace lost or damaged cards from the ACS as soon as you discover the loss or damage. For more information, see Replacing the Administrator Card Set.

> **❗** The security of the keys that you create within this Security World is wholly dependent on the security of these smart cards.

The Security World data is stored on the HSM and on the RFS. For more information, see Security World Files.

The HSM can now be used to create Operator Cards and keys for the new Security World.

8.2. Displaying information about your Security World

To display information about the status of your Security World:

- Select **Security World mgmt** > **Display World info** from the main menu
- Run the `nfkminfo` command-line utility. See Displaying information about a Security World with nfkminfo.
- Run the `kmfile-dump` command-line utility. See Displaying information about a Security World with kmfile-dump.
- Run the `nethsmadmin` command-line utility. See Using nethsmadmin to copy a Security World to an nShield HSM and check the current version.

You can also use KeySafe to view a summarized description of the Security World.

8.2.1. Displaying information about a Security World with nfkminfo

To display information about a Security World from the command line, run the command:

```
nfkminfo -w|--world-info [-r|--repeat] [-p|--preload-client-id]
```

In this command, the `-w|--world-info` option specifies that you want to display general information about the Security World. This option is set by default, so you do not need to include it explicitly.

Optionally, the command can also include the following:

Option	Description
`-r|--repeat`	This option repeats the information displayed. There is a pause at the end of each set of information. The information is displayed again when you press **Enter**.
`-p|--preload-client-id`	This option displays the preloaded client ID value, if any.

To output a detailed list of Security World information, run `nfkminfo` with the `-w|--world -info` option (with or without the other options). For a description of the fields in this list, and more information about using `nfkminfo`, see *nfkminfo: information utility*.

The following table maps there flags visible on the front panel when you select **3 Security World mgmt > 3-1 Display World Info** to the flags in the output of `nfkminfo`.

Front panel	nfkminfo
admin	k-out-of-n
nCore flags	slotlistflags
NFKM flags	flags
Module slots	nflags
Initialized	Initialised
ForeignTokenOpen	FTO

8.2.2. Displaying information about a Security World with kmfile-dump

To display information about a World from the command line, run the command:

```
kmfile-dump [<worldfile>]
```

where `<worldfile>` is the file storing the World data, usually `/opt/nfast/kmdata/local/world` (**Linux**) or `%NFAST_KMDATA%\local\world` (**Windows**).

If no **WorldVersion** is received as a result of the command then the World is either version 1 or version 2.

If a **WorldVersion** of either '2' or '3' is received then the World is version 3.

8.3. Adding or restoring an HSM to the Security World

When you have created your Security World, you can add additional HSMs to it. You can restore HSMs that have previously been removed from the same Security World in the same way.

You can also restore an HSM to a Security World to continue using existing keys and Operator Cards:

- After you upgrade the firmware
- If you replace the HSM.

> The additional HSMs can be any nShield HSMs.

To add an HSM to a Security World, you must:

- Have installed the additional HSM hardware, as described in the *Installation Guide*.

- Have a copy of the Security World data on the HSM's remote file system in the Key Management Data directory.

- Possess a sufficient number of cards from the ACS and the appropriate passphrases.

Adding or restoring an HSM to a Security World:

- Erases the Security World data on the HSM's internal file system

- Reads the required number of cards (*K*) from the ACS so that it can re-create the secret

- Reads the Security World data from the remote file system

- Uses the secret from the ACS to decrypt the Security World key

- Stores the Security World key in the HSM's nonvolatile memory

- Configures the HSM for audit logging if the Security World was created with audit logging selected.

After adding an HSM to a Security World:

- You cannot access any keys that were protected by a previous Security World that con tained that HSM.

- You have to sync the module file to the clients by one of the following methods:

    ◦ Copy the files manually to the clients.

    ◦ Run `rfs-sync -update`.

    See HSM and client configuration files.

> **ⓘ** It is not possible to program an HSM into two separate Security Worlds simultaneously.

8.3.1. Adding an HSM to a Security World using the nShield HSM front panel

To add an HSM to a Security World:

1. If the HSM already belongs to a Security World, erase it from the Security World to which it belongs, as described in Erasing a module from a Security World.

2. From the main menu, select **Security World mgmt** > **Module initialization** > **Load Security World**.

3. Specify whether the HSM can use the Remote Operator feature import slots. For more information, see Remote Operator.

4. At the prompt, insert an Administrator Card, and enter its passphrase if required.

5. Continue to insert Administrator Cards when prompted until you have inserted the number required to authorize HSM reprogramming.

8.3.2. Adding an HSM to a Security World with the CSP or CNG wizard (Windows)

To add an HSM to an existing Security World:

1. Ensure the HSM is in initialization mode and run the wizard by double-clicking its short cut in the Windows Start menu: **Start > Entrust nShield Security World**.
2. Click the **Next** button.

   The wizard allows you to configure HSM Pool mode for CAPI/CNG.

3. Click the **Next** button.

   If the wizard finds an existing Security World, it prompts you to specify whether you want to use the existing Security World or create a new Security World.

   If the wizard displays any other windows:

   a. Cancel the operation.
   b. Check that you have correctly set the environment variable `NFAST_KMDATA`.
   c. Copy the `local` sub-directory from the Key Management Data of another computer in the same Security World or from a backup tape of this computer to the Key Management Data directory of this computer.
   d. Run the wizard again.
4. Ensure that the **Use the existing security world** option is selected, and click the **Next** button.

You can then proceed to add HSMs in the same manner that you add multiple HSMs when you create a Security World.

8.3.3. Adding an HSM to a Security World with new-world

1. Open a command window and type the command:

```
new-world [-l|--program] [-S|--no-remoteshare-cert] [-m|--module=<MODULE>]
```

   In this command:

   ◦ `-l|--program`

This option adds an HSM to an existing Security World (in the Key Management Data directory). If you have multiple HSMs available, you can use the `-m|--module=‘MODULE` option to specify an HSM. If you do not specify an HSM `new-world` adds all available HSMs to the Security World.

- `-S|--no-remoteshare-cert`

These options prevent the HSM from becoming a target for remote shares.

- `-m|--module=<MODULE>`

This option specifies the HSM to use (by its `ModuleID`). If you have multiple HSMs and do not specify an HSM, `new-world` adds all available HSMs to the existing Security World.

If `new-world` cannot find the key-management data, it displays the message:

```
new-world: no existing world to load.
```

If you intend to initialize the HSM into a new Security World, run `new-world` with the `-i` option.

If the HSM is not in the pre-initialization state, `new-world` displays an error and exits.

The HSM must be in pre-initialization mode. See Checking and changing the mode on the HSM for more about changing the mode.

If the HSM is in the pre-initialization state, `new-world` prompts you for cards from the Security World's ACS and to enter their passphrases as required.

2. After `new-world` has reprogrammed the HSM, restart the HSM in the operational state.

   The HSM must be in pre-initialization mode. See Checking and changing the mode on the HSM for more about changing the mode.

3. Store the ACS in a safe place.

   > ℹ️ If any error occurs (for example, if you do not enter the correct passphrases), the HSM is reset to the factory state. The HSM does not form part of the Security World unless you run `new-world` again.

8.4. Security World migration

The current version of Security World software enables you to create a Security World that fully complies with the NIST Recommendations for the Transitioning of Cryptographic Algo

rithms and Key Sizes (SP800-131Ar1) or alternatively Common Criteria PP 419 221-5 (common-criteria-cmts) depending on the options selected at World creation. This is called World version 3.

We recommend that where compliance with the specifications above is required, you create a new World and create new keys within that World. However, the software also includes a `migrate-world` command-line utility that you can use for migrating existing keys into the new World. This is provided as a convenience for customers who require compliance with the specifications, and who need to continue using existing keys.

In the case of a Common Criteria World the specification prohibits the importing of assigned keys. Only general keys can be imported into a common-criteria-cmts World.

> ℹ️ Throughout the following sections, the terms `Source World` refers to the World from which you want to migrate keys, and `Destination World` refers to the World to which you want to migrate keys.

> ℹ️ The utility requires the use of two modules. One module is referred to as the source module. The other module is referred to as the destination module.

8.4.1. Pre-requisites for migrating keys

In order to use the `migrate-world` utility the following will be needed:

- Two HSMs. These can be any of the currently supported HSM types and the two HSMs do not need to be of the same type.
- A quorum of ACS cards for the source World.
- A quorum of ACS cards for the destination World.
- Sufficient blank cards to create new OCS cards for any keys that are OCS protected.
- Remote mode switching must be enabled on both HSMs used for the migration. See `enable_remote_mode` in the server_settings section or the *Top-level menu* chapter of the HSM *Install Guide*.

8.4.2. Restrictions on migrating keys

The following restrictions apply to the use of `migrate-world`:

- The source module must be running firmware version 12.50 or later.
- The destination module must be running firmware version 12.50 or later.

- Only recoverable keys can be migrated. If your source keys are non-recoverable, you cannot use the migration utility to migrate keys.

- You can change some, but not all, Security World properties during migration:

Property	Up to 13.3	13.4 and later
Key protection method whether softcard or OCS is used	Fixed	Fixed
Key protection name softcard name or cardset name	Fixed	Editable
Quorum	Editable	Editable

  If the name or quorum is to be changed, you must create the softcard or OCS in the destination world before migration begins.

- Replacement cards should be of the same or newer generation than the cards that they replace.

- The source and destination modules must both have KLF2 warrants in the correct location.

  The migration process directly downloads the warrants from the module for the nShield 5s and nShield 5c HSMs. You do not need to take any action.

  For pre-nShield5 HSMs:

  ◦ If one or both of the modules have a KLF warrant, you must request an upgrade to a KLF2 warrant from nshield.support@entrust.com before you start the migration.

  ◦ For Solo + and Solo XC, the default location is `NFAST_KMDATA/warrants/` (**Linux**) or `NFAST_KMDATA\warrants\` (**Windows**).

  ◦ For Connect + and Connect XC, the default location is `NFAST_KMDATA/hsm-<ESN>/warrants/` (**Linux**) or `NFAST_KMDATA\hsm-<ESN>\warrants\` (**Windows**).

  See Warrant Management.

- The operator running the migrate-world utility must have the access rights to create a privileged connection to the hardserver.

- The migration tool must have exclusive use of the modules during migration. Do not use them for any other purpose during migration and if either module is an nShield network-attached HSM, do not enter anything via the front panel during migration.

> 🛈 If the destination World is `fips-140-level-3`, then some keys that were usable in the source World may not be usable in the destination World due to those algorithms or key lengths being restricted. The migration

tool might not be able to successfully migrate these keys so they should be removed from the source World before attempting the migration. Any keys of this type that do migrate successfully will be restricted at the point of use.

If the destination World is `fips-140-level-3` or `common-criteria-cmts` the migration tool will automatically remove ExportAsPlain from the ACL of any migrated key during the migration process.

If the destination world does not support audit logging the migration tool will automatically remove LogKeyUsage from the ACL of any migrated key during the migration process.

8.4.3. About the migration utility

You can run the migration utility in the following modes:

- **Plan mode**: Returns a list of steps for migration and the required card sets and passphrases but does not migrate any keys.
- **Perform mode**: Runs the plan mode prior to presenting the option to proceed and migrate keys according to the plan.

8.4.3.1. Usage and options

```
migrate-world [OPTIONS] --src-module=<source_module> --dst-module=<dest_module> --source=<source-kmdata-path>
--debug --dst-warrant=<dst-warrant-path> --src-warrant=<src-warrant-path [--plan | --perform] --key-logging
```

Option	Enables you to...
`-k <KEYS>|--keys-at-once=<KEYS>`	Migrate no more than this number of keys per ACS loading. This is useful to prevent ACS time-outs if you have a large number of keys to migrate. (0=unlimited, default=0). It is recommended to limit the number of keys to be migrated at any one time to no more than 100.
`-h|--help`	Obtain information about the options you can use with the utility.
`-c <CARDSETS>|--cardsets-at-once=<CARDSETS>`	Migrate keys protected by this number of card sets or softcards per ACS loading. This is useful to prevent ACS time-outs if you have a large number of different card sets or softcards to migrate. (0=unlimited, default=0).
`--version`	View the version number of the utility.
`--src-warrant=<src-warrantfile>`	Specify the location of the warrant file of the source module.

Option	Enables you to…
`--src-module=<MODULE>`	Specify which module ID to use as the source module.
`--source=<SOURCE>`	Specify the path to the folder that contains the source World data.
`--plan`	View the list of steps that will be carried out.
`--perform`	Migrate keys interactively.
`--dst-warrant=<dst-warrantfile>`	Specify the location of the warrant file of the destination module.
`--dst-module=<ModuleId>`	Specify which module ID to use as the destination module.
`--debug`	Outputs debug messages and stack traces in case of errors. It is recommended to use this only for testing as it will slow down operation and make card timeouts more likely to occur. A large volume of output is produced for each key that is migrated, so it is recommended to migrate a single key at a time when using this option.
`--key-logging`	This option will enable key usage logging on all migrated keys. If the destination World does not support audit logging the keys will still be migrated but LogKeyUsage logging will not be set in the ACL of the migrated keys.
`--src-prots=<list of source protections>`	Specify a comma-separated list of OCS or softcard names in the source security world. The keys will be migrated to the corresponding protections specified with `--dst-prots`.
`--dst-prots=<list of destination protections>`	Specify a comma-separated list of OCS or softcard names in the destination security world. These will be the target protections for the keys that are protected with methods specified with `--src-prots` in the source security world.
`--prots-config=<PATH>`	Specify a configuration file that lists the source and destination protection pairs for migration. The file must contain pairs of tab-separated protection names `src_prot dst_prot`, one pair per line.

🛈 Do not terminate path names in the command parameters with a backslash character. If this is not possible then either terminate with a double backslash or insert a blank space between the backslash and the terminating quotation mark.

8.4.4. Migrating keys

8.4.4.1. Preparing for migration

Before you begin:

- Install the latest version of the Security World Software from the installation media. See the *Installation Guide* for more information.

- Ensure that the warrant files for the source and destination modules are stored in their default locations. If the warrant files are not at the default location, the `--src-warrant` and `--dst-warrant` parameters need to be specified in the `migrate-world` command.

  - For Solo +, or Solo XC, the default location is `NFAST_KMDATA/warrants/` (**Linux**) or `NFAST_KMDATA\warrants\` (**Windows**).

  - For Connect +, Connect XC modules, the default location is `NFAST_KMDATA/hsm-<ESN>/warrants/` (**Linux**) or `NFAST_KMDATA\hsm-<ESN>\warrants\` (**Windows**).

  - For nShield 5s and nShield 5c, you do not need to specify warrant locations because they store their warrants within the module.

- Copy the source World data to a location defined by the `--source=<SOURCE>` parameter of the migration tool.

- If the destination World does not exist already, create a new destination World. For instructions, see Creating a Security World

> ℹ️ You must enable all your features on the destination module before migration. Otherwise, the migration will fail.

8.4.5. Migrating keys process

> ℹ️ To ensure the security of your keys, we recommend that the migration process is overseen by ACS-holding personnel and the end-to-end migration process is completed continuously, without any breaks in the process. This will also reduce the possibility of your ACS experiencing a time-out.

> ℹ️ If the destination World supports audit logging you can choose whether the migrated keys will have key usage logging enabled or not by use of the `--key-logging` command line switch. If you only wish key usage logging to be enabled on a subset of the keys then you must separate the source keys into two groups and run the migrate-world command separately for each group.

To migrate keys to the destination World:

1. Run the migration utility in the perform mode with the required options. For information about the usage and options you can use, see About the migration utility.

2. Ensure that the data for the destination World is in the standard location for World

data, derived from one of the following:

- Either the environment variable `NFAST_KMLOCAL` or `NFAST_KMDATA`.
- The default directory:
  - (**Linux**) `/opt/nfast/kmdata/local/`
  - (**Windows**) `C:\ProgramData\Key Management Data\local`, or `C:\Documents and Settings\All Users\Application Data\nCipher\Key Management Data\local`, as appropriate.

3. If the module is not configured to use the destination World, the utility prompts you to program the module and supply the ACS of the destination World.

4. The utility guides you through specific steps, prompting you to supply the required card sets and passphrases.

5. At the end of the migration both the source and destination modules are cleared. If you wish to use the modules then you must reload them with an appropriate Security World.

> ℹ️ The utility will attempt to automatically change the module mode when needed. Should the automatic change of mode fail for any reason, then the utility will prompt you to change the module state to either initialization or operational at various points during the pro cedure. See Checking and changing the mode on the HSM for more about changing the mode.

8.4.6. Verifying the integrity of the migrated keys

To verify the integrity of the migrated keys, run the `nfkmverify` utility with the following options, as appropriate:

- If the keys are module-protected, run the utility with one of the following options:
  - `-L` option, which checks the ACL of a loaded key instead of the generation certificate.
  - `-R` option, which checks the ACL of a key loaded from a recovery blob.
- If the keys are protected by cardsets or softcards, run the `nfkmverify` utility with the `-R` option in combination with the preload utility.

Example:

```
preload --admin=RE nfkmverify -R -m1 <application> <key-ident>
```

> ℹ️ Do not use the `nfkmverify` utility with the default `-C` option. If you

> use this option, the utility returns errors because the ACL in the certificate will reflect the old world.

ℹ️ Note that if the destination World is `fips-140-level-3` then some keys that were usable in the source World may not be usable in the destination World due to those algorithms or key lengths being restricted. The migration tool will successfully migrate the keys but they will be restricted at the point of use.

8.4.7. Migrating keys using custom protection pairs

Regular security world migration will create new card sets and softcards in the destination world with the same names as the source protections or it will use existing destination protections if they share a name and type (card set or softcard) with the source protection.

You can specify custom protection pairs if you want to change the name, the quorum, or the properties of the protection. You can also combine multiple source protections of the same type into one destination protection. You cannot diffuse keys from one source protection to multiple destination protections.

The source-destination protection pairs can be selected either as:

- Two comma-separated lists `--src-prots <source protections>` and `--dst-prots <destination protections>`.
- Tab-separated pairs "source destination", one per line, in a configuration file `--prots-config <file path>`.

The protections can be referred to by their name, 40-character hash, or "c:name" and "s:name" when a source card set and softcard share a name. The source and destination protection types must match.

The following example shows the two ways of specifying a set of protection pairs and the different ways each protection can be referred to. The example hashes are shortened for readability.

Protection type	Source protection to be migrated	Target destination protection
card set	`ocs 1`	`ocstarget1`
softcard	`softcard 1`	`softcardtarget`
card set	`name1` (duplicate name)	`ocstarget1`
softcard	`name1` (duplicate name)	`softcardtarget`

Protection type	Source protection to be migrated	Target destination protection
card set	name2 (duplicate name and type) hash: XXXXXXX1	ocstarget1
card set	name2 (duplicate name and type) hash: XXXXXXX2	ocstarget2

By specifying the lists using the `--src-prots` and `--dst-prots` options:

```
migrate-world [OPTIONS] \
--src-prots "ocs 1,softcard 1,c:name1,s:name1,XXXXXXX1,XXXXXXX2" \
--dst-prots "ocstarget1,softcardtarget,ocstarget1,softcardtarget,ocstarget1,ocstarget2"
```

By using a configuration file specified with the `--prots-config` option:

```
migrate-world [OPTIONS] --prots-config=migration.cfg

--- migration.cfg ---
ocs 1     ocstarget1
softcard 1     softcardtarget
c:name1     ocstarget1
s:name1     softcardtarget
XXXXXXX1     ocstarget1
XXXXXXX2     ocstarget2
--------------------
```

8.4.8. Troubleshooting

ⓘ | If you encounter any errors that are not listed in the following table, contact Support.

Error	Explanation	Action
There are no keys requiring migration.	Any migrate-able keys found in the source World already exist in the destination World. The migration utility returns this error if: • The keys have already been migrated • All keys are non-recoverable and therefore cannot be migrated.	None.

Error	Explanation	Action
Source module must be specified. Destination module must be specified. Source and Destination modules must be different. Module is not usable.	This utility requires you to specify both a source and destination module which must be different modules and both must be usable.	Specify the correct modules.
Source World has indistinguishable cardsets or softcards. Destination World has indistinguishable keys.	There are irregularities in one of the Worlds, but these irregularities do not affect the migration process.	None.
Destination World has indistinguishable cardsets or softcards. Source World has indistinguishable keys. Cannot determine protection of keys.	There are problems with one of the Worlds.	Contact Support.
Source World not recoverable.	The source World is not recoverable and the keys therefore cannot be migrated.	If the source World is not recoverable, you cannot use the migration utility to migrate keys. Contact Support.
Missing security World at PATH. Source world must be specified.	The path for the source World is wrong. There is no World data at the location that was specified when running the migration utility.	Supply the correct path to the source World. If you have supplied the correct path to the directory that contains the source World data, the migration utility has not found a destination World.

Error	Explanation	Action
Source World is the same as the destination World.	An incorrect path was supplied for the source World data when running the utility. The destination World data does not exist in the default location defined by the environment variable `NFAST_ KMLOCAL` or `NFAST_KM-DATA`.	Run the utility with the correct path to the source World data. Move the source World data to a different location and then copy the destination World data to the default location. If the default location is defined by an environment variable, configure the variable to point to the location of the destination World, which then becomes the new default location.
Cannot find <NAME> utility, needed by this utility. <NAME> utility is too old, need at least version <VERSION NUM-BER>.	The software installation is partially completed. The path (in the environment variable for the operating system) might be pointing to an old version of the software.	Reinstall the software. Ensure that the path points to the latest version of the software.
nFast error: TimeLimitExceeded; in response to SetKM...	The ACS time-out limit has expired.	Restart the key migration process; see Security World migration.
Destination world does not support audit logging.	You have specified the `--key-log-ging` option but the destination world does not support audit logging.	None. The keys will be migrated but LogKeyUsage will not be set in the ACL of migrated keys.
Failed to load warrant file <FILE>.	There is a problem reading the warrant file.	Check that your warrant files are in the correct location and have not been edited in any way.

8.5. Migrating KMDATA (Windows)

To move KMDATA from the default location of `C:\ProgramData\nCipher`:

1. Open a command prompt window as an administrator.

2. Use `Xcopy` with the following parameters to copy the default folder to a new location:

```
Xcopy C:\ProgramData\nCipher <Destination> /e /v /o /i
```

3. Enter the new location for the following environment variables:

   a. In the Windows Control Panel, navigate to **Control Panel > System and Security > System > Advanced system settings**.

    b. In the **Advanced** tab, select **Environment Variables**.

    c. Update the following system variables:

- **NFAST_CERTDIR**: `<path\to\new\folder>\Feature Certificates`
- **NFAST_KMDATA**: `<path\to\new\folder>\Key Management Data`
- **NFAST_LOGDIR**: `<path\to\new\folder>\Log Files`

    d. If your Security World client is on or above v12.70.4, add the following environment variable in the same section:

- **NFAST_KNETIDIR**: `<path\to\new\folder>\hardserver.d`

4. In the Services tool, restart the nFast Server process.

5. After the service restarts, run the following command to check the migration was successful:

```
anonkneti -m 127.0.0.1
```

6. After confirming that the migration was successful, delete `C:\ProgramData\nCipher`.

8.6. Erasing a module from a Security World

Erasing a module from a Security World deletes from the module all of the secret information that is used to protect your Security World. This returns the module to the factory state. Provided that you still have the ACS and the host data, you can restore the secrets by adding the module to the Security World.

Erasing a module removes any data stored in its nonvolatile memory (for example, data for an SEE program or NVRAM-stored keys). To preserve this data, you must back it up before erasing the module. We provide the `nvram-backup` utility to enable data stored in nonvolatile memory to be backed up and restored.

> **(!)** You do not need the ACS to erase a module. However, unless you have a valid ACS and the host data for this Security World, you cannot restore the Security World after you have erased it.

After you have erased a module, it is in the same state as when it left Entrust (that is, it has a random module key and a known K_{NSO}).

8.6.1. Erasing a module from the unit front panel

To erase a module from a Security World, from the main menu, select **Security World mgmt** > **Module initialization** > **Erase Security World**.

When you erase a Security World in this way, the Security World files remain on the remote file system. Delete these files if you wish to remove Security World completely. For more information, see Security World Files.

8.6.2. Erasing a module with new-world

The `new-world` command-line utility can erase any modules that are in the pre-initialization mode.

To erase modules with the `new-world` utility, run the command:

```
new-world [-e|--factory] [-m|--module=<MODULE>]
```

In the `new-world` command:

Option	Description
`-e|--factory`	These options restore a module to its factory state.
`-m|--module=<MODULE>`	These options specify the `ModuleID` to use. `new-world` erases only one module at a time. To erase multiple modules, you must run `new-world` once for every module that you want to erase.

8.6.2.1. Output

If `new-world` successfully erased a module, it displays a message that it restored the module to factory state. Otherwise, `new-world` returns an error message.

8.6.3. Erasing a module with KeySafe

You can erase a module on a server with KeySafe by following these steps:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)
2. Click the **World** menu button, or select **World** from the **Manage** menu. KeySafe takes you to the **World Operations** panel.
3. Click the **Erase Module** button. KeySafe takes you to the **Erase Module** panel.
4. Select the module that you want to erase by clicking its listing on the **Security world status** tree, then click the **Commit** command button.

   KeySafe erases all secrets from the module, returning it to its factory state.

> **ⓘ** If you have any keys that were protected by an erased module, you cannot access them unless you restore these secrets. You cannot restore these secrets unless you have the appropriate ACS.

8.6.4. Erasing a module with initunit

The `initunit` command-line utility erases any modules that are in the pre-initialization state.

To erase modules with the `initunit` utility, run the command:

```
initunit [-m|--module=<MODULE>] [-s|--strong-kml]
```

In the `initunit` command, `--module=<MODULE>` specifies the ID of the module you want to erase. If you do not specify this option, all modules in the pre-initialization state are erased. `--strong-kml` specifies that the module generates an AES (SP800-131A) module signing key, rather than the default key.

> **ⓘ** The `--disablepkcs1pad` option will only work on SP800-131A Security Worlds.

8.6.4.1. Output

If `initunit` is successful, for each module that is in the pre-initialization state, it returns a message similar to this:

```
 Initialising Unit
>#  Setting dummy HKNSO Module Key Info:
 HKNSO
>################## HKM
>##################
```

Otherwise, `initunit` returns an error message.

8.7. Replacing an existing Security World

When you erase a Security World from the module's front panel, all long-term key material is deleted from the module's memory and all Security World data is removed from the module's internal file system.

This operation does not remove any files from the remote file system or client machines. You should remove the files manually from the `/opt/nfast/kmdata/local` (**Linux**) or

`%NFAST_KMDATA%\local` (**Windows**) directory on the remote file system and any client computers to which the Security World was copied.

> Any Operator Cards created in a previous Security World cannot be used in the new Security World. If you are replacing a Security World, you must erase all the Operator Cards created in the previous Security World before you create the new Security World. See Erasing cards and softcards.

8.8. Deleting a Security World

You can remove an existing Security World and replace it with a new one if, for example, you believe that your existing Security World has been compromised. However:

- You are not able to access any keys that you previously used in a deleted Security World

- It is recommended that you reformat any nShield Remote Administration Cards that were used as Operator Cards within this Security World *before* you delete it. For more information about reformatting (or erasing) Operator Cards, see Erasing cards and soft cards.

> Except for nShield Remote Administration Cards, if you do not reformat the smart cards used as Operator Cards before you delete your Security World, you must throw them away because they cannot be used, erased, or reformatted without the old Security World key.

> You can, and should, reuse the smart cards from a deleted Security World's ACS. If you do not reuse or destroy these cards, then an attacker with these smart cards, a copy of your data (for example, a weekly backup) and access to any nShield key management HSM can access your old keys.

To delete an existing Security World:

1. Remove all the HSMs from the Security World.
2. Delete the Security World data files, see Location of Security World files.

> There may be copies of the Security World data archive saved on your backup media. If you have not reused or destroyed the old ACS, an attacker in possession of these cards could access your old keys using this backup media.

> ℹ️ If audit logging was enabled for the Security World then audit logs can still be verified provided that the audit log data is maintained as this contains all the information needed to verify the logs. For further information see *Audit Logging*.

8.8.1. Deleting the Security World using the nShield HSM front panel

When you erase a Security World using the unit front panel, all long-term key material is deleted from the HSM's memory and all Security World data is removed from the HSM's internal file system.

- You will not be able to access any of the keys that you have previously used
- Before you remove an old Security World, you must reformat any smart cards that were used previously as Operator Cards within this Security World.

> ℹ️ If you do not reformat the smart cards used as Operator Cards before you reinitialize your HSM, you must throw them away because they can not be used, erased, or reformatted without the old Security World key.

You can, and should, reuse the smart cards from the old ACS. If you do not reuse or destroy these cards, then an attacker with these smart cards, a copy of your data (for example, a weekly backup) and access to any nShield key management HSM, can access your old keys.

To erase a Security World using the front panel of the unit, from the main menu select **Security World mgmt** > **Module initialization** > **Erase Security World**.

This operation does not remove any files from the RFS or client machines. You should remove the files manually from the `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KM-DATA%\local` (**Windows**) directory on the RFS and any client computers to which the Security World was copied.

9. Managing card sets and softcards

This chapter describes how to create and manage card sets and softcards, using a Security World.

When you create a Security World, an Administrator Card Set (ACS) is created at the same time. You use the ACS to:

- Control access to Security World configuration
- Authorize recovery and replacement operations.

The Security World is used to create and manage keys, and the Operator Card Sets (OCSs) and softcards you create with the Security World are used to protect those keys.

A Security World offers three levels of key protection:

Level of protection	Description
Direct protection	Keys that are directly protected by the Security World are usable at any time without further authorization.
Softcard	Keys that are protected by a softcard can only be used by the operator who possesses the relevant passphrases.
OCS	Keys that are protected by an OCS can only be used by the operator who possesses the OCS and any relevant passphrases (if set).

For more information about creating a Security World, see Creating and managing a Security World.

For more information about key management, see Working with keys.

After a Security World has been created, you can use it to create and manage OCSs and softcards (as described in this chapter), as well as to create and manage the keys it protects (see Working with keys).

> **!** To perform the tasks described in this chapter, we recommend using the unit front panel or a client on the same computer that contains the RFS. To perform these tasks on a different client, you must transfer the card data to the RFS.

> **!** If you are sharing the Security World across several client computers, you must ensure that the changes are propagated to all your computers. One way to achieve this is to use client cooperation. For more information, see Setting up client cooperation.

If you want to use the Remote Operator feature to configure smart cards for use with a remote unit, see Remote Operator.

9.1. Creating Operator Card Sets (OCSs)

You can use an Operator Card Set (OCS) to control access to application keys. OCSs are optional, but if you require one, create it before you start to use the hardware security module with applications. You must create an OCS before you create the keys that it is to protect.

You can create OCSs that have:

- Names for individual cards, as well as a name for the whole card set
- Specific *K*/*N* policies
- Optional passphrases for any card within a given set
- Formal FIPS 140 Level 3 compliance.

> **i** Some third-party applications impose restrictions on the OCS smart card quorums (*K*/*N*) or the use of smart card passphrases. For more information, see the appropriate integration guide for the application. Integration guides for third-party applications are available from https://nshieldsupport.entrust.com/.

OCSs belong to the Security World in which they are created. When you create an OCS, the smart cards in that set can only be read by hardware security modules belonging to the same Security World.

Creating (and managing) OCSs can be done with the unit front panel, as described in Creating an Operator Card Set using the nShield HSM front panel. However, you can also use the following tools to create an OCS:

- The `createocs` command-line utility, as described in Creating an Operator Card Set using the command line
- KeySafe, as described in Creating an Operator Card Set with KeySafe
- (**Windows**) The nShield CSP wizard, as described in Creating an Operator Card Set with the CSP or CNG wizard (Windows)
- (**Windows**) The nShield CNG wizard, as described in Microsoft Cryptography API: Next Generation (CNG).

9.1.1. Persistent Operator Card Sets

If you create a standard (non-persistent) OCS, the keys it protects can only be used while the last required card of the quorum remains loaded in the local slot of the HSM, or one of its Dynamic Slots. The keys protected by this card are removed from the memory of the device as soon as the card is removed from the smart card reader. If you want to be able to use the keys after you have removed the last card, you must make that OCS persistent.

Keys protected by a persistent card set can be used for as long as the application that loaded the OCS remains connected to the hardware security module (unless that application removes the keys).

For more information about persistent OCSs, see Using persistent Operator Card Sets.

An OCS to be used to authorize login on a unit must be persistent and not loadable remotely. It is recommended that such an OCS is not used to protect sensitive keys.

9.1.2. Time-outs

OCSs can be created with a time-out, so that they can only be used for limited time after the OCS is loaded. An OCS is loaded by most applications at start up or when the user supplies the final required passphrase. After an OCS has timed out, it is not loadable by another application unless it is removed and reinserted. Time-outs operate independently of OCS persistence.

9.1.3. FIPS 140 Level 3-compliant Security Worlds

When you attempt to create an OCS for a Security World that complies with FIPS 140 Level 3, you are prompted to insert an Administrator Card or Operator Card from an existing set. You may need to specify to the application the slot you are going to use to insert the card. You need to insert the card only once in a session.

9.1.4. Creating an Operator Card Set using the nShield HSM front panel

To create an OCS, follow these steps:

1. From the main menu, select **Security World mgmt** > **Cardset operations** > **Create OCS**.

   You are prompted to enter the name of the OCS.

2. Enter a name and press right-hand navigation button.

3. Enter the quorum for the OCS, using the touch wheel to move from one field to the other. The quorum consists of:

- ◦ The maximum number of cards from the OCS required by default for an operation. This number must be less than or equal to the total number of cards in the set.

- ◦ The total number of cards to be used in the OCS. This must be a value in the range 1 − 64.

4. Press the right-hand navigation button to move to the next screen.

5. If you wish to specify a time out for the card set, enter the time out in seconds.

6. Choose whether to create a persistent card set. You can select:

   - ◦ **Not persistent** (which is the default)

   - ◦ **Persistent**

   - ◦ **Remoteable/Persistent**

7. Choose whether to name individual cards and enable passphrase replacement by answering **Yes** or **No** to each question and then pressing the right-hand navigation but ton.

8. Insert a smart card to be formatted for the OCS.

   If the card is not blank, choose whether to overwrite it or to use a different card. (If the card is an Operator Card from another Security World, you cannot overwrite it and are prompted to enter a different card.)

9. If you have chosen to name individual cards, you are prompted to enter the name for the card.

10. You are asked whether you wish to specify a passphrase for the card. If you choose **Yes**, you are prompted to enter the passphrase twice.

    While the Operator Card is being created, the screen displays the message **Processing**.

    If there are further cards from this OCS to be processed, the screen changes to **Waiting**. Remove the card, and repeat steps 8 through 10 for each of the remaining cards.

When all the cards in the set have been processed, you are told that the card set has been created successfully.

9.1.5. Creating an Operator Card Set using the command line

To create an OCS from the command line:

1. Run the command:

```
createocs -m <MODULE>|--module=<MODULE> -Q|--ocs-quorum=<K>/<N> +
-N|--name=<NAME> [-M|--name-cards] +
[[-p|--persist]|[-P|--no-persist]] [[-R|--no-pp-recovery]|--pp-recovery] +
[-q|--remotely-readable] [-T|--timeout=<TIME>] [-e|--erase]
```

This command uses the following options:

Option	Description
`-m <MODULE>|--module=<MOD-ULE>`	This option specifies the number of the hardware security module to be used to create the token. If you only have one hardware security module, `<MODULE>` is 1.
`-Q|--ocs-quorum=<K>/<N>`	In this option, `<K>` is the minimum required number of cards. If you do not specify the value `<K>`, the default is 1. ℹ️ Some applications do not have mechanisms for requesting that cards be inserted. Therefore any OCSs that you create for use with these applications must have `<K>`=1. `<N>` is the total number of cards. If you do not specify the value `<N>`, the default is 1.
`-N|--name=<NAME>`	This option specifies a name for the card set. The card set must be named with this option before individual cards can be named using the `-M/--name-cards=<NAME>` options.
`-M|--name-cards`	Specifying this option allows you to name individual cards within the card set. You can only use this option after the card set has been named by using the `--name=`NAME option. `createocs prompts for the names of the cards as they are created. Not all applications can display individual card names.
`-p|--persist`	This option creates a persistent card set.
`-P|--no-persist`	This option creates a non-persistent card set.
`-R|--no-pp-recovery`	This option specifies that passphrase replacement for this OCS is disabled. Setting this option overrides the default setting, which is that the card passphrases are replaceable. You can specify the enablement of passphrase replacement explicitly by setting the `--pp-recovery` option.
`-q|--remotely-readable`	This option allows this card set to be read remotely. For information on configuring Remote OCSs, see Remote Operator. ℹ️ Not required for Remote Administration.
`-T|--timeout=<TIME>`	This option sets the time-out for the card set. Use the suffix `s` to specify seconds, `m` for minutes, `h` for hours, and `d` for days. If the time-out is set to 0, the OCS never times out. Otherwise, the hardware security module automatically unloads the OCS when the amount of time specified by `TIME` has passed since the OCS was loaded.
`-e|--erase`	Specifying this option erases a card (instead of creating a card set). You can specify this option twice in the form `-ee` to repeatedly erase cards.

> ℹ With Security World Software v11.72 and later, passphrases are limited to a maximum length of 254 characters, when using `createocs`. See Maximum passphrase length.

If you have created a FIPS 140 Level 3 compliant Security World, you must provide authorization to create new Operator Cards; `createocs` prompts you to insert a card that contains this authorization. Insert any card from the Administrator Card Set or any Operator Card from the current Security World.

When `createocs` has obtained the authorization from a valid card, or if no authorization is required, it prompts you to insert a card.

2. Insert the smart card to use.

   If you insert an Administrator Card from another Security World or an Operator Card that you have just created, `createocs` displays the following message:

   ```
   Module x slot n: unknown card +
   Module x slot n: Overwrite card ? (press Return)
   ```

   where x is the hardware security module number and n is the slot number. If you insert an Operator Card from another Security World, `createocs` displays the following message:

   ```
   Module x slot n: inappropriate Operator Card (TokenAuthFailed).
   ```

   When you insert a valid card, `createocs` prompts you to type a passphrase.

   > ℹ The nShield PKCS #11 library requires Operator Cards with passphrases.

   > ℹ Some applications do not have mechanisms for entering passphrases. Do not give passphrases to Operator Cards that are to be used with these applications.

3. Type a passphrase and press **Enter**. Alternatively, press **Enter** if you do not want this card to have a passphrase.

   A passphrase can be of any length and can contain any character that you can type.

   If you entered a passphrase, `createocs` prompts you to confirm it.

4. Type the passphrase again and press **Enter**.

If the passphrases do not match, `createocs` prompts you to input and confirm the passphrase again.

5. When the new card has been created, if you are creating a card set with more than one card in it, `createocs` prompts you to insert another card.

6. For each additional card in the OCS, follow the instructions from step 2 through 4.

9.1.6. Creating an Operator Card Set with KeySafe

KeySafe enables you to create OCSs with:

- Their own names
- *K/N* policies
- Optional passphrases for any card within the OCS
- Formal FIPS 140 Level 3 compliance.

To create an OCS with KeySafe:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the **Card sets** menu button, or select **Card sets** from the menu.

   The **List Operator Card Sets** panel is displayed.

3. Select an HSM within the Security World from the Security World status pane.

4. Click the **Create new card set** button to open the **Create Operator Card Set** panel. You can specify the following options:

   a. A name for the card set.

   b. Whether passphrase recovery will be enabled for the OCS. (Only available if the Security World has passphrase recovery enabled.)

   c. Whether the card set can be used remotely. (Only available if the Security World has remote sharing available.) For more information, see Remote Operator.

   d. Whether this OCS will be persistent.

   e. Whether this OCS will have a time-out (a period after which the card set must be inserted again).

   f. The value for the time-out, in seconds.

   g. The total number of Operator Cards (*N*) that you want this OCS to have. This must be a value in the range $1 - 64$.

   h. The number of Operator Cards needed to re-create a key (*K*). *K* must be less than or equal to *N*.

5. When you have entered all the details, click **Commit**. KeySafe takes you to a new **Create Operator Card Set** panel.

> **ⓘ** If *K* is equal to *N*, a message is displayed:
>
> **The total number of cards is equal to the required number of cards. – If the total and required number of cards are equal, losing one card will render any nonrecoverable keys unusable. Is this what you want**?
>
> Click **Yes** to confirm the values for *K* and *N*, or **No** to change them.

> **ⓘ** If you are creating the card set in a FIPS 140 Level 3 Security World, insert an Administrator Card or an existing Operator Card when prompted.

6. Insert a blank, unformatted card into the reader.

   A message is displayed, confirming that the card is blank. Click **OK** to open the **Set Card Protection passphrase** panel.

> **ⓘ** If you insert a card from another OCS, KeySafe asks whether you want to erase it. If you insert an Administrator Card from the current Security World, KeySafe prevents you from accidentally erasing it. If you insert an OCS card from another Security World you will get the message:
>
> **Error. Unreadable card - may be incorrectly inserted or be from another Security World's operator card set. Please check.**
>
> To overcome this you must replace the card you have inserted with another card that is readable (or blank).

> **ⓘ** When creating a card set, KeySafe recognizes cards that already belong to the set before the card set is complete. If you accidentally insert a card to be written again after it has already been written, you receive a warning.

7. Select whether or not you want to set a passphrase for the currently inserted card. Each card in a set can have an individual passphrase, and you can also create a set in which some cards have passphrases and others do not.

8. If setting a passphrase for the currently inserted card, enter the same passphrase in both text fields. A passphrase can contain any characters you can type except for tabs

or carriage returns (because these keys are used to move between data fields).

> ⓘ You can change a passphrase at any time. If you do not set a passphrase now, you can use the KeySafe `Change passphrase` option (on the **Examine/Change Card** panel) to add one later. Likewise, if you later decide that you do not need a passphrase on a card, you can use this option to remove it.

9. After entering your desired passphrase (if any) in both text fields, click the **OK** button. Unless you have entered details for the last card in the set, KeySafe returns you to the **Create Operator Card Set** panel and prompts you to enter the next card in the set to be written.

10. After KeySafe has written the details of the last smart card in the set, it displays a dialog indicating that the OCS has been successfully created. Click the **OK** button, and KeySafe returns you to the Create Operator Card Set panel, where you can create another OCS or choose a different operation by clicking one of the menu buttons.

9.1.7. Creating an Operator Card Set with the CSP or CNG wizard (Windows)

You can use the nShield CSP or CNG wizard to create a *K*/*N* OCS that is suitable for use with the nShield Cryptographic Service Provider (CSP) or Cryptography API: Next Generation (CNG), as appropriate. You can only create an OCS using the CSP or CNG wizard if you already have a Security World and have an ACS available for that Security World.

To create an OCS using the CSP or CNG wizard, follow these steps:

1. Ensure that you have created the Security World and that at least one HSM is in the operational state.

2. Run the wizard by double-clicking its shortcut in the Start menu: **Start > Entrust nShield Security World**.

3. The wizard displays the welcome screen.

4. Click the **Next** button. The wizard allows you to configure HSM Pool mode for CAPI/CNG.

> ⓘ Do not enable HSM Pool mode when creating an Operator Card Set because HSM Pool mode only supports module-protected keys.

5. Click the **Next** button.

The wizard determines what actions to take based on the state of the Security World and of the HSMs that are attached to your computer:

- ◦ If the wizard cannot find the Security World, it prompts you to create a new Security World or to install cryptographic acceleration only.

    In such a case, you should:

    - ▪ Cancel the operation
    - ▪ Check that the environment variable `NFAST_KMDATA` is set correctly
    - ▪ Copy the `local` sub-directory from the Key Management Data directory of another computer in the same Security World or from a backup tape of this computer to the Key Management Data directory of this computer.
    - ▪ run the wizard again.

- ◦ If there is an existing Security World, the wizard gives you the option of using the existing Security World, creating a new Security World or installing cryptographic acceleration only.

    - ▪ In order to use the existing Security World, ensure that the **Use the existing security world** option is selected, and click the **Next** button.
    - ▪ If there are any HSMs in the pre-initialization state, the wizard adds them to the Security World; see Adding or restoring an HSM to the Security World.

6. When at least one hardware security module is in the operational state, the wizard prompts you to select a method to protect private keys generated by the CSPs.

7. Ensure that the Operator Card Set option is enabled. If you are running the CNG wizard (not the CSP wizard) click the **Next** button. Then select the **Create a new Operator Card Set** option.

    If you want the OCS to be persistent, select the **Persistent** option. Persistence is described in Persistent Operator Card Sets.

8. Click the **Next** button, and if you have a FIPS world, the wizard prompts you to insert a card created with the current Security World.

    > **ℹ** This shows that your Security World is compliant with the roles and services of the FIPS 140 Level 3 standard. It is included for those customers who have a regulatory requirement for compliance.

    Under the constraints of level 3 of the FIPS 140 standard, Operator Cards cannot be created without authorization. To obtain authorization, insert any card from the ACS or any Operator Card belonging to the current Security World.

    The wizard does not enable the next world, the wizard warns you and prompts you for another card.

9. Click the **Next** button.

The wizard prompts you for a smart card to use as the first card in the OCS.

10. Insert a blank smart card to be used as the Operator Card, and click the **Next** button.

> Do not use a card from the ACS or an existing Operator Card.
>
> If you insert a card that is not blank, the wizard asks you if you want to erase it.

11. When you have inserted an appropriate card, the wizard prompts you for the name of the card and, if required, a passphrase.

    If you want to protect this card with a passphrase, turn on the **Card will require a passphrase** option, and enter the passphrase. You must enter the passphrase in both fields to ensure that you have typed it correctly.

    > Operator Cards with passphrases are required by the nShield PKCS #11 library.

12. If you have not yet written all the smart cards in the OCS, the wizard prompts you for another card. Repeat the appropriate preceding steps of the OCS creation process for all smart cards in the set.

13. When the wizard has finished creating the OCS, it displays a screen telling you this. If you want to create another OCS, click the **Back** button on this screen.

    When you have created all the OCSs that you require, click the **Next** button to install the CAPI CSP or register the CNG CSP. For more information, see Installing the CAPI CSP or Registering the CNG CSP.

9.2. Creating softcards

You must create a softcard before you create the keys that it is to protect.

A softcard is a file containing a logical token that cannot be loaded without a passphrase; its logical token must be loaded in order to authorize the loading of any key that is protected by the softcard. Softcard files are stored in the Key Management Data directory and have names of the form `softcard_<hash>` (where `<hash>` is the hash of the logical token share). Softcards belong to the Security World in which they are created.

A softcard's passphrase is set when you generate it, and you can use a single softcard to protect multiple keys. Softcards are persistent; after a softcard is loaded, it remains valid for loading the keys it protects until its `KeyID` is destroyed.

> It is possible to generate multiple softcards with the same name or

> passphrase. For this reason, the hash of each softcard is made unique (unrelated to the hash of its passphrase).

> **ℹ** Softcards are not supported for use with the nCipherKM JCA/JCE CSP in Security Worlds that are compliant with FIPS 140 Level 3.

> **ℹ** To use softcards with PKCS #11, you must have `CKNFAST_LOADSHARING` set to a nonzero value. When using pre-loaded softcards or other objects, the PKCS #11 library automatically sets `CKNFAST_LOADSHARING=1` (load-sharing mode on) unless it has been explicitly set to `0` (load-sharing mode off).

> **❗** As with OCSs, if debugging is enabled, a softcard's passphrase hash is available in the debug output (as a parameter to a `ReadShare` command).

You can create softcards from either:

- The command-line (see Creating a softcard with ppmk)
- KeySafe (see Creating softcards with KeySafe)

9.2.1. Creating a softcard with ppmk

To create a new softcard using the `ppmk` command-line utility:

1. Decide whether you want the new softcard's passphrase to be replaceable or non-replaceable. To create a softcard with a replaceable passphrase, run the command:

   ```
   ppmk --new --recoverable <NAME>
   ```

   To create a softcard with a non-replaceable passphrase, run the command:

   ```
   ppmk --new --non-recoverable <NAME>
   ```

   In these commands, `<NAME>` specifies the name of the new softcard to be created.

2. When prompted, type a passphrase for the new softcard, and press **Enter**.

   A passphrase can be of any length and contain any characters that you can type except for tabs or carriage returns (because these keys are used to move between data fields).

3. When prompted, type the passphrase again to confirm it, and press **Enter**.

If the passphrases do not match, ppmk prompts you to input and confirm the passphrase again.

After you have confirmed the passphrase, ppmk completes creation of the new softcard.

9.2.2. Creating softcards with KeySafe

To create a softcard with KeySafe:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the **Softcards** menu button, or select **Softcards** from the **Manage** menu. KeySafe takes you to the **List Softcards** panel.

3. Click **Create New Softcard** to open the **Create Softcard** panel.

4. Choose parameters for the softcard:

   a. Enter a name for the softcard. You must provide a valid name for each softcard.

   b. Choose whether you want passphrase replacement to be enabled for the softcard.

   > ⓘ In a Security World with passphrase recovery enabled the **Yes** radio button is selected as default and the selection can be changed between **Yes** and **No**. In a Security World with passphrase recovery disabled the **No** button is selected, and cannot be changed to **Yes**.

5. Click **Commit**.

   > ⓘ If you are creating the softcard in a FIPS 140 Level 3 Security World, insert an Administrator Card or an existing Operator Card when prompted.

   The **Set Softcard Protection passphrase** pane is displayed.

6. Set a passphrase for the softcard by entering the same passphrase in both text fields.

   A passphrase can contain any characters you can type except for tabs or carriage returns (because these keys are used to move between data fields) and can be up to 1024 characters long. You can change a passphrase at any time. You must provide a passphrase for each card.

7. After entering your desired passphrase in both text fields, click the **OK** button.

   KeySafe displays a dialog indicating that the softcard has been successfully created.

8. Click the **OK** button.

   KeySafe returns you to the **Create Softcard** panel, where you can create another soft-
   card or choose a different operation by clicking one of the menu buttons.

9.2.3. Creating a softcard with the CNG wizard (Windows)

You can use the nShield CNG wizard to create a Softcard that is suitable for use with the
nShield Cryptography API: Next Generation (CNG), as appropriate. You can only create an
Softcard using the CNG wizard if you already have a Security World and have an ACS avail-
able for that Security World.

To create an Softcard using the CNG wizard, follow these steps:

1. Ensure that you have created the Security World and that at least one HSM is in the
   operational state.
2. Run the wizard by double-clicking its shortcut in the Windows Start menu: **Start >
   Entrust nShield Security World**.
3. The wizard displays the welcome screen.
4. Click the **Next** button. The wizard allows you to configure HSM Pool mode for
   CAPI/CNG.

   > Do not enable HSM Pool mode when creating a Softcard because
   > HSM Pool mode only supports module-protected keys.

5. Click the **Next** button.

   The wizard determines what actions to take based on the state of the Security World
   and of the HSMs that are attached to your computer:

   ◦ If the wizard cannot find the Security World, it prompts you to create a new Secu-
     rity World or to install cryptographic acceleration only.

     In such a case, you should:

     ▪ Cancel the operation

     ▪ Check that the environment variable `NFAST_KMDATA` is set correctly

     ▪ Copy the `local` sub-directory from the Key Management Data directory of
       another computer in the same Security World or from a backup tape of this
       computer to the Key Management Data directory of this computer.

     ▪ Run the wizard again.

   ◦ If there is an existing Security World, the wizard gives you the option of using the

existing Security World, creating a new Security World or installing cryptographic acceleration only.

- In order to use the existing Security World, ensure that the **Use the existing security world** option is selected, and click the **Next** button.

- If there are any hardware security modules in the pre-initialization state, the wizard adds them to the Security World; see Adding or restoring an HSM to the Security World.

6. When at least one hardware security module is in the operational state, the wizard prompts you to select a method to protect private keys generated by the CSPs.

7. Ensure that the Softcard option is enabled. Click the **Next** button. Then select the **Create a new Softcard** option, and enter the name and passphrase of the Softcard in the boxes provided.

8. Click the **Next** button, and if you have a FIPS world, the wizard prompts you to insert a card created with the current Security World.

> ℹ️ This shows that your Security World is compliant with the roles and services of the FIPS 140 Level 3 standard. It is included for those customers who have a regulatory requirement for compliance.

Under the constraints of level 3 of the FIPS 140 standard, Softcards cannot be created without authorization. To obtain authorization, insert any card from the ACS or any OCS belonging to the current Security World.

9. On the Software Installation screen when you are informed **You now have a valid security world and key protection mechanism**, click the Back button if you want to create another Softcard, or if you want to change the default protection for new CNG keys to a different protection option. When you have created all the Softcards that you require, click the **Next** button on this screen to register the CNG providers. For more information, see Registering the CNG CSP.

9.3. Erasing cards and softcards

Erasing a card or softcard removes all the secret information from the card or softcard and deletes information about the card or softcard from the host.

> ❗ In the case of an OCS that uses nShield Remote Administration Cards, it is possible to reformat the cards at any time using `slotinfo --ignore-auth`. In the case of an OCS that uses standard nShield cards, it is only possible to erase or format the cards within the Security World in which they were created.

You can erase Operator Cards using the unit front panel, KeySafe or the `createocs` utility. You can also use these methods to erase Administrator Cards other than those in the current Security World's ACS (for example, you could use these methods to erase the remaining Administrator Cards from an incomplete set that has been replaced or Administrator Cards from another Security World).

> **ℹ** None of these tools erases cards from the current Security World's ACS.

If you erase an Operator Card that is the only card in an OCS, information about the card set is deleted. However, if you erase one card from an OCS of multiple cards, you must remove the card information from the `opt/nfast/kmdata/local/` (**Linux**) or `%NFAST_KM-DATA\local%` (**Windows**) directory after you have erased the last card.

> **ℹ** You can erase an entire card set at one time with the KeySafe **Remove OCS!** feature. For more information, see List an Operator Card Set.

9.3.1. FIPS 140 Level 3-compliant Security Worlds

When you attempt to erase cards for a Security World that complies with FIPS 140 Level 3, you are prompted to insert an Administrator Card or Operator Card from an existing set. You may need to specify to the application the slot you are going to use to insert the card. You need to insert the card only once in a session. You can therefore use one of the cards that you are about to erase.

9.3.2. Erasing card sets using the nShield HSM front panel

To erase a card set using the front panel, follow this procedure:

1. From the main menu select: **Security World mgmt** > **Card operations** > **Erase card**
2. Insert the card set that you want to erase. The card is read.
3. You are asked to confirm that you want to erase this card from the card set.
4. To confirm, press the right-hand navigation button.
5. You are asked once again if you want to erase this card.
6. To confirm, press the right-hand navigation button.

9.3.3. Erasing cards with KeySafe

To erase a card using KeySafe use the following procedure:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the **Card Sets** menu button. KeySafe takes you to the Card Operations panel.

3. Click the **Examine/Change Card** navigation button. KeySafe takes you to the **Examine/Change Card** panel.

4. Insert the card that you want to erase into the reader.

5. Click the **Erase Card** button. You do not need to supply the passphrase (if there is one) to erase an Operator Card.

6. KeySafe asks you to confirm that you want to erase this card. If you are sure that you want to erase it, click the Yes button.

> Erasing a card does not erase the keys protected by that card. The keys are still listed on the keys panel but are unusable.

If you erase an Operator Card that is the only card in an OCS, KeySafe deletes information about that card set. However, if you erase one card from an OCS of multiple cards, you must remove the card information from `opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA\local%` (**Windows**) after you have erased the last card.

7. After erasing a card, KeySafe displays a dialog to confirm that the card has been erased. Click **OK** to continue using KeySafe.

> You can erase an entire card set at one time with the KeySafe **Discard Card Set(s)** feature.

9.3.4. Erasing cards using the command line

To erase a card from the command line, run the command:

```
createocs -m|--module=<MODULE> -e|--erase
```

This command uses the following options:

Option	Description
`-m|--module=<MODULE>`	These options specify the module number of the module. If you only have one module, *MODULE* is 1.
`-e|--erase`	These options specify that you want to erase a card (rather than create an OCS).

> If you have more than one card reader and there is more than one card

> available, `createocs` prompts you to confirm which card you wish to
> erase. Use **[Ctrl][X]** to switch between cards.

If you have created a FIPS 140 Level 3 compliant Security World, you must provide autho-
rization in order to erase or create Operator Cards. You can obtain this authorization from
any card in the ACS or from any Operator Card in the current Security World, including
cards that are to be erased. After you insert a card containing this authorization, `createocs`
prompts you to insert the card to be erased.

As an alternative, you can reformat using `slotinfo --format`.

9.3.5. Erasing softcards

Erasing a softcard deletes all information about the softcard from the host. You can erase
softcards using KeySafe or with the `ppmk` command-line utility.

9.3.5.1. Erasing softcards with KeySafe

To erase softcards with KeySafe:

1. Start KeySafe.
2. Click the **Softcards** menu button. KeySafe takes you to the Softcard Operations panel.
3. Select the softcard you want to erase from the list.
4. Click the **Discard Softcard** button.
5. KeySafe asks you to confirm that you want to erase this card. Click **Yes** to confirm.
6. After erasing a softcard, KeySafe displays a dialog box to confirm that the card has
   been erased. Click **OK** to continue using KeySafe.

9.3.5.2. Erasing softcards with ppmk

To erase a softcard with `ppmk`, open a command window, and give the command:

```
ppmk --delete <NAME>|<IDENT>
```

In this command, you can identify the softcard to be erased either by its name (`NAME`) or by
its logical token hash as listed by `nfkminfo` (`<IDENT>`).

If you are working within a FIPS 140 Level 3 compliant Security World, you must provide
authorization to erase softcards; `ppmk` prompts you to insert a card that contains this autho-
rization. Insert any card from the ACS or any Operator Card from the current Security

---

World.

If you insert an Administrator Card from another Security World or an Operator Card that you have just created, `ppmk` displays an error message and prompts you to insert a card with valid authorization. When `ppmk` has obtained the authorization from a valid card or if no authorization is required, it completes the process of erasing the softcard.

9.4. Viewing cards and softcards

It is often necessary to obtain information from card sets, usually because for security reasons they are left without any identifying markings.

To view details of all the Operator Cards in a Security World or details of an individual Operator Card, you can use the front panel, KeySafe or the `nfkminfo` command-line utility. To check which passphrase is associated with a card, you can use the front panel or the `cardpp` command-line utility.

To list all softcards in a Security World or to show details of an individual softcard, you can use the `ppmk` or `nfkminfo` command-line utilities. To check which passphrase is associated with a softcard, you can use the `ppmk` command-line utility.

9.4.1. Viewing card sets using the nShield HSM front panel

You can use the unit front panel to view details of all the Operator Cards in a Security World or to view details of an individual Operator Card.

To view a list of all the card sets in the Security World, from the front panel select **Security World mgmt** > **Cardset operations** > **List cardsets**.

To view details of a single card using the unit front panel:

1. Insert the card into the unit.
2. From the main menu, select **Security World mgmt** > **Card operations** > **Card details**.
3. The type of the card (Administrator or Operator) is displayed with the number of the card in the card set.

9.4.2. Viewing card sets with KeySafe

You can use KeySafe to view details of all the Operator Cards in a Security World, details of individual OCSs or details of an individual Operator Card.

9.4.2.1. Examining a Card

In order to view information about individual cards with KeySafe, follow these steps:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the **Card Sets** menu button, or select the **Card sets** menu item from the **Manage** menu. KeySafe takes you to the **List Operator Card Sets** panel.

3. Click **Examine/Change Card** to open the **Examine/Change Card** panel.

4. Insert a card into the appropriate smart card slot. KeySafe displays information about the smart card currently in the slot. If there is no smart card in the slot, KeySafe displays a message **Card slot empty - please insert the card that you want to examine.**

From the **Examine/Change Card** panel, you can also:

- Change a card's passphrase (if it has one)
- Give a passphrase to a card that does not already have one
- Remove a passphrase from a card that currently has one
- Erase the card.

9.4.2.2. List an Operator Card Set

In order to view information about whole OCSs with KeySafe, follow these steps:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the `Card Sets` menu button, or select the **Card sets** menu item from the **Manage** menu. KeySafe takes you to the `List Operator Card Sets` panel, which displays information about all OCSs in the current Security World.

From the `List Operator Card Sets` panel, you can also:

- Examine / change a card (see Examining a Card)
- Create a new card set (see Creating an Operator Card Set with KeySafe)
- Replace an Operator Card Set (see Replacing OCSs with KeySafe)
- Discard a card set (see Erasing cards with KeySafe).

9.4.3. Viewing card sets using the command line

You can use the `nfkminfo` command-line utility to view details of either all the Operator Cards in a Security World or of an individual Operator Card.

To list the OCSs in the current Security World from the command line, open a command window, and give the command:

```
nfkminfo --cardset-list
```

In this command, `--cardset-list` specifies that you want to list the operator card sets in the current Security World.

`nfkminfo` displays output information similar to the following:

```
Cardset summary - 1 cardsets:          (in timeout, P=persistent, N=not)
 Operator logical token hash              k/n timeout name
hash                                     1/1 none-N name
```

To list information for a specific card, use the command:

```
nfkminfo <TOKENHASH>
```

In this command, `<TOKENHASH>` is the `Operator logical token hash` of the card (as listed when the command `nfkminfo --cardset-list` is run).

This command displays output information similar to the following:

```
name          "name"
k-out-of-n    1/1
flags         NotPersistent
timeout       none
card names    ""
hkltu         794ada39038fa8c4e9ea46a24136bbb2b8b337f2
```

> **ℹ** Not all software can give names to individual cards.

9.4.4. Viewing softcards

To view softcards, use KeySafe or the command line. The command line provides several options for viewing softcard information.

9.4.4.1. Viewing softcards with KeySafe

To view a softcard with KeySafe, follow these steps:

1. Start KeySafe.
2. Click the **Softcards** menu button. KeySafe takes you to the **Softcard Operations** panel.

3. Click the **List Softcards** navigation button. KeySafe takes you to the **List Softcards** panel, which displays information about all softcards in the current Security World.

From the **List Softcards** panel, you can also choose to remove a softcard from the Security World. For more information about this procedure, see Erasing cards and softcards.

9.4.4.2. Viewing softcards with nfkminfo

To list the softcards in the current Security World using the `nfkminfo` command-line utility, give the command:

```
nfkminfo --softcard-list
```

In this command `--softcard-list` specifies that you want to list the softcards in the current Security World.

To show information for a specific softcard using the `nfkminfo` command-line utility, give the command:

```
nfkminfo --softcard-list <IDENT>
```

In this command `<IDENT>` is the softcard's logical token hash (as given by running the command `nfkminfo --softcard-list`). This command displays output information similar to the following:

```
SoftCard
  name        "mysoftcard"
  hkltu       7fb95888ea2850d4e3ffcc8f0c22100937344308
Keys protected by softcard 7fb95888ea2850d4e3ffcc8f0c22100937344308:
  AppName simple              Ident mykey
  AppName simple              Ident myotherkey
```

9.4.4.3. Viewing softcards with ppmk

To list the softcards in the current Security World using the `ppmk` command-line utility, use the command:

```
ppmk --list
```

In this command `--list` specifies that you want to list the softcards in the current Security World.

In order to view the details of a particular softcard using the `ppmk` command-line utility, give the command:

```
ppmk --info <NAME>|<IDENT>
```

In this command, you can identify the softcard whose details you want to view either by its name (`<NAME>`) or by its logical token hash (as given by running the command `nfkminfo --softcard-list`).

9.4.5. Verifying the passphrase of a card or softcard

9.4.5.1. Verifying the passphrase of a card using the nShield HSM front panel

To verify the passphrase associated with a card using the unit front panel:

1. Insert the card into the unit.
2. From the main menu, select **Security World mgmt** > **Card operations** > **Check PIN**.

   The type of the card (Administrator or Operator) is displayed with the number of the card in the card set.

3. If this is the card that you want to check, press the right-hand navigation to confirm.
4. Enter the passphrase.

   If the passphrase that you entered is correct, a confirmation message is shown. Otherwise, an error is reported.

9.4.5.2. Verifying the passphrase of a card with cardpp

To verify the passphrase associated with a card using the `cardpp` command-line utility, use the command:

```
cardpp --check [-m|--module=<MODULE>]
```

This command uses the following options:

Option	Description
`--check`	This option checks the passphrase.
`--module=<MODULE>`	This option specifies the number of the module to use. If you only have one module, `<MODULE>` is 1. If you do not specify a module number, `cardpp` uses all modules by default.

The `cardpp` utility polls all available slots; if there is no card inserted, it prompts you to insert one. If the card belongs to this Security World, `cardpp` either tells you if no passphrase is set or prompts you to enter the passphrase and checks to see if it is correct.

9.4.5.3. Verifying the passphrase of a softcard with ppmk

In order to verify the passphrase of a particular softcard, open a command window, and give the command:

```
ppmk --check <NAME>|<IDENT>
```

In this command, you can identify the softcard whose passphrase you want to verify either by its name (`<NAME>`) or by its logical token hash (as given by running the command `nfk-minfo --softcard-list`).

`ppmk` prompts you to enter the passphrase and then tells you whether the passphrase you entered is correct for the specified softcard.

9.5. Changing card and softcard passphrase

Each softcard or card of a card set can have its own individual passphrase: you can even have a card set in which some cards have a passphrase and others do not, and you can have distinct softcards that nevertheless use the same passphrase. A passphrase can be of any length and can contain any characters that you can type.

Normally, in order to change the passphrase of a card or softcard, you need the card or soft card and the existing passphrase. Known card passphrase can be changed using the front panel, KeySafe or the `cardpp` command-line utility; softcard passphrase can be changed using KeySafe or the `ppmk` command-line utility. You can also add a passphrase to a card or softcard that currently does not have one or remove a passphrase from a card that does cur rently have one.

If you generated your Security World with the passphrase replacement option, you can also replace the passphrase of a card or softcard even if you do not know the existing passphrase. Such a passphrase replacement operation requires authorization from the ACS.

9.5.1. Changing known passphrase

To change a card passphrase, you need the card and the old passphrase.

Each card in a set can have its own individual passphrase. You can even have a set in which

some cards have a passphrase and others do not.

> **ℹ** Prior to Security World Software v11.72, we set no absolute limit on the length of a passphrase. However, some applications may not accept a passphrase longer than 255 characters. Likewise, the Security World does not impose restrictions on which characters you can use, although some applications may not accept certain characters. Entrust recommends that your password only contains 7-bit ASCII characters:
>
> `A-Z, a-z, 0-9, ! @ # $ % ^ & * - _ + = [ ] { } | \ : ' , . ? / ` `~ " < > ( ) ;`
>
> See Maximum passphrase length for more about passphrase length when using Security World Software v11.72.

9.5.1.1. Changing known passphrase from the unit front panel

To change the passphrase of a card using the unit front panel:

1. Insert the card.
2. From the main menu, select **Security World mgmt** > **Card operations** > **Change PIN**.
3. Select the card whose passphrase you want to change.
4. Enter the old passphrase, and then enter it again to confirm it.
5. Enter the new passphrase. If you do not want this card to have a passphrase, select **NO** at the prompt.

9.5.1.2. Changing known passphrase with KeySafe

To change a known passphrase for an Operator Card using KeySafe:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)
2. Click **Card sets**, or select **Card sets** from the **Manage** menu. The **List Operator Card Sets** panel is displayed.
3. Click **Examine / change card** to open the **Examine / Change Card** panel.
4. Click **Change passphrase**. The **Set Card Protection passphrase** panel is displayed.
5. Enter the old passphrase, and click the **OK** button.
6. A screen is displayed asking **Do you want to set a passphrase?**. Select **Yes**.
7. Enter your new passphrase, and enter it again in the second box as confirmation of the change.

8. Click **OK**.

9.5.1.2.1. Changing a known softcard passphrase with KeySafe

To change a known passphrase for a softcard using KeySafe:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the **Softcards** menu button, or select **Softcards** from the **Manage** menu. KeySafe takes you to the **List Softcards** panel.

3. Select the softcard for which you want to change the passphrase, and click the **Change passphrase** button. KeySafe takes you to the **Change/Recover Softcard passphrase** panel.

> **i** If a softcard is listed as PIN Recovery Enabled = No, then you will be unable to change the passphrase.

4. Select the softcard whose passphrase you want to change, and click the **Change passphrase** button. KeySafe takes you to the **Get Softcard Protection passphrase** panel.

5. Enter the old passphrase, and click the **OK** button.

   KeySafe either displays an error dialog (if the passphrase is not correct) or takes you to the `Set Softcard Protection` passphrase panel.

6. Enter your new passphrase, and enter it again in the second field to confirm the passphrase is correct.

7. Click the **OK** button.

   After changing a passphrase, KeySafe displays a dialog to confirm that the passphrase has been successfully changes.

8. Click the **OK** button to continue using KeySafe.

9.5.1.3. Changing known passphrase with cardpp

Each card in a card set can have its own individual passphrase. You can even have a set in which some cards have a passphrase and others do not. A passphrase can be of any length and can contain any characters that you can type.

> **i** With Security World Software v11.72 and later, passphrases are limited to a maximum length of 254 characters, when using `cardpp`. See Maxi-

mum passphrase length.

To change a known card's passphrase with the `cardpp` command-line utility, take the following steps:

1. Run the `cardpp` utility using the command:

```
cardpp --change [-m|--module=<MODULE>]
```

   If you only have one HSM, `<MODULE>` is 1. If you do not specify an HSM number, `cardpp` uses all HSMs by default.

2. If prompted, insert the card whose passphrase you want to change. (If there is a card already in the slot, you are not prompted.)

3. If prompted, enter the existing passphrase for the card (If the card has no current passphrase you are not prompted.) If you enter the passphrase correctly, `cardpp` prompts you to enter the new passphrase.

4. Enter a new passphrase, and then enter it again to confirm it.

   After you have confirmed the new passphrase, `cardpp` changes the card's passphrase.

9.5.1.4. Changing known softcard passphrase with ppmk

> **ℹ** With Security World Software v11.72 and later, passphrases are limited to a maximum length of 254 characters, when using `ppmk`. See Maximum passphrase length for more information.

To change a known softcard's passphrase when you know the passphrase, follow these steps:

1. Give the following command:

```
ppmk --change <NAME>|<IDENT>
```

   In this command, you can identify the softcard whose passphrase you want to change either by its name (`<NAME>`) or by its logical token hash as listed by `nfkminfo` (`<IDENT>`).

   `ppmk` prompts you to enter the old passphrase.

2. Type the old passphrase, and press **Enter**. If you enter the old passphrase correctly, `ppmk` prompts you to enter the new passphrase.

3. Type the old passphrase, and press **Enter**. Type the new passphrase again, and press

Enter to confirm it.

After you have confirmed the new passphrase, `ppmk` then changes the softcard's passphrase.

9.5.2. Changing unknown or lost passphrase

9.5.2.1. Changing unknown card passphrase with cardpp

If you generated your Security World with the passphrase replacement option, you can change the passphrase of a card even if you do not know its existing passphrase. Such a passphrase replacement operation requires authorization from the ACS.

To change an unknown card passphrase with the `cardpp` command-line utility:

- Run a command of the form:

```
cardpp --recover [--module=<MODULE>]
```

  In this command, `<MODULE>` specifies the number of the hardware security module to use. If you only have one hardware security module, `<MODULE>` is 1. If you do not specify a number, `cardpp` uses all hardware security modules by default.

- As prompted, insert the appropriate number of cards from the ACS required to authorize passphrase replacement.
- When prompted, insert the Operator Card whose passphrase you want to replace. To replace its passphrase:

  a. When prompted, type the new passphrase, and then press **Enter**.

  b. When prompted, type the new passphrase again to confirm it, and then press **Enter**.

     `cardpp` sets the new passphrase, and then prompts you for another Operator Card.

- Repeat the process in the previous step to change the passphrase on further cards, or press **Q** to quit.

> **(!)** Only insert Administrator Cards into a hardware security module that is connected to a trusted server.

9.5.2.2. Replacing unknown passphrase with ppmk

If you generated your Security World with the passphrase replacement option, you can

change the passphrase of a softcard even if you do not know its existing passphrase. Such a passphrase replacement operation requires authorization from the ACS.

To change an unknown softcard passphrase with the `ppmk` command-line utility:

1. Run a command of the form:

   ```
   preload --admin=p ppmk --recover <NAME>|<IDENT>
   ```

   In this command, you can identify the softcard by its `<NAME>` or by its `<IDENT>` (its logical token hash as shown in output from the `nfkminfo` command-line utility).

2. As prompted, insert the appropriate number of cards from the ACS required to authorize passphrase replacement.

3. When prompted, type the new passphrase, and then press **Enter**.

4. When prompted, type the new passphrase again to confirm it, and then press **Enter**.

   If the passphrase does not match, `ppmk` prompts you to input and confirm the passphrase again.

After you successfully confirm the new passphrase, `ppmk` finishes configuring the softcard to use the new passphrase.

> **!** Only insert Administrator Cards into a hardware security module that is connected to a trusted server.

9.6. Replacing Operator Card Sets

> **ℹ** Replacing an OCS requires authorization from the ACS of the Security World to which it belongs. You cannot replace an OCS unless you have the required number of cards from the appropriate ACS.

If you have lost a card from a card set, or you want to migrate from standard nShield cards to nShield Remote Administration Cards, you should use one of the following:

- The `rocs` utility
- The KeySafe **Replace Operator Card Set** option.

  Accessed from the `Card Operations` panel.

> **ℹ** You cannot mix standard nShield cards with nShield Remote Administration Cards in the same set.

We recommend that after you have replaced an OCS, you then erase the remaining cards in the old card set and remove the old card set from the Security World. For more information, see Erasing cards and softcards.

Deleting the information about an OCS from the client does not remove the data for keys protected by that card set. On the KeySafe **Key Operations** panel), such keys are listed as being protected by **Deleted Card Set**.

To prevent you from losing access to your keys if the smart card you are using as the Operator Card is lost or damaged, Entrust supplies several utilities that can recover the keys protected by the lost Operator Card to another OCS

Replacing one OCS with another OCS also transfers the keys protected by the first OCS to the protection of the new OCS.

When you replace an OCS or softcard and recover its keys to a different OCS or softcard, the key material is not changed by the process. The process deletes the original Security World data (that is, the encrypted version of the key or keys and the smart card or softcard data file) and replaces this data with host data protected by the new OCS or softcard.

To replace an OCS or softcard, you must:

- Have enabled OCS and softcard replacement when you created the Security World

  > If you did not enable OCS and softcard replacement, you cannot recover keys from lost or damaged smart cards or softcards.

- Have created the original OCS using the front panel of the unit, `createocs`, `createocs-simple`, KeySafe, or the nShield PKCS #11 library version 1.6 or later

  > If you initialized the token using `ckinittoken` from the nShield PKCS #11 library version 1.5 or earlier, you must contact Support to arrange for them to convert the token to the new format while you still possess a valid card.

- Have a sufficient number of cards from the ACS to authorize recovery and replacement

  > All recovery and replacement operations require authorization from the ACS. If any of the smart cards in the ACS are lost or damaged, immediately replace the entire ACS.

- Have initialized a second OCS using the front panel of the unit, `createocs`, `createocs-simple`, KeySafe, or the nShield PKCS #11 library version 1.6 or later.

> **ⓘ** | The new OCS need not have the same *K/N* policy as the old set.

If you are sharing the Security World across several client computers, you must ensure that the changes to the host data are propagated to all your computers. One way to achieve this is to use client cooperation. For more information, see Setting up client cooperation.

9.6.1. Replacing OCSs from the unit front panel

To replace an OCS from the unit front panel, follow these steps:

1. From the main menu, select **Security World mgmt** > **Admin operations** > **Recover keys**.

2. Select **all** to recover all keys in the Security World, or select the application for which you want to recover the keys.

3. If you selected an application, select the keys that you want to recover.

4. Insert the required number of Administrator Cards to recover keys, and enter their passphrases if required.

5. Insert the required number of Operator Cards, and enter their passphrases if required.

   When you have inserted the required number of cards, details of the recovered key are displayed.

6. Check the key details are correct and then scroll down and select **Recover key**.

   You can also select **More info** to see more information about the keys.

A message is displayed when the keys are recovered.

9.6.2. Replacing OCSs with KeySafe

In order to replace an OCS, you must have another OCS onto which to copy the first set's data. If you do not already have an existing second OCS, you must create a new one. For more information, see Creating Operator Card Sets (OCSs).

When you have a second OCS ready, follow these steps in order to replace the first OCS:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe.)

2. Click the **Card Sets** menu button, or select **Card Sets** from the **Manage** menu. KeySafe takes you to the **List Operator Card Sets** panel.

3. Click **Replace card set**. KeySafe takes you to the **Replace card set** panel.

This panel lists existing OCSs in tabular form. For each card set it displays:

Attribute	Description
Name	The name of the card set.
Required (K)	The number of cards needed to re-create a key.
Total (N)	The total number of cards in the set.
Persistent	Indicates whether or not the card set is persistent.
Timeout	The timeout value of the card, in seconds
Recoverable Key Count	The number of private keys protected by this card set that are recoverable.
Nonrecoverable Key Count	The number of private keys protected by this card set that are not recoverable.

You can click and drag with your mouse in order to resize the column widths and to rearrange the column order of this table. Clicking a column heading sorts the rows in ascending order based on that column heading.

4. Select an OCS that you want to replace, and click **Replace card set**.

> ℹ️ If an OCS does not have any recoverable keys, it cannot be replaced.

5. KeySafe takes you to the **Load Administrator Card Set** panel, where it prompts you to insert cards from the ACS in order to authorize the action. Each time you insert an Administrator Card into the smart card of the hardware security module slot, you must click the **OK** button to load the card.

> ℹ️ Only insert your ACS into a module that is connected to a trusted server.

6. When you have loaded enough cards from the ACS to authorize the procedure, KeySafe takes you to the **Load Operator Card Set** panel, where it prompts you to insert the OCS that is to protect the recoverable keys (this is the OCS onto which you are copying data from the OCS you are replacing). Each time you insert a card from the new OCS into the smart card slot of the hardware security module, you must click the **OK** button.

When you have loaded enough cards from the new OCS, KeySafe creates new working versions of the recoverable keys that are protected by this card set.

KeySafe deletes the original host data for all recovered keys and replaces this data with host data that is protected by the new OCS. If there are no nonrecoverable keys pro-

tected by the card set, KeySafe also removes the old card set from the Security World. However, if the OCS has nonrecoverable keys, the host data for the original card set and for the nonrecoverable keys is not deleted. These keys can only be accessed with the original OCS. If you want to delete these files, use the **Remove OCS** option.

7. When the process is complete, KeySafe displays a dialog indicating that the OCS has been successfully replaced. Click the **OK** button. KeySafe returns you the Replace Operator Card Set panel, where you may replace another OCS or choose a different operation.

9.6.3. Replacing OCSs or softcards with rocs

You can use the `rocs` command-line utility interactively, or you can supply all the parameters using the command line.

9.6.3.1. Using rocs interactively

To use the `rocs` command-line utility interactively, run it without any parameters:

```
rocs
```

`rocs` displays the following prompt:

```
'rocs' key recovery tool
Useful commands: 'help', 'help intro', 'quit'.
rocs >
```

In order to use `rocs` to replace an OCS or recover keys to a softcard, take the following steps:

1. You must select a hardware security module to use by using the `module` command, which is described in the section module <number>.
2. List the OCSs and softcards in the current Security World by using the `list cardsets` command, which is described in the section list cardsets.
3. Select the OCS or softcard to which you want to transfer the keys by using the `target` command, which is described in the section target <cardset-spec>.

> Keys protected by an OCS can only be recovered to another OCS, and not to a softcard. Likewise, softcard-protected keys can only be recovered to another softcard, and not to an OCS.

4. List the keys in the current Security World using the `list keys` command, which is described in the section list keys.

5. Select the keys that are to be recovered (from a different OCS or softcard than the one you selected for key transfer) by using the `mark` command, which is described in the section mark <key-spec>.

6. If you have selected any keys by mistake, deselect them by using the `unmark` command, which is described in the section unmark <key-spec>.

7. After you have selected the keys that are to be recovered, transfer these keys by using the `recover` command, which is described in the section recover.

   `rocs` prompts you to insert a card from the ACS.

8. Insert a card from the ACS.

   `rocs` prompts you for the passphrase for this card. This action is repeated until you have loaded the required number of cards from the ACS.

   If you do not have the required number of cards from the ACS, press **Q** and then **Enter**. The `rocs` utility returns you to the `rocs >` prompt without processing any keys.

   > ❗ Only insert Administrator Cards into a hardware security module that is connected to a trusted server.

9. If you are recovering keys to an OCS:

   a. `rocs` prompts you to insert a card from the first OCS that you have selected as the target. OCSs are processed in ascending numerical order as listed by the `list cardsets` command.

   b. Insert a card from this OCS.

   c. `rocs` prompts you for the passphrase for this card. This action is repeated until you have loaded the required number of cards from the OCS.

   If you are recovering keys to a softcard, `rocs` prompts you for the passphrase for the softcard that you have selected as the target.

   If you decide that you do not want to transfer the keys to the selected card set or softcard, press **Q** and then **Enter** to quit. `rocs` returns you to the `rocs >` prompt and does not process any further OCSs or softcards.

   When you have loaded the target softcard or the required number of cards from the target OCS, `rocs` transfers the selected keys to the target OCS or softcard.

   If you have selected other target OCSs or softcards, `rocs` prompts for a card from the next OCS.

10. Repeat step 9 for each selected target.

11. If you have transferred the correct keys, write the key blobs to disk by using the `save` function (described in the section `save <key-spec>`). If you have transferred a key by mistake, you can restore it to its original protection by using the `revert` command (described in the section `revert <key-spec>`).

At the `rocs` prompt, you can use the following commands:

- `help <topic>`
- `help intro`
- `list cardsets`
- `list keys`
- `mark <key-spec>`
- `module <number>`
- `quit`
- `recover`
- `rescan`
- `revert <key-spec>`
- `save <key-spec>`
- `status`
- `target <cardset-spec>`
- `unmark <key-spec>`

> ℹ️ You can specify a command by typing enough characters to identify the command uniquely. For example, for the `status` command, you can type `st` and then press **Enter**.

9.6.3.1.1. help

With no arguments specified, `help` shows a list of available commands with brief usage messages and a list of other help topics. With an argument, `help` shows detailed help information about a given topic.

`help intro` displays a brief step-by-step guide to using `rocs`.

9.6.3.1.2. list cardsets

This command lists the OCSs and softcards in the current Security World.

For example:

```
No.
Name                 Keys (recov) Sharing
 1 test                  6 (6)       3 of 5; 20 minute timeout
 2 test2                 3 (2)       2 of 3
 3 test3                 1 (1)       1 of 1; persistent
```

In this output:

Output	Description
No.	The card set or softcard number, which you can use to identify this card set in rocs commands.
Name	The OCS or softcard name.
Keys	The number of keys protected by this OCS or softcard.
(recov)	The number of keys protected by this OCS or softcard.
Sharing	The K of N parameters for this OCS.
persistent	The OCS is persistent and does not have a time-out set.
### minute timeout	The OCS is persistent and has a time-out set.

9.6.3.1.3. list keys

This command lists the keys in the current Security World, as in the following example:

```
No.
Name                  App       Protected by
 1 rsa-test           hwcrhk    module
 2 Id: uc63e0ca3cb032d71c1c pkcs11    test2
R 3 Server-Cert              pkcs11     test --> test2
 4 Id: uc63e0ca3cb032d71c1c pkcs11    test --> test3
 5 Server-Cert              pkcs11    module (test ---> fred2)
```

In this output:

Output	Description
No.	The key number, which you can use in mark and unmark commands.
Name	The key name.
App	The application with which the key is associated.
Protected by	This indicates the protection method (see table below).

In this output, the protection methods include:

Method	Description
module	Key protected by the Security World.
name	Key protected by the named OCS or softcard.
name-->name2	Key protected by the OCS or softcard *name1* marked for recovery to OCS or softcard *name2*.
module (name)	PKCS #11 public object. These are protected by the Security World but associated with a specific OCS or softcard.
module (*name-->name2*)	PKCS #11 public object marked for recovery.

9.6.3.1.4. mark <key-spec>

This command marks the listed keys that are to be recovered to the target OCS or softcard. You can mark one or more keys by number, *ident*, OCS or softcard, or hash. For more information, see Specifying keys.

To mark more than one key at a time, ensure that each *key-spec* is separated from the other by spaces, as in the following example:

```
mark key-spec1 key-spec2 key-spec3
```

If you have not selected a target OCS or softcard, or if `rocs` cannot parse the *key-spec*, then `rocs` displays an error message.

You can mark and remark the keys to be recovered to various target OCSs or softcards. Remarking a key displaces the first target in favor of the second target.

> ℹ️ Keys protected by an OCS can only be recovered to another OCS, and not to a softcard. Likewise, softcard-protected keys can only be recovered to another softcard, and not to an OCS.

9.6.3.1.5. module <number>

This command selects the hardware security module to be used. The module number must correspond to a hardware security module in the current Security World. If the hardware security module does not exist, is not in the Security World, or is otherwise unusable, then `rocs` displays an error message and does not change to the selected module.

9.6.3.1.6. quit

This command allows you to leave `rocs`. If you attempt to `quit` when you have recovered

---

keys but have not saved them, `rocs` displays a warning.

9.6.3.1.7. recover

This command transfers the marked keys to their target OCSs or softcards. This operation is not permanent until you save these keys by using the `save` command.

9.6.3.1.8. rescan

This command updates the card set and key information.

9.6.3.1.9. revert <key-spec>

This command returns keys that have been recovered, but not saved, to being protected by the original protection method. If the selected keys have not been recovered, `rocs` displays an error message.

9.6.3.1.10. save <key-spec>

This command writes the new key blobs to disk. If you specify a key-spec, only those keys are saved. Otherwise, all recovered keys are saved.

9.6.3.1.11. status

This command lists the currently selected hardware security module and target OCS or soft card.

9.6.3.1.12. target <cardset-spec>

This command selects a given OCS or softcard as the target. You can specify the card set or softcard name, the number returned by `list cardsets`, or the hash.

9.6.3.1.13. unmark <key-spec>

This command unmarks the listed keys. Unmarked keys are not recovered.

9.6.3.2. Using rocs from the command line

You can select all the options for `rocs` using the command line by running a command of the form:

```
rocs -m|--module=<MODULE> [-t|--target=<CARDSET-SPEC>] [-k|--keys=<KEYS-SPEC>] [-c|--cardset=<CARDSET-SPEC>] [-
i|--interactive]
```

In this command:

Option	Description
-m|--module=<MODULE>	These options specify the number of the hardware security module to use.
-t|--target=<CARDSET-SPEC>	These options specify the OCS or softcard to be used to protect the keys. For more information, see Specifying card sets.
-k|--keys=<KEYS-SPEC>	These options select the keys to be recovered. For more information, see Specifying keys.
-c|--cardset=<CARDSET-SPEC>	These options select all keys that are protected by the given OCS or softcard. For more information, see Specifying card sets.
-i|--interactive	These options force rocs to start interactively even if you have already selected keys.

You must specify the target before you specify keys.

You can use multiple --keys=<KEYS-SPEC> and --cardset=<CARDSET-SPEC> options, if neces-
sary.

You can specify multiple targets on one command line by including separate --keys=<KEYS
-SPEC> or --cardset=<CARDSET-SPEC> options for each target. If a key is defined by
--keys=<KEYS-SPEC> or --cardset=<CARDSET-SPEC> options for more than one target, it is
transferred to the last target for which it is defined.

If you have selected a hardware security module, a target OCS or softcard, and keys to
recover but have not specified the --interactive option, rocs automatically recovers the
keys. rocs prompts you for the ACS and OCS or softcard. For more information, see Using
rocs interactively.

> If you use rocs from the command line, all keys are recovered and saved
> automatically. You cannot revert the keys unless you still have cards
> from the original OCS.

If you do not specify the target and keys to recover, or if you specify the --interactive
option, rocs starts in interactive mode with the selections you have made. You can then use
further rocs commands to modify your selection before using the recover and save com-
mands to transfer the keys.

9.6.3.3. Specifying card sets

The value of `<CARDSET-SPEC>` identifies one or more OCSs or softcards. It may have any of the following forms:

Value	Description
`[number]` *cardset-number*	A value of this form selects the OCS or softcard with the given number from the list produced by the `list cardsets` command.
`[name]` *cardset-name*	A value of this form selects card sets or softcards by their names (the card set or softcard name may be a wildcard pattern in order to select all matching OCSs or softcards).
`hash` *cardset-hash*	A value of this form selects the OCS or softcard with the given hash.

In order to specify multiple OCSs or softcards, include several *CARDSET-SPEC*'s using the command line.

> ℹ Keys protected by an OCS can only be recovered to another OCS, and not to a softcard. Likewise, softcard-protected keys can only be recovered to another softcard, and not to an OCS.

9.6.3.4. Specifying keys

The `--keys=<KEYS-SPEC>` option identifies one or more keys. It may have any of the following forms:

Value	Description

mark `key-number`	A value of this form selects the key with the given number from the list produced by the list keys command. Examples of usage are:
	
	`rocs -t <target_OCS> -k <key_number>`
	
	and
	
	`rocs -t <target_OCS> -k "mark 56"`

appname_:`keyident`	A value of this form selects keys by their internal application name and `ident`. You must supply at least one of `appname` or `keyident`, but you can use wildcard patterns for either or both in order to select all matching keys. An example of usage is: `rocs -t <target_OCS> --keys="simple:simplekey"`

Value	Description
hash *keyhash*	A value of this form selects the key with the given key hash. An example of usage is: ``` rocs -t <target_OCS> --keys="hash e364[...]" ```
--cardset *cardset-spec*	A value of this form selects all keys protected by a given card set.

9.7. Replacing the Administrator Card Set

Replacing the ACS requires a quorum of cards from the current ACS (*K/N*) to perform the following sequence of tasks:

1. loading the secret information that is to be used to protect the archived copy of the Security World key.
2. creating a new secret that is to be shared between a new set of cards.
3. creating a new archive that is to be protected by this secret.

If you discover that one of the cards in the current ACS has been damaged or lost, or you want to migrate from standard nShield cards to nShield Remote Administration Cards, you should use one of the following to create a new set:

- The `racs` utility.

> When using the `racs` utility, you cannot redefine the quantities in a K of N relationship for an ACS. The K of N relationship defined in the original ACS persists in the new ACS.

- The front panel of the nShield HSM.
- The KeySafe **Replace Administrator Card Set** option.

  Accessed from the `Card Operations` panel.

> If further cards are damaged, you may not be able to re-create your Security World.

> You cannot mix nShield cards with nShield Remote Administration Cards in the same set.

> Replacing the ACS modifies the `world` file. In order to use the new ACS on other machines in the Security World, you must copy the updated `world` file to all the machines in the Security World after replacing the

> ACS. Failure to do so could result in loss of administrative access to the Security World.

(!) We recommend that you erase your old Administrator Cards as soon as you have created the new ACS. An attacker with the old ACS and a copy of the old host data could still re-create all your keys. With a copy of a current backup, they could even access keys that were created after you replaced the ACS.

(i) Before you start to replace an ACS, you must ensure that you have enough blank cards to create a complete new ACS. If you start the procedure without enough cards, you will have to cancel the procedure part way through.

9.7.1. Replacing an Administrator Card Set using the nShield HSM front panel

To replace an ACS:

1. From the main menu, select **Security World mgmt** > **Admin operations** > **Replace ACS**.

2. Insert one of the remaining cards from the card set that you want to replace and press the right-hand navigation button.

   Continue to insert cards until you have inserted the number of cards required to authorize the process.

3. When prompted, insert a card for the replacement card set and press the right-hand navigation button.

4. If required, specify a passphrase for the card.

5. Insert cards until the card set is complete. A message confirms that the card set has been created.

6. At this point the modified **world** file has been pushed to the RFS, so make a backup of the modified **world** file on the RFS, preferably in a way that distinguishes it from previous backups.

7. Copy the **world** file to any other HSMs in the same Security World, either using the **Security World mgmt** > **RFS operations** > **Update World files** option on the HSM concerned or using the `nethsmadmin` utility, see Using nethsmadmin to copy a Security World to a nShield HSM and check the current version.

8. If client cooperation is not enabled, copy the modified `world` file onto any client

machines where it is needed.

9. Check that the new Administrator Cards are usable and that their passphrases have been set as intended, see Verifying the passphrase of a card or softcard

10. Erase the Administrator Cards from the old card set. For more information, see Erasing cards and softcards.

9.7.2. Replacing an ACS with KeySafe

When you have enough cards to create a complete new ACS ready and a quorum of the ACS you want to replace, follow these steps:

1. Start KeySafe. (For an introduction to KeySafe and information on starting the software, see Using KeySafe).

2. Click the **Card sets** menu button, or select **Card sets** from the **Manage** menu. KeySafe takes you to the **List Operator Card Sets** panel.

3. Click the **Replace ACS** navigation button, and KeySafe takes you to the **Replace Administrator Card Set** panel.

4. If you are sure that you want to replace the ACS, click the **Replace ACS** command button

5. KeySafe takes you to the **Load Administrator Card Set** panel, where it prompts you to insert cards from the ACS in order to authorize the action. Each time you insert an Administrator Card into the module's smart card slot, you must click the **OK** button to load the card.

> **ℹ** Only insert cards from your ACS into a module that is connected to a trusted server.

6. When you have loaded enough Administrator Cards to authorize the action, KeySafe takes you to the Create Administrator Card Set panel, where it prompts you to insert the cards that are to form the ACS. These must be blank cards or cards that KeySafe can erase. KeySafe will not let you use cards from the existing ACS. If you do not have enough cards to form a complete new ACS, cancel the operation now.

> **ℹ** When creating a card set, KeySafe recognizes cards that belongs to the set even before the card set is complete. If you accidentally insert a card to be written again after it has already been written, KeySafe displays a warning.

7. When you insert a blank card, KeySafe takes you to the **Set Card Protection passphrase** panel.

8. If you want to set a passphrase for this Administrator Card:

   a. Select the **Yes** option.

   b. Enter the same passphrase in both text fields.

   c. Click the **OK** button.

   KeySafe then prompts you for the next card (if any). A given passphrase is associated with a specific card, so each card can have a different passphrase. You can change these passphrases at any time by using the KeySafe **Examine/Change Card** option (available from the **List Operator Card Sets** panel) or the `cardpp` command-line utility.

9. If you do not want to set a passphrase for this Administrator Card:

   a. Select the **No** option.

   b. Click the **OK** button.

10. After you have created all the Administrator Cards, KeySafe displays a message confirming that the ACS has been successfully replaced.

11. Click the **OK** button, and KeySafe returns you to its introduction panel.

When you have finished replacing the ACS:

1. If you ran KeySafe on a client machine, ensure that there is a copy of the modified **world** file on the RFS.

2. Make a backup of the **world** file, preferably in a way that distinguishes it from previous backups.

3. Copy the **world** file to any other HSMs in the same Security World, for example using the `nethsmadmin` utility, see Using nethsmadmin to copy a Security World to an nShield HSM and check the current version.

4. If client cooperation is not enabled, copy the modified **world** file onto any other client machines where it is needed.

5. Check that the new Administrator Cards are usable and that their passphrases have been set as intended, see Verifying the passphrase of a card or softcard.

6. Erase the old Administrator Cards; for more information, see Erasing cards and softcards.

9.7.3. Replacing an Administrator Card Set using racs

The `racs` utility creates a new ACS to replace a set that was created on the module with the new-world utility.

When using the `racs` utility, you cannot redefine the quantities in a K of

> N relationship for an ACS. The K of N relationship defined in the original ACS persists in the new ACS.

1. Ensure the hardware security module is in operational mode.

2. Run a command of the form:

```
racs [-m|--module=<MODULE>]
```

In this command, the `-m|--module`=<MODULE>` option specifies the `ModuleID` (`<MODULE>`) of the module to use.

3. When prompted, insert the appropriate quorum of Administrator Cards to authorize the replacement.

4. When prompted that `racs` is writing the new ACS, insert blank cards as necessary on which to write the replacement Administrator Cards.

5. If you ran `racs` on a client machine, ensure that there is a copy of the modified **world** file on the RFS.

6. Make a backup of the **world** file, preferably in a way that distinguishes it from previous backups.

7. Copy the **world** file to any other HSMs in the same Security World, for example using the `nethsmadmin` utility, see Using nethsmadmin to copy a Security World to an nShield HSM and check the current version.

8. If client cooperation is not enabled, copy the modified **world** file onto any other client machines where it is needed.

9. Check that the new Administrator Cards are usable and that their passphrases have been set as intended, see Verifying the passphrase of a card or softcard.

10. Erase the old Administrator Cards; for more information, see Erasing cards and softcards.

10. Application interfaces

This chapter explains how to use an HSM with various types of application:

- nCipherKM JCA/JCE CSP
- PKCS #11 applications
- nShield native and Custom applications
- Microsoft CAPI CSP
- Microsoft CNG CSP.

> **i** For information about using the Microsoft Cryptographic API, see the appropriate third-party integration guide at: https://www.entrust.com/documentation.

You can use KeySafe or the `generatekey` utility to generate or import keys for use with your applications (see Working with keys). By default, KeySafe uses the same mechanisms and supports the same applications as the `generatekey` utility.

> **i** On **Linux**, you must add the user of any application that uses an nShield HSM to the group **nfast** before the application runs. On **Windows**, by default any user is allowed to use any application that uses an nShield HSM.

> **i** If you create keys on a client that is not on the same computer as the RFS, you must copy the key data to the RFS before the nShield HSM can use these keys.

10.1. nCipherKM JCA/JCE CSP

The nCipherKM JCA/JCE CSP (Cryptographic Service Provider) allows Java applications and services to access the secure cryptographic operations and key management provided by Entrust hardware. This provider is used with the standard JCE (Java Cryptographic Extension) programming interface.

To use the nCipherKM JCA/JCE CSP, you must install:

- the nShield Java package which includes the nShield Java jars and KeySafe.

For more information about the bundles and components supplied on your Security World Software installation media, see the User Guide.

The following versions of Java have been tested to work with, and are supported by, your

nShield Security World Software:

- Java 8 (or Java 1.8x)
- Java 11
- Java 17

We recommend that you ensure Java is installed before you install the Security World Software. The Java executable must be on your system path.

If you can do so, please use the latest Java version currently supported by Entrust that is compatible with your requirements. Java versions before those shown are no longer supported. If you are maintaining older Java versions for legacy reasons, and need compatibility with current nShield software, please contact Entrust nShield Technical Support, https://nshieldsupport.entrust.com.

To install Java you may need installation packages specific to your operating system, which may depend on other pre-installed packages to be able to work.

Suggested links from which you may download Java software as appropriate for your operating system:

- http://www.oracle.com/technetwork/java/index.html
- http://www.oracle.com/technetwork/java/all-142825.html

> 🛈 Detailed documentation for the JCE interface can be found on the Oracle Technology web page https://docs.oracle.com/en/java/javase/11/security/java-cryptography-architecture-jca-reference-guide.html.

> 🛈 Softcards are not supported for use with the nCipherKM JCA/JCE CSP in Security Worlds that are compliant with FIPS 140 Level 3.

10.1.1. Installing the nCipherKM JCA/JCE CSP

To install the nCipherKM JCA/JCE CSP:

1. In the hardserver configuration file, ensure that:
   - `priv_port` (the port on which the hardserver listens for local privileged TCP connections) is set to 9001
   - `nonpriv_port` (the port on which the hardserver listens for local nonprivileged TCP connections) is set to 9000.

     If you need to change either or both of these port settings, you restart the hardserver before continuing the nCipherKM JCA/JCE CSP installation process. For

more information, see Stopping and restarting the hardserver.

2. For Java 8 only. Copy the `nCipherKM.jar` file to the extensions folder of your local Java Virtual Machine installation from `/opt/nfast/java/classes` (**Linux**) or `%NFAST_HOME%\java\classes` (**Windows**).

The location of the extensions folder depends on the type of your local Java Virtual Machine (JVM) installation:

JVM type	Extensions folder (Linux)	Extensions folder (Windows)
Java Developer Kit (JDK)	`$JAVA_HOME/jre/lib/ext`	`%JAVA_HOME%\jre\lib\ext`
Java Runtime Environment (JRE)	`$JAVA_HOME/lib/ext`	`%JAVA_HOME%\lib\ext`

In these paths, `$JAVA_HOME` (**Linux**) and `%JAVA_HOME%` (**Windows**) are the home directory of the Java installation (commonly specified in the `JAVA_HOME` environment variable). If you are using Java 11 or Java 17 you do not need to copy the jar file.

3. Add `$JAVA_HOME/bin` (**Linux**) or `%JAVA_HOME%\bin` (**Windows**) to your PATH system variable

4. For Java 8 only. Install the unlimited strength JCE jurisdiction policy files that are appropriate to your version of Java. JDK 9 and later ship with, and use by default, the unlimited policy files.

The Java Virtual Machine imposes limits on the cryptographic strength that may be used by default with JCE providers. Replace the default policy configuration files with the unlimited strength policy files.

To install the unlimited strength JCE jurisdiction policy files:

a. If necessary, download the archive containing the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files from the Web site of your Java Virtual Machine vendor.

> The Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files are covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. We recommend that you take legal advice before downloading these files from your Java Virtual Machine vendor.

b. Extract the files `local_policy.jar` and `US_export_policy.jar` from Java Virtual Machine vendor's Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy File archive.

c. Copy the extracted files `local_policy.jar` and `US_export_policy.jar` into the security directory for your local Java Virtual Machine (JVM) installation:

JVM type	Extensions folder (Linux)	Extensions folder (Windows)
Java Developer Kit (JDK)	`$JAVA_HOME/jre/lib/security`	`%JAVA_HOME%\jre\lib\security`
Java Runtime Environment (JRE)	`$JAVA_HOME/lib/security`	`%JAVA_HOME%\lib\security`

In these paths, `$JAVA_HOME` (**Linux**) and `%JAVA_HOME%` (**Windows**) are the home directory of the Java installation (commonly specified in the `JAVA_HOME` environment variable).

> **ⓘ** Copying the files `local_policy.jar` and `US_export_policy.jar` into the appropriate folder must overwrite any existing files with the same names.

5. Add the nCipherKM provider to the `java.security` file located in the security directory for your local Java Virtual Machine (JVM) installation: `security.provider.<n>=com.ncipher.provider.km.nCipherKM`, where `<n>` is the position in the list of providers, for example:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=sun.security.rsa.SunRsaSign
security.provider.3=com.sun.net.ssl.internal.ssl.Provider
security.provider.4=com.sun.crypto.provider.SunJCE
security.provider.5=sun.security.jgss.SunProvider
security.provider.6=com.sun.security.sasl.Provider
security.provider.7=com.ncipher.provider.km.nCipherKM
```

For Java 11 and Java 17 you do not need to specify the fully qualified class name for the provider. Instead you can just use the provider name: `security.provider.<n>=nCipherKM`.

The JVM uses this file to select the provider from which to request a mechanism instance. If your JCE application does not request the nCipherKM provider by name, or if it fails to load keys, you might need to move the nCipherKM provider to the top of the list: `security.provider.1=com.ncipher.provider.km.nCipherKM`. Do not change the relative order of the other providers in the list.

> **ⓘ** Ensure you do not list multiple providers with the same number (for example, ensure your list of providers does not include two instances of `security.provider.1`, both `com.ncipher.provider.km.nCipherKM` and another provider). If you add the

> nCipherKM provider as `security.provider.1`, ensure that the subsequent providers are re-numbered correctly.

6. Save your updates to the file `java.security`.

   When you have installed the nCipherKM JCA/JCE CSP, you must have created a Security World before you can test or use it. For more information about creating a Security World, see Creating a Security World.

   > **ℹ** If you have a Java Enterprise Edition Application Server running, you must restart it before the installed nCipherKM provider is loaded into the Application Server virtual machine and ready for use.

10.1.1.1. Testing the nCipherKM JCA/JCE CSP installation

After installation, you can test that the nCipherKM JCA/JCE CSP is functioning correctly by running the command.

For Java 8:

```
java com.ncipher.provider.InstallationTest
```

For Java 11 and Java 17:

Linux

```
java --module-path /opt/nfast/java/classes com.ncipher.provider.InstallationTest
```

Windows

```
java --module-path %NFAST_HOME%\java\classes com.ncipher.provider.InstallationTest
```

> **ℹ** For this command to work, you must have added `$JAVA_HOME` (**Linux**) or `%JAVA_HOME%` (**Windows**) to your `PATH` system variable.

If the nCipherKM JCA/JCE CSP is functioning correctly, output from this command has the following form:

```
Installed providers:
1: nCipherKM
2: SUN
3: SunRsaSign
4: SunJSSE
5: SunJCE
6: SunJGSS
```

```
7: SunSASL
Unlimited strength jurisdiction files are installed.
The nCipher provider is correctly installed.
nCipher JCE services:
Alg.Alias.Cipher.1.2.840.113549.1.1.1
Alg.Alias.Cipher.1.2.840.113549.3.4
Alg.Alias.Cipher.AES
Alg.Alias.Cipher.DES3
```

If the `nCipherKM` provider is installed but is not registered at the top of the providers list in the `java.security` file, the `InstallationTest` command produces output that includes the message:

```
The nCipher provider is installed, but is not registered at the top of the providers list in the java.security
file.
See the user guide for more information about the recommended system configuration.
```

In such a case, edit the `java.security` file (located in the security directory for your local JVM installation) so that the nCipherKM provider is registered in the first position in that file's list of providers. For more information about the `java.security` file, see Installing the nCipherKM JCA/JCE CSP.

If the nCipherKM provider is not installed at all, or you have not created a Security World, or if you have not configured ports correctly in the hardserver configuration file, the `InstallationTest` command produces output that includes the message:

```
The nCipher provider is not correctly installed.
```

In such case:

- Check that you have configured ports correctly, as described in Installing the nCipherKM JCA/JCE CSP. For more information about hardserver configuration file settings, see server_startup.
- Check that you have created a Security World. If you have not created a Security World, create a Security World. For more information, see Creating a Security World.
- If you have already created a Security World, repeat the nCipherKM JCA/JCE CSP installation process as described in Installing the nCipherKM JCA/JCE CSP.

After making any changes to the nCipherKM JCA/JCE CSP installation, run the `InstallationTest` command again and check the output.

Whether or not the nCipherKM provider is correctly installed, if the unlimited strength jurisdiction files are not installed or (not correctly installed), the `InstallationTest` command produces output that includes the message:

```
Unlimited strength jurisdiction files are NOT installed.
```

This message means that, because the Java Virtual Machine imposes limits on the cryptographic strength that you can use by default with JCE providers, you must replace the default policy configuration files with the unlimited strength policy files. For information about how to install the unlimited strength jurisdiction files, see Installing the nCipherKM JCA/JCE CSP.

10.1.2. Named Modules in Java 11 and Java 17

The nCipherKM Provider has been implemented as a named module. This means that, for Java 11 and Java 17, if you have added the provider to your `java.security` file, then you can run your application with the `nCipherKM.jar` on the module-path and the Java Service-Loader class will automatically find it. For example:

Linux

```
java --module-path /opt/nfast/java/classes com.ncipher.provider.InstallationTest
```

Windows

```
java --module-path %NFAST_HOME%\java\classes com.ncipher.provider.InstallationTest
```

Alternatively, you can specify the location of the nCipherKM jar on the classpath:

Linux

```
java --class-path /opt/nfast/java/classes/nCipherKM.jar com.ncipher.provider.InstallationTest
```

Windows

```
java --class-path %NFAST_HOME%\java\classes\nCipherKM.jar com.ncipher.provider.InstallationTest
```

10.1.3. keytool

Use the Java `keytool` utility to read and edit an nShield KeyStore. You must specify the correct `nCipher.sworld` KeyStore type when you run the `keytool` utility.

To generate a new key in an OCS-protected KeyStore with the Java `keytool` utility, run the following command:

```
keytool -genkeypair -storetype nCipher.sworld -keyalg RSA -sigalg SHA1withRSA -storepass <KeyStore_passphrase>
```

```
-keystore <KeyStore_path>
```

In this example command, `KeyStore_passphrase` is the passphrase for the OCS protecting the KeyStore. `KeyStore_path` is the path to the KeyStore.

To generate a new key in a module-protected KeyStore with the Java `keytool` utility, run the following command:

```
keytool -J-Dprotect=module -J-DignorePassphrase=true -genkeypair -storetype nCipher.sworld -keyalg RSA -sigalg
SHA1withRSA -keystore <KeyStore_path>
```

In this example command, `<KeyStore_path>` is the path to the KeyStore.

By default, the `keytool` utilities use the `MD5withRSA` signature algorithm to sign certificates used with a KeyStore. This signature mechanism is unavailable on modules with firmware version 2.33.60 or later.

10.1.4. Using keys

Only the nCipherKM provider can use keys stored in an nShield KeyStore because the under lying key material is held separately in the Security World.

You can always store nShield keys in an nShield KeyStore. You can also store keys gener-ated by a third-party provider into an nShield KeyStore if both of the following conditions apply:

- the key type is known to the nCipherKM provider
- the Security World is *not* compliant with FIPS 140 Level 3.

When you generate an nShield key (or create it from imported key material), that key is asso ciated with an ACL (Access Control List). This ACL prevents the key from being used for operations for which it is unsuited and enforces requirements that certain tokens be pre-sented; for example, the ACL can specify that signing key cannot be used for encryption.

10.1.5. System properties

You can use system properties to control the provider. You set system properties when starting the Java Virtual Machine using a command such as:

```
java -D<property>=<value> <MyJavaApplication>
```

In this example command, `<property>` represents any system property, `<value>` represents

---

the value set for that property, and `<MyJavaApplication>` is the name of the Java application you are starting. You can set multiple system properties in a single command, for example:

```
java -Dprotect=module -DignorePassphrase=true <MyJavaApplication>
```

The available system properties and their functions as controlled by setting different values for a property are described in the following table:

Property	Function for different values
`JCECSP_DEBUG`	This property is a bit mask for which different values specify different debugging functions; the default value is `0`. For details about the effects of setting different values for this property, see JCECSP_DEBUG property values.
`JCECSP_DEBUGFILE`	This property specifies a path to the file to which logging output is to be written. Set this property if the `JCECSP_DEBUG` property is set to a value other than the default of `0`. For details about the effects of setting different values for this property, see JCECSP_DEBUG property values. In a production environment, we recommend that you disable debug logging to prevent sensitive information being made available to an attacker.
`protect`	This property specifies the type of protection to be used for key generation and nCipherKM KeyStore instances. You can set the value of this property to one of `module`, `<SOFTCARD_NAME:SOFTCARD_IDENT>`, or `cardset`. OCS protection (`cardset`) uses the card from the first slot of the first usable hardware security module. To find the logical token hash `<IDENT>` of a softcard, run the command `nfkminfo --softcard-list`.
`module`	This property lets you override the default HSM and select a specific HSM to use for HSM and OCS protection. Set the value of this property as the ESN of the HSM you want to use.
`slot`	This property lets you override the default slot for OCS-protection and select a specific slot to use. Set this the value of this property as the number of the slot you want to use.
`ignorePassphrase`	If the value of this property is set to `true`, the nCipherKM provider ignores the passphrase provided in its KeyStore implementation. This feature is included to allow the Oracle or IBM `keytool` utilities to be used with module-protected keys. The `keytool` utilities require a passphrase be provided; setting this property allows a dummy passphrase to be used.
`seeintegname`	Setting the value of this property to the name of an SEE integrity key causes the provider to generate SEE application keys. These keys may only be used by an SEE application signed with the named key.

Property	Function for different values
com.nci-pher.provider.announcemode	The default value for this property is `auto`, which uses firmware auto-detection to disable algorithms in the provider that cannot be supported across all installed HSMs. Setting the value of this property to `on` forces the provider to advertise all mechanisms at start-up. Setting the value of this property to `off` forces the provider to advertise no mechanisms at start-up.
com.ncipher.provider.enable	For the value of this property, you supply a comma-separated list of mechanism names that are to be forced on, regardless of the announce mode selected.
com.ncipher.provider.dis-able	For the value of this property, you supply a comma-separated list of mechanism names that are to be forced off, regardless of the announce mode selected. Any mechanism supplied in the value for the `com.nci-pher.provider.disable` property overrides the same mechanism if it is supplied in the value for the `com.ncipher.provider.enable` property.

10.1.5.1. JCECSP_DEBUG property values

The `JCECSP_DEBUG` system property is a bit mask for which you can set different values to control the debugging functions. The following table describes the effects of different values that you can set for this property:

JCECSP_DEBUG value	Function
0	If this property has no bits set, no debugging information is reported. This is the default setting.
1	If this property has the bit 1 set, minimal debugging information (for example, version information and critical errors) is reported.
2	If this property has the bit 2 set, comprehensive debugging information is reported.
4	If this property has the bit 3 set, debugging information relating to creation and destruction of memory and HSM resources is reported.
8	If this property has the bit 4 set, `debugFunc` and `debugFuncEnd` generate debugging information for functions that call them.
16	If this property has the bit 5 set, `debugFunc` and `debugFuncEnd` display the values for all the arguments that are passed in to them.
32	If this property has the bit 6 set, context information is reported with each debugging message (for example, the `ThreadID` and the current time).
64	If this property has the bit 7 set, the time elapsed during each logged function is calculated, and information on the number of times a function is called and by which function it was called is reported.

`JCECSP_DEBUG` value	Function
128	If this property has the bit 8 set, debugging information for NFJAVA is reported in the debugging file.
256	If this property has the bit 9 set, the call stack is printed for every debug message.

To set multiple logging functions, add up the `JCECSP_DEBUG` values for the debugging functions you want to set, and specify the total as the value for `JCECSP_DEBUG`. For example, if you want to set the debugging to use both function tracing (bit 4) and function tracing with parameters (bit 5), add the `JCECSP_DEBUG` values shown in the table for these debugging functions (8 + 16 = 24) and specify this total (24) as the value to use for `JCECSP_DEBUG`.

10.1.6. Compatibility

The nCipherKM JCA/JCE CSP supports both module-protected keys and OCS-protected keys. The CSP currently supports 1/*N* OCSs and a single protection type for each nCipherKM JCE KeyStore.

You can use the nCipherKM JCA/JCE CSP with Security Worlds that comply with FIPS 140 at either Level 2 or Level 3.

> **ⓘ** In a Security World that complies with FIPS 140 Level 3, it is not possible to import keys generated by other JCE providers.

The nCipherKM JCA/JCE CSP supports load-sharing for keys that are stored in the nCipherKM KeyStore. This feature allows a server to spread the load of cryptographic operations across multiple connected HSMs, providing greater scalability.

> **ⓘ** We recommend that you use load-sharing unless you have existing code that is designed to run with multiple HSMs. To share keys with load-sharing, you must create a 1/*N* OCS with at least as many cards as you have HSMs. All the cards in the OCS must have the same passphrase.

> **ⓘ** The nCipherKM JCA/JCE CSP does not support HSM Pool mode. If you want to use HSM Pool mode with a Java application that only uses module protected keys, one option may be to use the Sun PKCS #11 provider to access the nShield PKCS #11 library instead of using nCipherKM JCA/JCE CSP.

Keys generated or imported by the nCipherKM JCA/JCE CSP are not recorded into the

Security World until:

1. The key is added to an nCipherKM KeyStore (by using a call to `setKeyEntry()` or `setCertificateEntry()`).
2. That nCipherKM KeyStore is then stored (by using a call to `store()`).

The passphrase used with the KeyStore must be the passphrase of the card from the OCS that protects the keys in the KeyStore.

10.2. nShield PKCS #11 library

> **!** Do not use PKCS #11 to perform any task that requires an Administrator Card. Use the equivalent nShield utilities instead.

To use the nShield PKCS #11 library, you must tell the application the name and location of the library. The exact method for doing this depends on the application.

Instructions for using the nShield PKCS #11 library with specific applications are available from Entrust nShield Technical Support, https://nshieldsupport.entrust.com.

Depending on the application, you may need to set the path and library name `/opt/nfast/toolkits/pkcs11/libcknfast.so` (**Linux**) or `%NFAST_HOME%\toolkits\pkcs11\cknfast.dll` (**Windows**) in a dialog or configuration file.

The nShield PKCS #11 library has security options which you must configure before you use the PKCS #11 library. For more information, see PKCS #11 library with Security Assurance Mechanism.

From version 1.7, the nShield PKCS #11 library can be used with FIPS 140 Level 3 compliant Security Worlds. This version of the library also introduces load-sharing mode. This feature provides support for multiple hardware security modules that are connected to a single server, spreading the load of cryptographic operations between the HSMs in order to provide scalability in terms of performance.

To share OCS protected keys with load-sharing mode, you must create a 1/*N* OCS that contains at least as many cards as you have HSMs. All the cards on the OCS must have the same passphrase.

With module firmware version 2.65.2 or later, if your application only uses module protected keys, you can use HSM Pool mode as an alternative to using load-sharing mode. HSM Pool mode supports returning or adding a hardware security module to the pool without restarting the system.

> ℹ️ If you are using the `preload` command-line utility in conjunction with the nShield PKCS #11 library, you can create *K/N* OCSs.

10.2.1. Choosing functions

Some PKCS #11 applications enable you to choose which functions you want to perform on the PKCS #11 token and which functions you want to perform in your application.

The following paragraphs in this section describe the functions that an nShield HSM can provide.

10.2.1.1. Generating random numbers and keys

The nShield HSM includes a hardware random number generator. A hardware random number generator provides greater security than the pseudo-random number generators provided by host computers. Therefore, always use the nShield HSM to generate random numbers and keys.

10.2.1.2. Digital signatures

The nShield PKCS #11 library can use the nShield HSM to sign and verify messages using the following algorithms:

- DSA
- RSA
- DES3_MAC
- AES
- ECDSA (if the appropriate feature is enabled)

An nShield hardware security module is specifically optimized for public key algorithms, and therefore it will provide significant acceleration for DSA, RSA and ECDSA signature generation and verification. You should always choose to perform asymmetric signature generation and verification with an nShield HSM.

10.2.1.3. Asymmetric encryption

The nShield PKCS #11 library can use an nShield HSM to perform asymmetric encryption and decryption with the RSA algorithm.

The nShield HSM is specifically optimized for asymmetric algorithms, so you should always

choose to perform asymmetric operations with the nShield HSM.

10.2.1.4. Symmetric encryption

The nShield PKCS #11 library can use the nShield HSM to perform symmetric encryption with the following algorithms:

- DES
- Triple DES
- AES

Because of limitations on throughput, these operations can be slower on the nShield HSM than on the host computer. However, although the nShield HSM may be slower than the host under a light load, you may find that under a heavy load the advantage gained from off-loading the symmetric cryptography (which frees the host CPU for other tasks) means that you achieve better overall performance.

10.2.1.5. Message digest

The nShield PKCS #11 library can perform message digest operations with MD5, SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 algorithms. However, for reasons of through-put, the library performs these operations on the host computer.

10.2.1.6. Mechanisms

The mechanisms currently supported by the nShield PKCS #11 library, including some vendor-supplied mechanisms, are listed in the *Cryptographic API Integration Guide*.

10.2.1.7. Key wrapping

The nShield PKCS #11 library can use an nShield HSM to wrap (encrypt) a private or secret key, or to unwrap (decrypt) a wrapped key.

The mechanisms supported by the nShield PKCS #11 library, including some vendor-supplied mechanisms, are listed in the *Cryptographic API Integration Guide*.

10.2.2. PKCS #11 library with Security Assurance Mechanism

It is possible for an application to use the PKCS #11 API in ways that do not necessarily provide the expected security benefits, or which might introduce additional weaknesses. For

example, the PKCS #11 standard requires the nShield library to be able to generate keys that are extractable from the HSM in plaintext. An application could use this ability in error, when a secure key would be more appropriate.

The PKCS #11 library with the Security Assurance Mechanism (SAM), `libcknfast`, can help users to identify potential weaknesses, and help developers create secure PKCS #11 applica tions more easily.

The SAM in the PKCS #11 library is intended to detect operations that reveal questionable behavior by the application. If these occur, the application fails with an explanation of the cause of failure.

After a review of your security policy and the way the application uses the PKCS #11 library with the SAM, if there are questionable operations that are considered to be acceptable and pose no security risk, the PKCS #11 library can be configured to permit some, or all, of them by means of the `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` environment variable (described in CKNFAST_OVERRIDE_SECURITY_ASSURANCES).

> To ensure the security of your keys, you must review any messages returned by the PKCS #11 library before changing the settings of the `CKN FAST_OVERRIDE_SECURITY_ASSURANCES` environment variable.

The `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` environment variable uses a semicolon sepa-rated list of parameters, with associated values, to explicitly allow operations that could compromise the security of cryptographic keys if the operations are not well understood.

If no parameters, or the `none` parameter, are supplied to the `CKNFAST_OVERRIDE_SECURITY_AS SURANCES`, the PKCS #11 library fails to perform the operation in question, and issues a warn-ing, when the following operations are detected:

- Creating short-term session keys as long-term objects
- Creating keys that can be exported as plain text
- Importing keys from external sources
- Creating or importing wrapping keys
- Creating or importing unwrapping keys
- Creating keys with weak algorithms (such as DES)
- Creating keys with short key lengths.

For more information about parameters and diagnostic warnings, see CKNFAST_OVER-RIDE_SECURITY_ASSURANCES.

10.2.2.1. Key security

Questionable operations largely relate to the concept of a key being *secure*. A private or secret key is considered insecure if there is some reason for believing that its value may be available outside the HSM. Public keys are never considered insecure; by definition they are intended to be public.

An explicitly insecure PKCS #11 key is one where `CKA_SENSITIVE` is set to false. If an application uses a key that is insecure but `CKA_SENSITIVE` is not set to false, it is possible that the application is using an inadequate concept of key security, and that the library disallows use of that key by default. Use of insecure keys should, by default, be restricted to short-term session keys, and applications should explicitly recognize the insecurity.

10.2.3. Using the nShield PKCS #11 library

After you have loaded the nShield PKCS #11 library, it is added to your application's list of cryptographic HSMs or PKCS #11 slots.

Whether or not the library uses load-sharing mode depends on the value of the `CKN-FAST_LOADSHARING` environment variable, described in CKNFAST_LOADSHARING. Whether or not the library uses HSM Pool mode depends on the value of the `CKNFAST_HSM_POOL` environment variable, described in CKNFAST_HSM_POOL.

10.2.3.1. nShield PKCS #11 library with load-sharing mode

If load-sharing mode is enabled, the nShield PKCS #11 library creates a virtual slot for each OCS in the Security World (returning the name of the card set) unless you have set `CKN-FAST_CARDSET_HASH` (as described in CKNFAST_CARDSET_HASH).

An additional virtual slot may be returned (with the label of `accelerator`), depending on the value given to the variable `CKNFAST_NO_ACCELERATOR_SLOTS` (described in CKNFAST_NO_AC-CELERATOR_SLOTS). Accelerator slots can:

- Be used to support session objects
- Be used to create module-protected keys
- Not be used to create private objects.

When you insert a smart card from an OCS in the current Security World, the nShield PKCS #11 library treats this card as a PKCS #11 token that is present in the virtual slot for that OCS.

After the PKCS #11 token is present, you can open a session to that token. Until you log in, a session can only access public objects that belong to that PKCS #11 token.

The PKCS #11 token is present until you remove the last card belonging to the OCS. When you remove the token, the nShield PKCS #11 library closes any open sessions.

Logging in gives access to the private objects that are protected by the PKCS #11 token. Logging in requires the passphrase for the OCS. The exact mechanism for supplying the passphrase depends on the application that you are running.

The PKCS #11 token is shared across all the HSMs that have a smart card from the OCS in the reader at the point that you log in. After you have logged in, inserting additional cards from this OCS has no effect.

If you remove a smart card that belongs to a logged-in token, the nShield PKCS #11 library closes any open sessions and marks the token as being not present (unless the OCS is persistent). Removing a card from a persistent OCS has no effect, and the PKCS #11 token remains present until you log out.

10.2.3.2. nShield PKCS #11 library with HSM Pool mode

If HSM Pool mode is enabled, the nShield PKCS #11 library exposes a single pool of HSMs and a single virtual slot for a fixed token with the label `accelerator`. This accelerator slot can be used to create module protected keys and to support session objects. HSM Pool mode does not support token protected keys, any pre-existing OCS or softcard protected keys are hidden from PKCS #11. In FIPS 140 Level 3 Security Worlds, keys cannot be created in HSM Pool mode, however keys created outside HSM Pool mode, for example using `generatekey` or a non-Pool mode PKCS #11 application, can be used in HSM Pool mode.

10.2.3.3. nShield PKCS #11 library without load-sharing

There will be two entries for each HSM, unless you have set `CKNFAST_NO_ACCELERATOR_SLOTS`.

> The entry called `accelerator` cannot be used to create private objects. It can be used to create module-protected keys.

Use the second of the two entries (which has the same name as the Operator Card that is currently in a smart card reader) to protect your keys or token objects.

PKCS #11 does not allow two tokens to be present in the same slot. Therefore, when you insert a smart card into a reader, the nShield PKCS #11 library logs out any previously logged-in token from the slot and closes any open sessions.

10.2.3.4. nShield PKCS #11 library with the preload utility

You can use the `preload` command-line utility to preload *K/N* OCSs before actually using PKCS #11 applications. The `preload` utility loads the logical token and then passes it to the PKCS #11 utilities.

You must provide any required passphrase for the tokens when using `preload` to load the card set. However, because the application is not aware that the card set has been pre-loaded, the application operates normally when handling the login activity (including prompting for a passphrase), but the PKCS #11 library will not actually check the supplied passphrase. `preload` must be also used with the `cksotool` utility to perform operations that require the PKCS #11 Security Officer role.

Normally, `preload` uses environment variables to pass information to the program using the preloaded objects, including the PKCS #11 library. Therefore, if the application you are using is one that clears its environment before the PKCS #11 library is loaded, you must set the appropriate values in the `cknfastrc` file (see nShield PKCS #11 library environment variables). The current environment variables remain usable. The default setting for the `CKNFAST_LOADSHARING` environment variable changes from specifying load-sharing as disabled to specifying load-sharing as enabled. Moreover, in load-sharing mode, the loaded card set is used to set the environment variable `CKNFAST_CARDSET_HASH` so that only the loaded card set is visible as a slot.

The `NFAST_NFKM_TOKENSFILE` environment variable must also be set in the `cknfastrc` file to the location of the preload file (see Environment variables).

A logical token preloaded by `preload` for use with the nShield PKCS #11 library is the only such token available to the application for the complete invocation of the library. You can use more than one HSM with the same card set.

If the loaded card set is non-persistent, then a card must be left in each HSM on which the set has been loaded during the start-up sequence. After a non-persistent card has been removed, the token is not present even if the card is reinserted.

If load-sharing has been specifically switched off, you see multiple slots with the same label.

10.2.4. nShield PKCS #11 library environment variables

The nShield PKCS #11 library uses the following environment variables:

- `CKNFAST_ASSUME_SINGLE_PROCESS`
- `CKNFAST_ASSURANCE_LOG`
- `CKNFAST_CARDSET_HASH`

- CKNFAST_CONCATENATIONKDF_X963_COMPLIANCE

- CKNFAST_DEBUG

- CKNFAST_DEBUGDIR

- CKNFAST_DEBUGFILE

- CKNFAST_DH_LSB

- CKNFAST_FAKE_ACCELERATOR_LOGIN

- CKNFAST_HSM_POOL

- CKNFAST_JCE_COMPATIBILITY

- CKNFAST_LOADSHARING

- CKNFAST_LOAD_KEYS

- CKNFAST_NO_ACCELERATOR_SLOTS

- CKNFAST_NO_SYMMETRIC

- CKNFAST_NO_UNWRAP

- CKNFAST_NONREMOVABLE

- CKNFAST_NVRAM_KEY_STORAGE

- CKNFAST_OVERRIDE_SECURITY_ASSURANCES

- CKNFAST_RELOAD_KEYS

- CKNFAST_SEED_MAC_ZERO

- CKNFAST_SESSION_THREADSAFE

- CKNFAST_SHARE_SESSION_KEYS

- CKNFAST_TOKENS_PERSISTENT

- CKNFAST_USE_THREAD_UPCALLS

- CKNFAST_WRITE_PROTECTED

If you used the default values in the installation script, you should not need to change any of these environment variables.

You can set environment variables in the file `cknfastrc`.

Linux

This file must be in the `/opt/nfast/` directory of the client.

Windows

If the `NFAST_HOME` environment variable is not set, or if environment variables are cleared by your application, the file `cknfastrc` must be in the `%NFAST_HOME%` directory of the client.

> ℹ️ The `cknfastrc` file should be saved without any suffix (such as `.txt`).

Each line of the file `cknfastrc` must be of the following form:

```
<variable>=<value>
```

> ℹ️ Variables set in the environment are used in preference to those set in the resource file.

Changing the values of these variables after you start your application has no effect until you restart the application.

If the description of a variable does not explicitly state what values you can set, the values you set are normally `1` or `0`, `Y` or *N*.

> ℹ️ For more information concerning Security World Software environment variables that are not specific to PKCS #11 and which are used to configure the behavior of your nShield installation, see the Security World Software installation instructions.

10.2.4.1. CKNFAST_ASSUME_SINGLE_PROCESS

By default, this variable is set to `1`. This specifies that only token objects that are loaded at the time `C_Initialize` is called are visible.

Setting this variable to `0` means that token objects created in one process become visible in another process when it calls `C_FindObjects`. Existing objects are also checked for modification on disc; if the key file has been modified, then the key is reloaded. Calling `C_SetAttributeValues` or `C_GetAttributeValues` also checks whether the object to be changed has been modified in another process and reloads it to ensure the most recent copy is changed.

Setting the variable to `0` can slow the library down because of the additional checking needed if a large number of keys are being changed and a large number of existing objects must be reloaded.

10.2.4.2. CKNFAST_ASSURANCE_LOG

This variable is used to direct all warnings from the Security Assurance Mechanism to a specific log file.

10.2.4.3. CKNFAST_CARDSET_HASH

This variable enables you to specify a specific card set to be used in load-sharing mode. If this variable is set, only the virtual smart card slot that matches the specified hash is

present (plus the accelerator slot). The hash that you use to identify the card set in `CKN-FAST_CARDSET_HASH` is the SHA-1 hash of the secret on the card. Use the `nfkminfo` command-line utility to identify this hash for the card set that you want to use: it is listed as `hkltu`. For more information about using `nfkminfo`, see *nfkminfo: information utility*.

10.2.4.4. CKNFAST_CONCATENATIONKDF_X963_COMPLIANCE

Sets the correct use of ECDH derive with concatenate KDF using the ANSI X9.63 specification as per the PKCS#11 standard.

> 🛈 | The default is ANSI X9.63 to match that of the PKCS #11 Specification.

> 🛈 | ECDH derive with concatenate KDF SP800-56a can use the standard PKCS #11 v3 `CKD_SHA[x]_SP800_KDF` values.

10.2.4.5. CKNFAST_DEBUG

This variable is set to enable PKCS #11 debugging. The values you can set are in the range `0` - `11`. If you are using `NFLOG_*` for debugging, you must set `CKNFAST_DEBUG` to `1`.

Value	Description
0	None (default setting)
1	Fatal error
2	General error
3	Fix-up error
4	Warnings
5	Application errors
6	Assumptions made by the nShield PKCS #11 library
7	API function calls
8	API return values
9	API function argument values
10	Details
11	Mutex locking detail

10.2.4.6. CKNFAST_DEBUGDIR

If this variable is set to the name of a writeable directory, log files are written to the specified directory. The name of each log file contains a process ID. This can make debugging easier for applications that fork a lot of child processes.

10.2.4.7. CKNFAST_DEBUGFILE

You can use this variable to write the output for `CKNFAST_DEBUG` (`Path name > file name`).

10.2.4.8. CKNFAST_DH_LSB

If this variable is set the least significant bytes of the result of DH/ECDH key agreement using the `CKM_DH_PKCS_DERIVE`, `CKM_X9_42_DH_DERIVE` or `CKM_ECDH1_DERIVE` mechanisms are taken. This is in line with the PKCS#11 specification. If this variable is not set the most significant bytes will be used. The latter behavior is consistent with Security World software prior to v12.81.

10.2.4.9. CKNFAST_FAKE_ACCELERATOR_LOGIN

If this variable is set, the nShield PKCS #11 library accepts a PIN for a module-protected key, as required by Sun Java Enterprise System (JES), but then discards it. This means that a Sun JES user requesting a certificate protected by a load-shared HSM can enter an arbitrary PIN and obtain the certificate.

`CKNFAST_FAKE_ACCELERATOR` slots allow the creation of objects with `CKA_PRIVATE=TRUE` in the template even though the login is "fake" and the objects are not private.

- Examining the attributes shows `CKA_PRIVATE` as `FALSE`.
- A search for the object will not find it if the search criteria includes `CKA_PRIVATE=TRUE`.

10.2.4.10. CKNFAST_HSM_POOL

HSM Pool mode is determined by the state of the `CKNFAST_HSM_POOL` environment variable.

Set the environment variable to *1*, *y* or *Y* to enable HSM Pool mode for the PKCS #11 application, or set to *0*, *n* or *N* to explicitly disable HSM Pool mode for the PKCS #11 application.

HSM Pool mode takes precedence over load-sharing mode. HSM Pool mode only supports module protected keys so do not use `CKNFAST_NO_ACCELERATOR_SLOTS` to disable the accelerator slot.

10.2.4.11. CKNFAST_JCE_COMPATIBILITY

This property is included to allow the saving of objects when using Java PKCS#11 providers.

It is possible, using `C_CopyObject`, to change a key's `CKA_TOKEN` value from `CK_FALSE` to `CK_TRUE`. This requires the `CKNFAST_JCE_COMPATIBILITY` environment variable to be set to `1`. The original key's `CKA_TOKEN` value will remain unchanged.

10.2.4.12. CKNFAST_LOADSHARING

Load-sharing mode is determined by the state of the `CKNFAST_LOADSHARING` environment variable.

To enable load-sharing mode, set the environment variable `CKNFAST_LOADSHARING` to a value that starts with something other than *0*, *n*, or *N* and ensure that the `CKNFAST_HSM_POOL` environment variable is not set. The virtual slot behavior then operates.

> To use softcards with PKCS #11, you must have `CKNFAST_LOADSHARING` set to a nonzero value. When using pre-loaded softcards or other objects, the PKCS #11 library automatically sets `CKNFAST_LOADSHARING=1` (load-sharing mode on) unless it has been explicitly set to `0` (load-sharing mode off).

10.2.4.13. CKNFAST_NO_ACCELERATOR_SLOTS

If this variable is set, the nShield PKCS #11 library does not create the accelerator slot, and thus the library only presents the smart card slots (real or virtual, depending on whether load-sharing is in use).

Do not set this environment variable if you want to use the accelerator slot to create or load module-protected keys.

> Setting this environment variable has no effect on `ckcheckinst` because `ckcheckinst` needs to list accelerator slots.

10.2.4.14. CKNFAST_NO_SYMMETRIC

If this variable is set, the nShield PKCS #11 library does not advertise any symmetric key operations.

10.2.4.15. CKNFAST_NO_UNWRAP

If this variable is set, the nShield PKCS #11 library does not advertise the `c_wrap` and `c_unwrap` commands. You should set this variable if you are using Sun Java Enterprise System

(JES) or Netscape Certificate Management Server as it ensures that a standard SSL hand-shake is carried out. If this variable is not set, Sun JES or Netscape Certificate Management Server make extra calls, which reduces the speed of the library.

10.2.4.16. CKNFAST_NONREMOVABLE

When this environment variable is set, the state changes of the inserted card set are ignored by the nShield PKCS #11 library.

> **ⓘ** Since protection by non-persistent cards is enforced by the HSM, not the library, this variable does not make it possible to use keys after a non-persistent card is removed, or after a timeout expires.

10.2.4.17. CKNFAST_NVRAM_KEY_STORAGE

When this environment variable is set, the PKCS #11 library generates only keys in non-volatile memory (NVRAM). You must also ensure this environment variable is set in order to delete NVRAM-stored keys.

10.2.4.18. CKNFAST_OVERRIDE_SECURITY_ASSURANCES

This variable can be assigned one or more of the following parameters, with an associated value where appropriate, to override the specified security assurances in key operations where this is deemed acceptable:

- `all`
- `none`
- `tokenkeys`
- `longterm [=<days>]`
- `explicitness`
- `import`
- `unwrap_mech`
- `unwrap_kek`
- `derive_kek`
- `derive_xor`
- `derive_concatenate`
- `unwrap_rsa_aes_kwp`
- `weak_<algorithm>`

- `shortkey_<algorithm>=<bitlength>`
- `silent`.

Each parameter specified is separated by a semicolon. Using the command line, enter the following to set the variable:

Linux

```
CKNFAST_OVERRIDE_SECURITY_ASSURANCES="<parameter1>;<parameter2>=<value3>"
```

Windows

```
set CKNFAST_OVERRIDE_SECURITY_ASSURANCES=<parameter1>;<parameter2>=<value3>
```

In the configuration file, enter the following to set the variable:

```
CKNFAST_OVERRIDE_SECURITY_ASSURANCES=<parameter1>;<parameter2>=<value3>
```

Unknown parameters generate a warning; see Diagnostic warnings about questionable operations.

The meaning of these parameters is described in the rest of this section.

10.2.4.18.1. all

The `all` parameter overrides all security checks and has the same effect as supplying all the other `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` parameters except the `none` parameter. Using the `all` parameter prevents the library from performing any of the security checks and allows the library to perform potentially insecure operations. This parameter cannot be used with any other parameters.

10.2.4.18.2. none

The `none` parameter does not override any of the security checks and has the same effect as supplying no parameters. Using the `none` parameter allows the library to perform all security checks and warn about potentially insecure operations without performing them. This parameter cannot be used with any other parameters.

10.2.4.18.3. tokenkeys

The `tokenkeys` parameter permits applications to request that insecure keys are stored long-term by the cryptographic hardware and library.

Some PKCS #11 applications create short-term session keys as long-term objects in the cryptographic provider, for which strong protection by the HSM is not important. There-fore, provided that you intend to create long-term keys, the need to set this token does not always indicate a potential problem because the `longterm` keys restriction is triggered auto-matically. If you set the `tokenkeys` parameter, ensure that your Quality Assurance process tests all of your installation's functionality at least 48 hours after the system was set up to check that the key lifetimes are as expected.

When the `tokenkeys` parameter is set, the effect on the PKCS #11 library is to permit inse-cure Token keys. By default, any attempts to create, generate, or unwrap insecure keys with `CKA_TOKEN=true` fails with `CKR_TEMPLATE_INCONSISTENT` and a log message that explains the insecurity. When tokenkeys is included as a parameter for `CKNFAST_OVERRIDE_SECURITY_AS-SURANCES`, attempts to create, generate, or unwrap insecure keys with `CKA_TOKEN=true` are allowed.

10.2.4.18.4. longterm[=*days*]

The `longterm` parameter permits an insecure key to be used for `days` after it was created. Usually insecure keys may not be used more than 48 hours after their creation. If `days` is not specified, there is no time limit.

> ℹ️ A need to set this variable usually means that some important keys that should be protected by the HSM's security are not secure.

When the `longterm` parameter is set, the PKCS #11 API permits the use of the following func tions with an insecure key up to the specified number of `days` after its creation:

- `C_Sign` and `C_SignUpdate`
- `C_Verify` and `C_VerifyUpdate`
- `C_Encrypt` and `C_EncryptUpdate`
- `C_Decrypt` and `C_DecryptUpdate`.

By default these functions fail with `CKR_FUNCTION_FAILED`, or `CKR_KEY_FUNCTION_NOT_PERMIT-TED`, and a log message that explains the insecurity of these functions when used with an insecure private or secret key more than 48 hours after the creation of the key as indicated by `time()` on the host.

When the `longterm` parameter is set, the functions `C_SignInit`, `C_VerifyInit`, `C_Encryp-tInit`, and `C_DecryptInit` check the `CKA_CREATION_DATE` against the current time.

10.2.4.18.5. explicitness

The `explicitness` parameter permits applications to create insecure keys without explicitly recognizing that they are insecure. An insecure key is a key that is deemed sensitive, but can be wrapped and extracted from the HSM by any untrusted key. A secure key must have the `CKA_WRAP_WITH_TRUSTED` attribute.

> ℹ️ A need to set the `explicitness` parameter does not necessarily indicate a problem, but does usually indicate that a review of the application's security policies and use of the PKCS #11 API should be carried out.

Unless the `explicitness` parameter is set, attempts to create, generate, or unwrap insecure keys with `CKA_SENSITIVE=true`, or to set `CKA_SENSITIVE=true` on an existing key, fail by default with `CKR_TEMPLATE_INCONSISTENT` and a log message explaining the insecurity. However, when the `explicitness` parameter is set, these operations are allowed.

10.2.4.18.6. import

The `import` parameter allows keys that are to be imported into the HSM's protection from insecure external sources to be treated as secure, provided that the application requests security for them. Usually, the library treats imported keys as insecure for the purposes of checking the security policy of the application. Even though the imported copy may be secure, insecure copies of the key may still exist on the host and elsewhere.

If you are migrating from software storage to hardware protection of keys, you must enable the `import` parameter at the time of migration. You can disable `import` again after migrating the keys.

> ℹ️ Setting this variable at any other time indicates that the library regards the key as secure, even though it is not always kept within a secure environment.

When the `import` parameter is set, the PKCS #11 API treats keys that are imported through `C_CreateObject` or `C_UnwrapKey` as secure (provided there is no other reason to treat them as insecure). By default, keys which are imported through `C_CreateObject` or `C_UnwrapKey` without this option in effect are marked as being insecure. Only the setting of the parameter at the time of import is relevant.

10.2.4.18.7. unwrap_mech

The `unwrap_mech` parameter allows you to create keys with `CKA_UNWRAP=true` and `CKA_DECRYPT=false`.

By default, when `unwrap_mech` is not supplied as a parameter for `CKNFAST_OVERRIDE_SECURI-`

`TY_ASSURANCES`, trying to create a key with `CKA_UNWRAP=true` and `CKA_DECRYPT=false` fails with `CKR_TEMPLATE_INCONSISTENT`.

When `CKA_UNWRAP` is set to `true` and `CKA_DECRYPT` is not specified in the template, then `CKA_DECRYPT` is automatically set to `true`.

10.2.4.18.8. unwrap_kek

When a key is transferred into the HSM in encrypted form, the key is usually treated as inse cure unless the key that was used for the decryption only allows the import and export of keys and not the decryption of arbitrary messages. This behavior is necessary to prevent an unauthorized application from simply decrypting the encrypted key instead of importing it. However, because PKCS #11 wrapping mechanisms are insecure, all unwrapping keys have `CKA_DECRYPT=true`.

By default, keys that are unwrapped with a key that has `CKA_DECRYPT` permission are considered insecure. When the `unwrap_kek` parameter is set, the PKCS #11 API considers keys that are unwrapped with a key that also has `CKA_DECRYPT` permission as secure (provided there is no other reason to treat them as insecure).

10.2.4.18.9. derive_kek

By default, keys that have been derived by using `CKM_DES3_ECB_ENCRYPT_DATA` with a key that has `CKA_ENCRYPT` permission are considered insecure. However, when the `derive_kek` parameter is set, the PKCS #11 API considers keys that are derived with a key that has `CKA_ENCRYPT` permission as secure (provided that there is no other reason to treat them as insecure).

10.2.4.18.10. derive_xor

Normally, you can only use only extractable keys with `CKM_XOR_BASE_AND_DATA` and, on unextractable keys, only `CKM_DES3_ECB_ENCRYPT_DATA` is allowed by `CKA_DERIVE`. However, when the `derive_xor` parameter is set, the PKCS #11 API also allows such functions with keys that are not extractable and treats them as secure (provided that there is no other reason to treat them as insecure).

10.2.4.18.11. derive_concatenate

Normally, you can only use session keys with `CKM_CONCATENATE_BASE_AND_KEY` for use with the operation `C_DeriveKey`. However, when the `derive_concatenate` parameter is set, the PKCS #11 API also allows such functions with keys that are long term (token) keys. The

PKCS #11 API treats these keys as secure, provided there is no other reason to treat them as insecure. Even if the `all` parameter is set, if you do not include the `CKA_ALLOWED_MECHA-NISMS` with `CKM_CONCATENATE_BASE_AND_KEY`, this `C_DeriveKey` operation will not be allowed.

10.2.4.18.12. unwrap_rsa_aes_kwp

The `unwrap_rsa_aes_kwp` parameter only applies to firmware version 13.3 or earlier. It is not needed in later versions.

The `C_UnwrapKey` operation with `CKM_RSA_AES_KEY_WRAP` imports the temporary AES key with an nCore API ACL that permits unwrapping of the wrapped target key by the temporary AES key. When using the `C_UnwrapKey` operation with only a user supplied template (`pTemplate`) it is possible to create this ACL such that it permits a one-time unwrap of only the wrapped target key. When the RSA unwrapping key has `CKA_UNWRAP_TEMPLATE` set it is necessary to construct the ACL when the RSA key is created in order to setup the partitioning guarantees from the `CKA_UNWRAP_TEMPLATE`. The intended wrapped target keys are unknown at this time, which means the ACL must permit a one-time unwrap of any key.

The Security Assurance Mechanism (SAM) considers this scenario insecure by default and therefore the use of the `C_UnwrapKey` operation with `CKM_RSA_AES_KEY_WRAP` is disabled when the RSA unwrapping key has `CKA_UNWRAP_TEMPLATE` set. When the `unwrap_rsa_aes_kwp` parameter is set the SAM enables the `C_UnwrapKey` operation with `CKM_RSA_AES_KEY_WRAP` in this scenario. The RSA unwrapping key must also explicitly allow the `CKM_RSA_AES_KEY_WRAP` mechanism via `CKA_ALLOWED_MECHANISMS` in addition to setting the `unwrap_rsa_aes_kwp` (or `all`) parameter; otherwise, the `C_UnwrapKey` operation will remain disabled when the RSA unwrapping key has `CKA_UNWRAP_TEMPLATE` set.

10.2.4.18.13. weak_<algorithm>

The `weak_<algorithm>` parameter allows you to treat keys used with a weak algorithm as secure. For example, DES is not secure, but setting the parameter `weak_des` means that such keys are considered secure. You can apply the `weak_<algorithm>` parameter to all keys that have a short fixed key length or whose algorithms have other security problems. As a guide, weak algorithms are those whose work factor to break is less than approximately 80 bits.

10.2.4.18.14. shortkey_<algorithm=bitlength>

The `shortkey_<algorithm=bitlength>` parameter permits excessively short keys for the specified `<algorithm>` to be treated as secure. The parameter `<bitlength>` specifies the minimum length, in bits, that is to be considered secure. For example, RSA keys must usually be

at least 1024 bits long in order to be treated as secure, but `shortkey_rsa=768` would allow 768-bit RSA keys to be treated as secure.

10.2.4.18.15. silent

The `silent` parameter turns off the warning output. Checks are still performed and still return failures correctly according to the other variables that are set.

10.2.4.18.16. Diagnostic warnings about questionable operations

When the `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` environment variable is set to a value other than `all`, diagnostic messages are always generated for questionable operations. Each message contains the following elements:

- The PKCS #11 label of the key, if available
- The PKCS #11 identifier of the key, if available
- The hash of the key
- A summary of the problem.

If the problem is not that a questionable operation has been permitted because of a setting in `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` it could be that an operation has failed. In such a case, the setting required to authorize the operation is noted.

By default, these messages are sent to `stderr`. On Windows platforms, they are also always sent to the Event Viewer. If a file name has been specified in the `CKNFAST_ASSURANCE_LOG` environment variable, diagnostic messages are also written to this file.

If `CKNFAST_DEBUG` is `1` or greater and a file is specified in `CKNFAST_DEBUGFILE`, the PCKS #11 library Security Assurance Mechanism log information is sent to the specified file. These variables must be set whenever `generatekey` or KeySafe are used.

> **ℹ** If a file is specified in `CKNFAST_ASSURANCES_LOG` and no file is specified in `CKNFAST_DEBUGFILE` (or if `CKNFAST_DEBUG` is `0`), diagnostic messages are sent to `stderr` as well as to the file specified in `CKNFAST_ASSURANCES_LOG`.

10.2.4.19. CKNFAST_SEED_MAC_ZERO

Set this variable to use zero padding for the Korean SEED MAC mechanisms (`CK_SEED_MAC` and `CKM_SEED_MAC_GENERAL`). If this variable is not set, or is set to `n`, then the SEED MAC mechanisms will use the default PKCS #5 padding scheme.

10.2.4.20. CKNFAST_SESSION_THREADSAFE

You must set this environment variable to `yes` if you are using the Sun PKCS #11 provider when running nCipherKM JCA/JCE code.

10.2.4.21. CKNFAST_SHARE_SESSION_KEYS

This variable can take a list of one or more semicolon (;) separated values to improve performance through loadsharing when session keys are used. See CKNFAST_LOADSHARING.

Loadsharing improves performance and adds resilience in the case of module failure. However, if the key is used only a few times, the overhead of sharing it may be greater than the performance benefit. If a key will be used many times or if it has a long lifespan, sharing is recommended.

- `all` (default)
- `copy`
- `derive`
- `generate`
- `import`
- `none`
- `unwrap`

If the origin of the session key matches a selected category, then the key is automatically shared to all HSMs when it is created.

10.2.4.22. CKNFAST_TOKENS_PERSISTENT

This variable controls whether or not the Operator Cards that are created by your PKCS #11 application are persistent. If this variable is set when your application calls the PKCS #11 function that creates tokens, the Operator Card created is persistent.

> **ⓘ** Use of the nShield PKCS #11 library to create tokens is deprecated, because it can only create 1/1 tokens in FIPS 140 Level 2 Security Worlds. Use KeySafe or one of the command-line utilities to create OCSs.

10.2.4.23. CKNFAST_USE_THREAD_UPCALLS

If this variable is set and `CKF_OS_LOCKING_OK` is passed to `C_Initialize`, `NFastApp_SetThread Upcalls` is called by means of `nfast_usencthreads` and only a single `NFastApp_Connection` is

used, shared between all threads.

If this variable is set and mutex callbacks are passed to `C_Initialize` but `CKF_OS_LOCK-ING_OK` is not passed, `C_Initialize` fails with `CKR_FUNCTION_FAILED`. (`NFastApp_SetThreadUp-calls` requires more callbacks than just the mutex ones that PKCS #11 supports.)

If neither mutex callbacks nor `CKF_OS_LOCKING_OK` is passed, this variable is ignored. Only a single connection is used because the application must be single threaded in this case.

10.2.4.24. CKNFAST_LOAD_KEYS

This variable will load private objects at `C_Login` time, rather than at the first cryptographic operation.

10.2.4.25. CKNFAST_WRITE_PROTECTED

Set this variable to make your OCS or softcard (token) write-protected. If a token is write-protected, you cannot:

- Generate certificate, data, and key objects for that token.
- Modify attributes of an existing object.

> This environment variable does not prevent you from deleting an object from your token.

10.2.4.26. CKNFAST_RELOAD_KEYS

Set this variable to enable PKCS #11 key reloading. See section *PKCS _11 with key reloading* in the *Cryptographic API Integration Guide*.

Key reloading requires load sharing-mode to operate, and enables it automatically if `CKN-FAST_LOADSHARING` is not set.

10.2.5. Checking the installation of the nShield PKCS #11 library

After you have created a Security World, ensure that the nShield PKCS #11 library has been successfully installed by using the `ckcheckinst` command-line utility.

To verify the installation of the nShield PKCS #11 library, follow these steps:

1. Give the command `ckcheckinst`.

   If you have an invalid Security World (for example, if all your HSMs are in the initializa-

---

tion state), `ckcheckinst` quits with the following error message:

```
ckcheckinst: C_Initialize failed rv = 00000006
Is the security world initialized? (Use nfkminfo to check)
```

If your Security World is valid, `ckcheckinst` displays information similar to the following:

```
PKCS#11 library interface version 2.40
 flags 0
 manufacturerID "nCipher Corp. Ltd "
 libraryDescription "nCipher PKCS#11 1.#.# "
 implementation version 1.##
 Load sharing and Failover enabled

slot Status Label
===== ====== ===== 0 Fixed token "accelerator "
1 Operator card "card2 "
2 Operator card "card3 "
Select slot Number to run library test or 'R'etry or to 'E'xit:
```

In this example output:

- `PKCS #11 library interface version 2.40` refers to the version of the PKCS #11 specification supported

- `implementation version 1.##` refers to the version of the nCipher PKCS #11 library

- `Loadsharing and Failover enabled` is shown if load-sharing has been enabled. Alternatively `Pool mode enabled` is shown if Pool mode has been enabled.

Slots that contain a valid Operator Card are indicated by the status `Operator card` and the card's label. A fixed token is always available and is listed as slot 0.

If you insert a blank card or an unrecognized card (for example, an Operator Card from a different Security World or an Administrator Card), this is indicated in the `Status` column. The corresponding slot number is not available.

> 🛈 If you are using the `preload` command-line utility in conjunction with the nShield PKCS #11 library, you can only see the token that you loaded with the `preload` utility. In load-sharing mode, the loaded card set is used to set the environment variable `CKNFAST_-CARDSET_HASH`, so only this card set is visible as a slot.

If there is no card in a slot, `ckcheckinst` displays `No token present` beside the relevant slot numbers.

`ckcheckinst` gives you the following choices:

```
No removable tokens present.
Please insert an operator card into at least one available slot and
enter 'R' retry.
If you have not created an operator card or there are no physical slots, enter a fixed token slot number,
or 'E' to exit this program and create a card set before continuing.
```

2. If there are no available slots with cards in them, you can choose one of the following actions:

   ◦ Insert a valid Operator Card, and press **R**

   ◦ choose a fixed token slot

   ◦ Press **E** to quit, then create an OCS, and run `ckcheckinst` again.

   When there is at least one slot with a valid token, input a slot number, and press **Enter**. In a FIPS 140 Level 3 compliant Security World, `ckcheckinst` prompts you to enter the passphrase for the selected Operator Card.

3. Type the passphrase, and press **Enter**.

   `ckcheckinst` displays the results of the tests:

```
Test Pass/Failed
---- -----------
1 Generate RSA key pair Pass
2 Generate DSA key pair Pass
3 Encryption/Decryption Pass
4 Signing/Verify Pass
Deleted test keys ok
PKCS11 Library test successful.
```

If any tests fail, `ckcheckinst` displays a message indicating the failure and quits. It does not run any subsequent tests.

If `ckcheckinst` fails:

   ◦ Check that the hardserver is running

   ◦ Use the `enquiry` and `nfkminfo` world.

If all seems in order, reinstall the nShield library.

10.2.6. How the nShield PKCS #11 library protects keys

Session objects are created on an HSM and never leave that HSM. The following table lists the protection for different types of PKCS #11 token objects:

	Smart card Slot	Accelerator Slot
Private Token Object	Operator Card Set	not supported

	Smart card Slot	**Accelerator Slot**
Public Token Object	Security World	Security World
Public key	well known HSM key	well known HSM key

`Operator Card Set`

The object is stored as an nShield key blob encrypted by the OCS key. You must log in to this OCS before you can load this object.

`security world`

The object is stored as an nShield key blob encrypted by the Security World key. This object can be loaded on to any HSM in the Security World. The nShield PKCS #11 library only allows access if a card from this OCS is present.

`well-known module key`

Public keys are encrypted under a well-known HSM key. This encryption is for programming convenience only and does not provide security. These keys can be loaded on any nShield HSM.

10.3. nShield native and custom applications

Use the nShield native option for applications that were written using nShield key management software and that expect keys to be both protected by the Security World and stored in the Security World data structure.

Use the `custom` external application option for applications that were written using nShield key management software and that expect their keys to be in standalone files.

> ℹ️ KeySafe does not place any restrictions on the OCS that is used to protect nShield native or `custom` application keys. You must make sure that your application is capable of loading the card set.

10.4. Microsoft CAPI CSP

We provide a Cryptographic Service Provider (CSP) that implements the Cryptographic API (CAPI) supported in Windows 2003 and later.

10.4.1. Installing the CAPI CSP

A shortcut to the CSP installation wizard is placed in the Start menu: **Start > Entrust nShield Security World** when installing the Security World Software. If you want to use 32-bit applications with the nShield CAPI provider run the 32-bit installation wizard to install the CAPI CSP, and if you want to use 64-bit applications with the nShield CAPI provider run the 64-bit CSP installation wizard to install the CAPI CSP.

You can also use the CSP installation wizard to load existing Security Worlds, see Adding an HSM to a Security World with the CSP or CNG wizard, generate new Operator Card Sets, see Creating an Operator Card Set with the CSP or CNG wizard, and configure the set-up parameters of the CAPI CSP including HSM Pool mode.

With module firmware version 2.65.2 or later, if your application only uses module protected keys, you can use HSM Pool mode with multiple hardware security modules. HSM Pool mode exposes a single pool of HSMs and supports returning or adding a hardware security module to the pool without restarting the system. With a FIPS 140 Level 3 Security World, keys cannot be created in HSM Pool mode, however keys created outside HSM Pool mode can be used in HSM Pool mode.

> ℹ️ The CSP installation wizard is not suitable for creating complex Security World setups. When creating such Security Worlds, or if you require more flexibility than the CSP installation wizard provides, we recommend following the instructions in Creating a Security World using new-world.

Use the standard Security World utility `nfkmverify` to check the security of all stored keys in the Security World; `nfkminfo`, `nfkmcheck` and other standard utilities can also be used to assist in this process.

The CSP installation wizard registers the CAPI CSP as a key provider on your system.

10.4.2. Importing a key

Use the `cspimport` utility to move keys between containers or to import a pre-generated NFKM key into a container. For more information about using the `cspimport` utility, run `cspimport` specifying either the `--help` or `--usage` options.

10.4.3. Supported algorithms

The nShield CSPs support a similar range of algorithms to the Microsoft CSP.

10.4.3.1. Symmetric algorithms

- `CALG_DES`
- `CALG_3DES_112` (double-DES)
- `CALG_3DES`
- `CALG_RC4`
- `CALG_AES_128`
- `CALG_AES_192`
- `CALG_AES_256`

10.4.3.2. Asymmetric algorithms

- `CALC_RSA_SIGN` (only Enhanced RSA and AES Cryptographic Provider)
- `CALC_RSA_KEYX` (only Enhanced RSA and AES Cryptographic Provider)
- `CALC_DSA_SIGN` (only Enhanced DSS and Diffie-Hellman Cryptographic Provider and DSS Signature Cryptographic Provider)
- `CALC_DSS_SIGN` (only Enhanced DSS and Diffie-Hellman Cryptographic Provider)
- `CALC_DH_KEYX` (only Enhanced DSS and Diffie-Hellman Cryptographic Provider)
- `CALC_DH_SF` (only Enhanced DSS and Diffie-Hellman Cryptographic Provider)
- `CALC_DH_EPHEM` (only Enhanced DSS and Diffie-Hellman Cryptographic Provider)

10.4.3.3. Hash algorithms

- `CALG_SHA1`
- `CALG_SHA256`
- `CALG_SHA384`
- `CALG_SHA512`
- `CALG_SSL3_SHAMD5`
- `CALG_MD5`
- `CALG_MAC`
- `CALG_HMAC`

In addition, the Enhanced SChannel Cryptographic Provider and the Enhanced DSS and Diffie-Hellman SChannel Cryptographic Provider support all the internal algorithm types necessary for SSL3 and TLS1 support.

The nShield CSPs do not support SSL2.

10.4.4. Container storage format

> **ℹ** Versions of the CSP later than 1.11.0 have an updated container storage mechanism. CSP containers are now stored as part of the Security World instead of in the Windows registry file.

> **⚠** Versions of the CSP later than 1.11.0 use a non-backwards-compatible container and key storage format. If you are installing version 1.11.0 or later of the CSP over older versions, you must run the `cspmigrate` utility in order to convert containers and keys from the old system to the new system.

CSP versions 1.11.0 and later have a number of advantages over older versions:

- The CSP state is easily mirrored between multiple machines simply by copying the contents of the Key Management Data directory or by sharing the Key Management Data directory across a network.
- The CSP key files can have arbitrary names (previously, the names of key files were linked to their key type and their container name). This new method facilitates the importation of existing Security World keys into the CSP.
- Every different container is now guaranteed to have a distinct storage location. There were circumstances in CSP versions older than 1.11.0 in which two containers with similar names could have shared the same keys wrongly.

However, there are some points to bear in mind concerning CSP versions 1.11.10 and later:

- If you want to share the same key between multiple computers, we supply the `cspimport` utility for transferring keys between containers.
- Any existing containers with older versions of the CSP must be migrated to the new format. We provide a utility, `cspmigrate`, to migrate containers from the old to the new system.

10.4.5. Utilities for the CAPI CSP

To help you migrate from Windows registry-based CSP container storage to the new CSP format, CSP version 1.11.0 and later provides you with a set of utilities. The new CSP format stores all information about a Security World in the Key Management Data directory. There are also utilities to manage the interfaces between the MSCAPI library and the module.

These utilities are:

Utility	Description
cspcheck	This utility checks that CSP container files are intact and uncorrupted, and also that referenced key files exist. Use cspcheck in conjunction with nfkmcheck, but run nfkmcheck first in order to test the integrity of your Security World files.
cspimport	This utility allows you to insert keys manually into existing CSP containers. This utility has two modes that either allow you to change a container's key association to that of an arbitrary Security World key or to copy CSP keys between containers.
cspmigrate	This utility moves the CSP container information from the registry into the Security World. If a new container already exists and has a key in it, and an identically-named old container exists with the same key, the utility asks you which key to keep. You can either: Enter -q to keep the new keys. Enter -f to overwrite new keys with old keys.
cspnvfix	Regenerate the NVRAM key counter area for a specified nShield CSP key.
csptest	Test the installed Cryptographic Service Providers.
csputils	This utility lists CSP containers and provides detailed information about them. It can also be used to delete container files if the current user has administrative privileges.
configure-csp-poolmode	The --mscapi option allows HSM Pool mode to be enabled or disabled for the nShield CAPI CSP without using the CSP wizard.
keytst	This utility displays information about existing CSP key containers by using the Microsoft CryptoAPI. If you have the appropriate permissions, keytst also allows you to create containers and their keys, as well as delete containers.

> ℹ️ Each of these commands has an -h option that displays the usage message for the command.

10.4.6. Uninstalling the CAPI CSP

To uninstall the CAPI CSP and unregister it as a cryptographic provider on your system, run the cngregister and cnginstall commands with the -U option. For more information, see Utilities for CNG.

10.5. Microsoft Cryptography API: Next Generation (CNG)

Cryptography API: Next Generation (CNG) is the successor to the Microsoft Cryptographic

API (CAPI) and its long-term replacement. CNG is designed to be extensible at many levels and cryptography agnostic in its behavior.

The Security World Software implementation of Microsoft CNG is supported on Microsoft Windows Server 2016 and later releases. The nShield CNG CSP provides the benefits of hardware-based encryption accessed through the standard Microsoft API, and supports the National Security Agency (NSA) classified Suite B algorithms.

10.5.1. Configuring the nShield CNG CSP

The DLL files that support the nShield CNG CSP are installed during product installation. However, you need to register the CNG CSP without removing the provider DLL files from your system.

You can unregister the nShield CNG CSP without removing the provider DLL files from your system. After unregistering, you can reregister the nShield CNG CSP, removing the files from your system. For more information, see Unregistering or reregistering the CNG CSP.

You can completely uninstall the nShield CNG CSP, removing the files from your system. After uninstalling, you must reinstall the files and then reregister the CNG CSP before you can use it. For more information, see Unregistering or reregistering the CNG CSP.

10.5.1.1. Registering the CNG CSP

You can register the nShield CNG CSP with:

- CNG Configuration Wizard
- The `cngregister` command-line utility

To register the nShield CNG CSP, the hardserver must be running and able to communicate with at least one module. This requirement is normally fulfilled during the product installation process. You can check that this requirement is fulfilled by running the `enquiry` command-line utility and checking the output for details about the module.

10.5.1.1.1. Registering the CNG CSP with the CNG Configuration Wizard

We recommend using the CNG Configuration Wizard to register the nShield CNG CSP. The product installation process places a shortcut to the CNG Configuration Wizard in the Windows Start menu: **Start > Entrust nShield Security World**.

> You can also use the CNG Configuration Wizard to load existing Security Worlds, see Adding an HSM to a Security World with the CSP or

> CNG wizard, generate new OCSs, see Creating an Operator Card Set
> with the CSP or CNG wizard, and configure the set-up parameters of
> the CNG CSP including HSM Pool mode.

With module firmware version 2.65.2 or later, if your application only uses module protected keys, you can use HSM Pool mode with multiple hardware security modules. HSM Pool mode exposes a single pool of HSMs and supports returning or adding a hardware security module to the pool without restarting the system. With a FIPS 140 Level 3 Security World, keys cannot be created in HSM Pool mode, however keys created outside HSM Pool mode can be used in HSM Pool mode.

To register the CNG CSP with the CNG Configuration Wizard, you must have already created a Security World and chosen a key protection method, either module-protection or OCS-protection. If you chose OCS-protection, you must also have already created an OCS before you can register the nShield CNG CSP with the CNG Configuration Wizard.

> 🛈 The CNG Configuration Wizard is not suitable for creating complex
> Security World setups or for creating Security Worlds with the unit.
> When creating such Security Worlds, or if you require more flexibility
> than the CNG configuration wizard provides, we recommend following
> the instructions in Creating a Security World using new-world.

If you use the CNG Configuration Wizard to create a Security World (and, if appropriate, an OCS), the wizard automatically prompts you to register the CNG CSP after you have fulfilled the necessary prerequisites.

You can also use the CNG Configuration Wizard to change an existing configuration at any time by running the wizard as usual and choosing the **Use the existing security world** option on the **Initial setup** screen.

To register the CNG CSP with the CNG Configuration Wizard after the necessary key-protection prerequisites have been fulfilled:

1. If the wizard is not already running:

    a. Run the wizard by double-clicking its shortcut in the Windows Start menu: **Start > Entrust nShield Security World**.

        The wizard displays the welcome window.

    b. Click the **Next** button.

        The wizard allows you to configure HSM Pool mode for CNG.

    c. Click the **Next** button.

If the prerequisite to create a Security World has been fulfilled, the wizard displays a confirmation screen.

d. Click the **Next** button.

The wizard displays a screen confirming that your Security World and (if you chose to create an OCS) an OCS have been created.

> **ⓘ** If you chose module-protection for your keys, the wizard does not confirm that an OCS has been created.

2. When the wizard has confirmed that it is ready to register the nShield CNG providers, click the **Next** button.

The wizard registers the nShield CNG CSP.

> **ⓘ** You cannot use the CNG Configuration Wizard to configure the nShield CNG providers for use as defaults. We recommend that you always use the nShield CNG providers by selecting them directly with the application that is using CNG.

When configuration of your nShield CNG CSP is complete, the wizard displays a confirmation screen.

10.5.1.1.2. Registering the CNG CSP with cngregister

You can use the `cngregister` command-line utility to register the nShield CNG CSP manually even if you have not already created a Security World (or, if you choose OCS-protection for your keys, even if you have not already created an OCS).

To register the nShield CNG CSP with the `cngregister` command-line utility, run the command without specifying any options:

```
cngregister
```

> **ⓘ** You cannot use the `cngregister` command-line utility to configure the nShield CNG providers for use as defaults. We recommend that you always use the cngregister command-line utility, see cngregister.

10.5.1.2. Unregistering or reregistering the CNG CSP

You can use the `cngregister` command-line utility to unregister or reregister the nShield CNG CSP manually.

To unregister the nShield CNG CSP, run the command:

```
cngregister -U
```

This command unregisters the CNG CSP, but does not remove the provider DLL files from your system. For information about removing these files, see Uninstalling or reinstalling the CNG CSP.

> **(!)** If any applications or services are using the nShield CNG providers for key storage or cryptography, unregistering the CNG CSP, you can register it at any time as long as the files have not been uninstalled from your system.

After unregistering the nShield CNG CSP, you can reregister it at any time as long as the files have not been uninstalled from your system. To reregister the nShield CNG CSP on your system, run the command:

```
cngregister
```

> **(i)** You cannot use the `cngregister` command-line utility to configure the nShield CNG providers for use as defaults. We recommend that you always use the nShield CNG providers by selecting them directly with the application that is using CNG.

For more information about these command-line utilities, see Utilities for CNG.

10.5.1.3. Uninstalling or reinstalling the CNG CSP

To uninstall the nShield CNG CSP:

1. To remove any and all dependencies that you have set, run the command:

   ```
   ncsvcdep -x
   ```

   > **(i)** Always run `ncsvcdep` as a user with full administrative privileges.

2. Unregister the nShield CNG CSP on your system by running the command:

   ```
   cngregister -U
   ```

   This command unregisters the CNG CSP, but does not remove the provider DLL files

from your system.

3. Uninstall the nShield CNG DLLs from your system:

   ◦ On 32-bit versions of Windows, run the command:

   ```
   cnginstall32 -U
   ```

   ◦ On 64-bit versions of Windows, run the command:

   ```
   cnginstall -U
   ```

To reinstall the nShield CNG CSP after you have previously uninstalled it:

1. Reinstall the nShield CNG CSP files on your system:

   ◦ On 32-bit versions of Windows, run the command:

   ```
   cnginstall32 -i
   ```

   ◦ On 64-bit versions of Windows, run the command:

   ```
   cnginstall -i
   ```

2. Reregister the nShield CNG CSP on your system by running the command:

   ```
   cngregister
   ```

For more information about these command-line utilities, see Utilities for CNG

10.5.2. Supported algorithms for CNG

This section lists the National Security Agency (NSA) classified Suite B algorithms supported by the nShield CNG providers.

> 🛈 | The MQV algorithm is not supported by the nShield CNG providers.

10.5.2.1. Signature interfaces (key signing)

Interface name	Type of support
RSA PKCS#1 v1	Hardware
RSA PSS	Hardware

Interface name	Type of support
DSA	Hardware
ECDSA_P224	Hardware
ECDSA_P256	Hardware
ECDSA_P384	Hardware
ECDSA_P521	Hardware

> ℹ️ Hashes used with ECDSA must be of the same length or shorter than the curve itself. If you attempt to use a hash longer than the curve the operation returns NOT_SUPPORTED. In FIPS 140 Level 3 Security Worlds ECDSA signing is only supported where the length of the curve is approximately the length of the hash.

10.5.2.2. Hashes

Hash name	Type of support
SHA1	Hardware (HMAC only)/software
SHA256	Hardware (HMAC only)/software
SHA384	Hardware (HMAC only)/software
SHA512	Hardware (HMAC only)/software
SHA224	Hardware (HMAC only, requires firmware version 2.33.60 or later)/software
MD5	Hardware (HMAC only)/software

> ℹ️ MD5 is not supported in FIPS 140 mode.

10.5.2.3. Asymmetric encryption

Algorithm name	Type of support
RSA Raw (NCRYPT_NO_PADDING_FLAG)	Hardware
RSA PKCS#1 v1 (NCRYPT_PAD_PKCS1_FLAG)	Hardware
RSA OAEP (NCRYPT_PAD_OAEP_FLAG)	Hardware

10.5.2.4. Symmetric encryption

Algorithm name	Type of support
RC4	Hardware and Software (not supported in FIPS 140 Level 3 mode)
AES ECB,CBC	Hardware and Software
DES ECB,CBC	Hardware and Software (DES is not supported in FIPS 140 Level 3 mode)
3DES ECB,CBC	Hardware and Software
3DES_112 ECB,CBC	Hardware and Software

10.5.2.5. Key exchange

Protocol name	Type of support
DH	Hardware
ECDH_P224	Hardware
ECDH_P256	Hardware
ECDH_P348	Hardware
ECDH_P521	Hardware

> ℹ️ Elliptic curve cryptography algorithms must be enabled before use. Use the `fet` command-line utility with an appropriate certificate to enable a purchased feature. If you enable the elliptic curve feature on your modules after you first register the CNG providers, you must run the configuration wizard again for the elliptic curve algorithm providers to be registered.

10.5.2.6. Random Number Generation

Name	Type of support
RNG	Hardware

10.5.3. Migrating keys for CNG

We provide functionality for migrating existing keys from other providers into the Security World Key Storage Provider. To identify installed providers, run the command:

```
cnglist --list-providers
```

To identify the keys that are available from a particular provider, run the command:

```
cnglist --list-keys --provider="ProviderName"
```

In this command, *ProviderName* is the name of the provider. The following command provides an example of identifying keys from the *Security World Key Storage Provider*:

```
cnglist --list-keys --provider="nCipher Security World Key Storage Provider"
MyApp Personal Data Key: RSA
CertReq-5eb45f6d-6798-472f-b668-288bc5d961da: ECDSA_P256 machine
WebServer Signing Key: DSA machine
ADCS-Root-Key: ECDSA_P521 machine
```

> **i** To list the keys available from the *Security World Key Storage Provider*, run the command `cnglist --list-keys` (without specifying the `--provider` option).

10.5.3.1. Importing a Microsoft CAPI key into the Security World Key Storage Provider

To import a Microsoft CAPI key into the Security World Key Storage Provider, first run the CAPI utility `csputils` to identify the existing CAPI containers and their key contents.

CAPI containers can contain either a signing key or a key exchange key, or both. The following example shows how to import both a signing key and a key exchange key from a Microsoft CAPI container:

```
cngimport -m --csp="Microsoft Strong Cryptographic Provider"
  -k "EXAMPLE_CAPICONTAINER"
    "EXAMPLE_IMPORTED_SIGNATURE_CAPICONTAINER"
    "EXAMPLE_IMPORTED_KEYEXCHANGE_CAPICONTAINER"
```

To check the success of the import, list the keys present in the *Security World Key Storage Provider*:

```
cnglist --list-keys
EXAMPLE_IMPORTED_SIGNATURE_CAPICONTAINER: RSA
EXAMPLE_IMPORTED_KEYEXCHANGE_CAPICONTAINER: DH
```

The following example command shows how to import a single signing key:

```
cngimport -m -s --csp="Microsoft Strong Cryptographic Provider"
  --key="EXAMPLE_CAPICONTAINER"
       "EXAMPLE_IMPORTED_SIGNATURE_ONLY_CAPICONTAINER"
```

Run the `cnglist` command with the `--list-keys` option to check the success of the key import:

```
cnglist --list-keys
EXAMPLE_IMPORTED_SIGNATURE_ONLY_CAPICONTAINER: RSA
```

> **ℹ** The `cngimport` option `-m/--migrate` cannot be used to migrate nShield CAPI container keys to CNG. For information about importing nShield CAPI container keys into CNG, see Importing a Microsoft CNG key into the Security World Key Storage Provider.

10.5.3.2. Importing a Microsoft CNG key into the Security World Key Storage Provider

To import a Microsoft CNG key into the Security World Key Storage Provider, run the `cngimport` command as shown in the following example:

```
cngimport -m
  -k "EXAMPLE_RSA_1024"
     "IMPORTED_RSA_1024"
```

Run the `cnglist` command with the `--list-keys` option to check the success of the key import:

```
cnglist --list-keys
  IMPORTED_RSA_1024: RSA
```

The original key is not deleted from the provider from which it was imported:

```
cnglist --list-keys --provider="Microsoft Software Key Storage Provider"
  EXAMPLE_RSA_1024
```

> **ℹ** Certain applications, such as Certificate Services, create keys using the Microsoft Software Key Storage Provider which cannot be exported. Attempting to import such a key into the nShield provider results in the following message:

```
cngimport -m -k WIN-KQ1Z6JMCUTB-CA WIN-ncipher-CA
Unable to continue.
This key can not be exported from Microsoft Software Key Storage Provider.
```

10.5.3.3. Importing a Security World key into the Security World Key Storage Provider

To import a Security World key into the Security World Key Storage Provider, run the `cngim-`

`port` utility as shown in the following example:

```
cngimport --import --key=nfkmsimple1 --appname=simple nfkmsimple1
Found key 'nfkmsimple1'
Importing NFKM key.. done
```

Run `cnglist` with the `--list-keys` option to confirm that the key has been successfully imported:

```
cnglist --list-keys
nfkmsimple1: RSA
```

To import an nShield CAPI container into the Security World Key Storage Provider, run the `csputils` command to identify the container name:

```
csputils -l
File ID     Container name        Container owner       DLL name    S X
=========   ==================    ==================    =========   = =
31e994f07   CONTAINER2            SYWELL\Administrato   ncsp        * *
3a2b082a8   CAPICONTAINER         SYWELL\Administrato   ncsp        * *
2 containers and 4 keys found.
```

> ℹ️ Run the `csputils` command with the `-l` and `-m` options to migrate an nShield CAPI machine container.

Identify the Security World key names of the keys in the container by running the `csputils` command as follows:

```
csputils -d -n CAPICONTAINER
Detailed report for container ID #3a2b082a8f2ee1a5acb756d5e95b09817072807a
Filename:       key_mscapi_container-3a2b082a8f2ee1a5acb756d5e95b09817072807a
Container name: CAPICONTAINER
User name:      SYWELL\Administrator
User SID:       s-1-5-21-352906761-2625708315-3490211485-500
CSP DLL name:   ncsp.dll
Filename for signature key is key_mscapi_ce51a0ee0ea164b993d1edcbf639f2be62c53222
  Key was generated by the CSP
  Key hash:    ce51a0ee0ea164b993d1edcbf639f2be62c53222
  Key is recoverable.
Key is cardset protected.
Cardset name:            nopin
     Sharing parameters:     1 of 1 shares required.
Cardset hash:            d45b30e7b60cb226f5ade5b54f536bc1cc465fa4
     Cardset is non-persistent.
Filename for key exchange key is key_mscapi_dbd84e8155e144c59cf8797d16e7f8bd19ac446a
  Key was generated by the CSP
  Key hash:    dbd84e8155e144c59cf8797d16e7f8bd19ac446a
  Key is recoverable.
Key is cardset protected.
Cardset name:            nopin
     Sharing parameters:     1 of 1 shares required.
Cardset hash:            d45b30e7b60cb226f5ade5b54f536bc1cc465fa4
     Cardset is non-persistent.
1 container and 2 keys found.
```

The key name to pass to the `cngimport` command `--key` option is the part of the key name that follows `key_mscapi_` in the output line that starts `Filename for signature key is key_mscapi_`.

For example, the signature key file name for `CAPICONTAINER` in the example shown above is `key_mscapi_ce51a0ee0ea164b993d1edcbf639f2be62c53222`, so `ce51a0ee0ea164b993d1ed-cbf639f2be62c53222` is the key name that should be passed to `cngimport`:

```
cngimport --import --key="ce51a0ee0ea164b993d1edcbf639f2be62c53222" --
appname="mscapi" Signature_Key_Imported_From_nCipher_CAPI
Found unnamed key
Importing NFKM key.. done
```

Run `cnglist` with the `--list-keys` option to confirm that the key has been successfully imported:

```
cnglist --list-keys
Signature_Key_Imported_From_nCipher_CAPI: RSA
cngsoak: ECDH_P256
```

Follow the same procedure for importing the key exchange key from the nShield CAPI container.

10.5.4. Using CAPI keys in CNG

We now provide the capability to use keys generated by CAPI in CNG applications. This is provided through the standard `NCryptOpenKey` CNG API call. Passing either `AT_SIGNATURE` or `AT_KEYEXCHANGE` as the `dwLegacyKeySpec` parameter and the CAPI container name as the `pszKeyName` parameter will invoke this mode of operation. The CAPI key will be loaded into the CNG provider and will behave as if it was a CNG key. Any key authorization required will be handled with a user interface being invoked to prompt the application user to insert the smart card or enter appropriate passphrases. There is support for Key Usage and Key Count ing properties.

The CNG application has to be written such that it calls `NCryptOpenKey` to open a CAPI key explicitly.

10.5.5. Utilities for CNG

Use the `nfkmverify` command-line utility to check the security of all stored keys in the Secu rity World. Use `nfkminfo`, `nfkmcheck`, and other command-line utilities to assist in this process.

The following table lists the utilities specific to the nShield CNG CSP:

x86	x64	Utility description
`cngimport32.exe`	`cngimport.exe`	This key migration utility is used to migrate Security World, CAPI, and CNG keys to the Security World Key Storage Provider.
`cnginstall32.exe`	`cnginstall.exe`	This utility is the nShield CNG CSP installer. Only use this utility to remove or reinstall the provider DLLs and associated registry entries manually.
`cnglist32.exe`	`cnglist.exe`	This utility lists information about CNG CSP.
`cngregister32.exe`	`cngregister.exe`	This is the nShield CNG CSP registration utility. You can use it to unregister and re-register the CNG CSP.
`ncsvcdep32.exe`	`ncsvcdep.exe`	This utility is the service dependency tool. You can configure some service based applications, such as Microsoft Certificate Services and IIS, to use the nShield CNG CSP.

These utilities are located in the `bin` directory of your Security World Software installation (for example, `%NFAST_HOME%\bin`).

> On 64-bit versions of Windows, both the 32-bit and 64-bit versions of the listed utilities are installed. When working on an 64-bit version of Windows, always ensure that you use the 64-bit version of the utility (if one is available).

10.5.5.1. cngimport

Use `cngimport` to migrate keys to the Security World Key Storage Provider. For more information, see Migrating keys for CNG.

10.5.5.2. cnginstall

The `cnginstall` utility is used by the Security World Software installation wizard. You can also use this utility to manually uninstall (or reinstall) the nShield CNG DLLs and registry entries.

To uninstall the nShield CNG DLL files, run the command:

```
cnginstall -U
```

This command removes the provider DLL files from your system. It produces output of the form:

```
nckspppt.dll removed.
nckspsw.dll removed.
ncpp.dll removed.
```

Before you uninstall the nShield CNG DLL files, ensure that you unregister the CNG CSP. For more information, see:

- cngregister

- Unregistering or reregistering the CNG CSP

After unregistering the nShield CNG CSP, you can reregister it at any time as long as the files have not been uninstallted from your system. To reregister the nShield CNG CSP on your system, run the command:

```
cngregister
```

For more information about uninstalling and reinstalling the nShield CNG CSP with `cnginstall`, see Uninstalling or reinstalling the CNG CSP.

10.5.5.3. cngregister

Use `cngregister` to unregister the nShield CNG CSP manually.

To unregister the nShield CNG CSP, run the command:

```
cngregister -U
```

This command produces output for the form:

```
Unregistered provider 'nCipher Primitive Provider'
Unregistered provider 'nCipher Security World Key Storage Provider'
```

This command unregisters the CNG CSP, but does not remove the provider DLL files from your system. For information about removing these files, see:

- cnginstall

- Uninstalling or reinstalling the CNG CSP.

> ! If any applications or services are using the nShield CNG CSP for key storage or cryptography, unregistering it can cause system instability.

After unregistering the nShield CNG CSP, you can reregister it at any time as long as the files have not been uninstalled from your system. To reregister the nShield CNG CSP on your system, run the command:

```
cngregister
```

> **ℹ** You cannot use the `cngregister` command-line utility to configure the nShield CNG providers for use as defaults. We recommend that you always use the nShield CNG providers by selecting them directly with the application that is using CNG.

10.5.5.4. cngsoak

Use `cngsoak` to obtain statistics about the performance of the nShield CNG CSP. Specifically, use `cngsoak` to determine the speed of:

- Signing a hash (`cngsoak --sign`)
- encryption (`cngsoak --encrypt`)
- key exchange (`cngsoak --keyx`)
- key generation (`cngsoak --generate`).

The output from `cngsoak` displays information as columns of information. From left to right, these columns display:

- The time in second that `cngsoak` has been running
- the total number of operations completed
- the number of operations completed in last second
- the average number of operations completed each second.

10.5.5.5. ncsvcdep

Use the `ncsvcdep` utility to ensure that the nShield `nFast Server` service is running before certain services are enabled. For example, Active Directory Certificate Services or Internet Information Services require that the hardserver is running in order to use the nShield CNG CSP. Failure to set this dependency can lead to system instability.

To list installed services, run the `ncsvcdep` command with the `-l` option:

```
ncsvcdep -l
```

Output from this command has the form:

```
Installed Services (Count - "Display Name" - "Service Name")
0 - "Application Experience" - "AeLookupSvc"
1 - "Application Layer Gateway Service" - "ALG"
2 - "Application Information" - "Appinfo"
3 - "Application Management" - "AppMgmt"
4 - "Windows Audio Endpoint Builder" - "AudioEndpointBuilder"
.
.
.
108 - "nFast Server" - "nFast Server"
109 - "Active Directory Certificate Services" - "CertSvc"
```

ℹ | Always run `ncsvcdep` as a user with full administrative privileges.

To set a dependency, run the command:

```
ncsvcdep -a "DependentService"
```

In this command, *DependentService* is the service that has the dependency. The following example shows how to make the Active Directory Certificate Services dependent on the nFast Server:

```
ncsvcdep -a "CertSvc"
Dependency change succeeded.
```

To remove a specific dependency relationship, run `ncsvcdep` with the `-r` option, for example:

```
ncsvcdep -r "CertSvc"
Dependency change succeeded.
```

To remove all dependencies, run `ncsvcdep` with the `-x` option:

```
ncsvcdep -x
```

ℹ | Microsoft Certificate Services require that the `certsvc` service is made dependent on the hardserver.

ℹ | Microsoft Internet Information Services require that the `http` service is made dependent on the hardserver.

10.5.5.6. cnglist

Use `cnglist` to display details of CNG providers, keys, and algorithms.

To list details of the CNG providers, run the `cnglist` command with the `--list-providers`

---

option:

```
cnglist --list-providers
```

Output from this command is of the form:

```
Microsoft Primitive Provider
Microsoft Smart Card Key Storage Provider
Microsoft Software Key Storage Provider
Microsoft SSL Protocol Provider
nCipher Primitive Provider
nCipher Security World Key Storage Provider
```

To list details of the algorithms, run the `cnglist` command with the `--list-algorithms` option:

```
cnglist --list-algorithms
```

Output from this command has the form:

```
BCryptEnumAlgorithms(BCRYPT_CIPHER_OPERATION):
  Name                       Class      Flags
  AES                        0x00000001 0x0
  RC4                        0x00000001 0x0
  DES                        0x00000001 0x0
  DESX                       0x00000001 0x0
  3DES                       0x00000001 0x0
  3DES_112                   0x00000001 0x0
BCryptEnumAlgorithms(BCRYPT_HASH_OPERATION):
  Name                       Class      Flags
  SHA1                       0x00000002 0x0
  MD2                        0x00000002 0x0
  MD4                        0x00000002 0x0
  MD5                        0x00000002 0x0
  SHA256                     0x00000002 0x0
  SHA384                     0x00000002 0x0
  SHA512                     0x00000002 0x0
  AES-GMAC                   0x00000002 0x0
  SHA224                     0x00000002 0x0
BCryptEnumAlgorithms(BCRYPT_ASYMMETRIC_ENCRYPTION_OPERATION):
  Name                       Class      Flags
  RSA                        0x00000003 0x0
```

To list details of the algorithms for the Security World Key Storage Provider, run the `cnglist` command with the `--list-algorithms`, `--keystorage`, and `--nc` options:

```
cnglist --list-algorithms --keystorage --nc
```

Output from this command has the form:

```
NCryptEnumAlgorithms(NCRYPT_CIPHER_OPERATION) no supported algorithms
NCryptEnumAlgorithms(NCRYPT_HASH_OPERATION) no supported algorithms
```

```
NCryptEnumAlgorithms(NCRYPT_ASYMMETRIC_ENCRYPTION_OPERATION):
  Name                           Class     Operations Flags
  RSA                            0x00000003 0x00000014 0x0
NCryptEnumAlgorithms(NCRYPT_SECRET_AGREEMENT_OPERATION):
  Name                           Class     Operations Flags
  DH                             0x00000004 0x00000008 0x0
  ECDH_P224                      0x00000004 0x00000008 0x0
  ECDH_P256                      0x00000004 0x00000008 0x0
  ECDH_P384                      0x00000004 0x00000008 0x0
  ECDH_P521                      0x00000004 0x00000008 0x0
NCryptEnumAlgorithms(NCRYPT_SIGNATURE_OPERATION):
  Name                           Class     Operations Flags
  RSA                            0x00000003 0x00000014 0x0
  DSA                            0x00000005 0x00000010 0x0
  ECDSA_P224                     0x00000005 0x00000010 0x0
  ECDSA_P256                     0x00000005 0x00000010 0x0
  ECDSA_P384                     0x00000005 0x00000010 0x0
  ECDSA_P521                     0x00000005 0x00000010 0x0
```

To list details of the algorithms for a specific named key storage provider, run the `cnglist` command with the `--list-algorithms` and `--provider`="*ProviderName*" options:

```
cnglist --list-algorithms --provider="Microsoft Software Key Storage Provider"
```

Output from this command has the form:

```
Microsoft Software Key Storage Provider
NCryptEnumAlgorithms(NCRYPT_CIPHER_OPERATION) no supported algorithms
NCryptEnumAlgorithms(NCRYPT_HASH_OPERATION) no supported algorithms
NCryptEnumAlgorithms(NCRYPT_ASYMMETRIC_ENCRYPTION_OPERATION):
  Name                           Class     Operations Flags
  RSA                            0x00000003 0x00000014 0x0
NCryptEnumAlgorithms(NCRYPT_SECRET_AGREEMENT_OPERATION):
  Name                           Class     Operations Flags
  DH                             0x00000004 0x00000008 0x0
  ECDH_P256                      0x00000004 0x00000018 0x0
  ECDH_P384                      0x00000004 0x00000018 0x0
  ECDH_P521                      0x00000004 0x00000018 0x0
NCryptEnumAlgorithms(NCRYPT_SIGNATURE_OPERATION):
  Name                           Class     Operations Flags
  RSA                            0x00000003 0x00000014 0x0
  DSA                            0x00000005 0x00000010 0x0
  ECDSA_P256                     0x00000005 0x00000010 0x0
  ECDSA_P384                     0x00000005 0x00000010 0x0
  ECDSA_P521                     0x00000005 0x00000010 0x0
```

10.5.5.6.1. configure-csp-poolmode

The `configure-csp-poolmode` utility allows HSM Pool mode to be enabled or disabled for the nShield CNG CSP without using the CNG wizard.

To enable HSM Pool mode for CNG run the command:

```
configure-csp-poolmode --cng --enable
```

To disable HSM Pool mode for CNG run the command:

```
configure-csp-poolmode --cng --disable
```

To remove HSM Pool mode setting for CNG from the registry, use the command:

```
configure-csp-poolmode --cng --remove
```

11. Working with CodeSafe

11.1. CodeSafe applications

To run CodeSafe applications on your system, you must have enabled the Secure Execution Engine (SEE) by purchasing and enabling an appropriate SEE activation licence as described in Enabling optional features.

You must also have loaded a valid certificate, known as a CodeSafe developer ID certificate, that can be used to verify the signature of each CodeSafe application that you want to run. These certificates will be supplied by the developer of the CodeSafe application and should be delivered to you together with the application.

CodeSafe developer ID certificates have a limited lifetime. If a certificate expires, applications that are already running will be unaffected but it will not be possible to start an application until a new certificate is loaded.

For this reason you should keep track of the expiry dates of your certificates and request a new certificate from the CodeSafe developer before your certificate expires.

You can view the expiry date of certificates by using `csadmin ids list` and looking at the `notAfter` field for each certificate.

New certificates can be loaded whilst an application is running without interrupting service.

If you want to develop your own CodeSafe applications, you must also purchase the CodeSafe developer kit.

An SEE application is typically a standalone SEE machine that is loaded automatically by the hardserver (for example, a CodeSafe C application).

Check the documentation that your CodeSafe application vendor supplies for information about how to set up and use the application, as well as for any other installation and configuration information.

CodeSafe applications are standalone applications, but each CodeSafe C application can consist of multiple parts, and its installation can include several configuration steps. For instructions on installing and configuring each application, see your application vendor's documentation.

11.2. CodeSafe setup for the nShield 5c

To use CodeSafe applications with the nShield 5c you must exchange launcher service keys

between the client and the nShield 5c by following the procedures in this section.

To use multiple nShield 5c modules simultaneously from the same client machine, see Configure client to access multiple nShield 5c HSMs.

11.2.1. Initial Client Configuration: Retrieve or generate the client's public key

If the client's launcher key pair has already been generated, view the client's key type and public key data by issuing the following command:

```
hsmadmin cs5 clientinfo
```

The following is an example of the output expected for `clientinfo`:

```
ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBDmcFY2wxtNj+di478pvOK3czNeXJoqw+itu/Hw6KRcGjDO60Zqr8eaWJcDS7
9DHncvdUZv0ZPbLqTCnNjN0ECQ=
```

> **i** You must copy the key type public key. You will need these in a later step for setting the client info in the nShield 5c CLI.

If the client does not already have a launcher key pair, generate a new one on the client machine with the following command:

```
hsmadmin cs5 genclientinfo
```

And then view and copy the key type and public key as in the previous step.

11.2.2. Configure nShield5c for Client Authorization:

> **i** To configure the nShield 5c to deploy CodeSafe applications, you need access to its serial console. For more information, see Configuring an nShield HSM using the Serial Console.

1. Put the module into Maintenance mode (this will temporarily stop nCoreapi and launcher services):

```
(cli)maintmode enable
```

> **i** The module will not be able to perform any tasks and will appear

| unavailable until it is put back into Operational mode.

2. Remove any existing client information in the nShield 5c to ensure a clean setup:

```
(cli)cs5 clearclientinfo
```

3. Set the client's public key for the launcher service by running the following command :

```
(cli)cs5 setclientinfo <key_type> <public_key>
```

Where `<key_type>` and `<public_key>` are the output from the `hsmadmin cs5 clientinfo` command that you copied earlier. For example:

```
(cli)cs5 setclientinfo ecdsa-sha2-nistp256
ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ012345678
9ABCDEFGHIJKLMNOPQRSTUVWXY
```

4. Enable ports for Codesafe 5, using the nShield 5c's IP address:

```
(cli)cs5 ports enable <nshield-5c-ip-address>
```

Where `<nshield-5c-ip-address>` is the address on the HSM's network interface that CodeSafe 5 will access.

> 🛈 | This prevents KeySafe 5 from being used until `cs5 ports disable` is run.

5. Retrieve all the information the client needs to connect to the remote launcher service (ESN, IP address, and launcher service public key):

```
(cli)cs5 getserverinfo
```

Copy the ESN, IP address, and launcher service public key for use in the final client con figuration. The following is an example of the output expected for `getserverinfo`:

```
Codesafe 5 server info in format <ESN> <IP Addresss> <Port> <Key Type> <Public Key>
B0A7-B88C-D3F4 10.194.147.134 2205 ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBOEJMTeBR1546vmQ4YwT7WGlUv57PyCIShw66c9BAuD+WP5lrE8yjwJ
R/aBo7hrBxV95Vlgi6+4Up4fDI0uqO6s=
```

6. Put the module back into Operational mode:

```
(cli)maintmode disable
```

11.2.3. Finalize the Client Configuration:

On the client machine, enroll the nShield 5c for CodeSafe so that the client can connect to the remote launcher service with the following command:

```
hsmadmin cs5 enroll <ESN> <IP_Address> <Port> <Key_Type> <Public_Key>
```

Where `ESN`, `IP_Address`, `Port`, `Key_Type` and `Public_key` is the server information in the output for the `cs5 getserverinfo` command that you copied earlier.

You can verify the setup by running the following command:

```
csadmin list
```

The following is an example of the output expected when running `csadmin list` with no SEE machines installed:

```
B0A7-B88C-D3F4        No SEE machines currently installed
```

The nShield 5c is now ready for use with CodeSafe. For more information about building, signing, and running SEE machines with CodeSafe, see The csadmin utility tool.

11.2.4. Configure client to access multiple nShield 5c HSMs

To use multiple nShield 5c HSMs from the same client you should configure them with the same launcher client key. To do this repeat the procedures in Configure nShield5c for Client Authorization: for each HSM using the same `<key_type>` and `<public_key>` obtained from `cs5 clientinfo` each time.

The information used in Finalize the Client Configuration: will be unique to each nShield 5c HSM and will be that returned by the `cs5 getservinfo` command for the individual HSM.

11.2.5. Removal of client configuration

When a client machine is finished with an nShield 5c, disconnect from it using:

```
hsmadmin cs5 unenroll <ESN>
```

Where `<ESN>` is the ESN of the HSM from which you wish to remove configuration.

Then disable ports for Codesafe 5:

```
(cli)cs5 ports disable
```

> **ℹ** If you disable the ports and try to run `csadmin list` on the client you will notice that it will be unable to connect to the launcher.

The following is an example of the output expected when the ports are disabled:

```
% csadmin list
B0A7-B88C-D3F4          CONNECTION ERROR: Unable to connect to 'launcher': [Errno None] Unable to connect to port
2205 on 10.194.147.134
Failed to get SEE machines list for module(s): ['B0A7-B88C-D3F4']
```

11.3. The csadmin utility tool

The `csadmin` tool is used to manage CodeSafe 5 applications.

> **ℹ** The `csadmin` tool covers CodeSafe application deployment from both the perspective of a CodeSafe application developer and a CodeSafe application user. The help text displays the complete set of commands available. This document lists the commands that are relevant to the CodeSafe application user.

You must be logged in as an Administrator or a user with local administrator rights to execute `csadmin` commands.

You must have `/opt/nfast/bin (Linux)` or `NFAST_HOME (Windows)` in your `PATH` environment variable to use `csadmin`.

To view the help text included here while using `csadmin`, run a command or sub-command with the `-h|--help` option.

> **ℹ** The following examples use a Linux machine for the deployment of CodeSafe applications. The same commands can be applied to a Windows machine.

```
$ csadmin --help
usage: csadmin [-h] [-v] {image,load,start,stop,list,destroy,sshd,ids,log,config,stats} ...

Tool to manage nShield 5s CodeSafe.

positional arguments:
  {image,load,start,stop,list,destroy,sshd,ids,log,config,stats}
    image               Perform tasks related to loadable nShield 5s CodeSafe images
    load                Load a CodeSafe 5 image onto an nShield 5s HSM
    start               Start a previously loaded image on an nShield 5s HSM
    stop                Stop an SEE machine running on an nShield 5s HSM
    list                List the SEE machines loaded on an nShield 5s HSM
```

```
    destroy           Destroy an SEE machine loaded on an nShield 5s HSM
    sshd              Manage the SSH daemon dedicated to a specific SEE machine on an nShield 5s HSM
    ids               Manage the CodeSafe Developer Authentication Certificates on an nShield 5s HSM
    log               Manage the SEE machine logging state
    config            Manage the SEE machine configuration
    stats             Manage the SEE machine statistics

options:
  -h, --help          show this help message and exit
  -v, --version       Print csadmin version
```

11.3.1. Manage loadable CodeSafe images

CodeSafe 5 images have the extension `.cs5` and are often referred to as "CS5 files". Operations such as generating, signing, and inspecting loadable CodeSafe images are handled by `csadmin image`:

```
$ csadmin image --help
usage: csadmin image [-h] {generate,sign,inspect} ...

positional arguments:
  {generate,sign,inspect}
    generate          generate a loadable CodeSafe image.
    sign              sign a loadable CodeSafe image. Note: images are signed with an HSM.
    inspect           inspect a loadable CodeSafe image's metadata.

options:
  -h, --help          show this help message and exit
```

11.3.1.1. Generate loadable images

Only applicable to CodeSafe developers. See the *CodeSafe 5 Developer Guide*

11.3.1.2. Inspect images

The `csadmin` tool provides the `inspect` operation so developers can view the details of a `.cs5` file.

The `inspect` operation usage:

```
csadmin image inspect --help
usage: csadmin image inspect [-h] [--json] CS5FILE

positional arguments:
  CS5FILE     The cs5 file being inspected

optional arguments:
  -h, --help  show this help message and exit
  --json      Show details in json format
```

For example:

```
$ csadmin image inspect myapp.cs5
Type            codesafe-container
Platform        nShield5
PackageName     MyCodeSafeApp
Version         1.0
EntryPoint      /home/see/launch_see.sh
Format          1
```

11.3.1.3. Sign images

Only applicable to CodeSafe developers. See the *CodeSafe 5 Developer Guide*

11.3.2. Load a CodeSafe application

> ❗ If you are using a third-party developed CodeSafe application, ensure that you trust the developer of the CodeSafe application, and verify that the image you are using has a genuine certificate issued by the trusted organization. You must have loaded a valid certificate prior to loading the application.

Loading a CodeSafe image onto the module creates the SEE container and the environment needed to run the application. This is done using `csadmin load`:

```
$ csadmin load --help

usage: csadmin load [-h] [--timeout TIMEOUT] [--esn ESN] [--verbose] [--json] CS5FILE

positional arguments:
  CS5FILE           The CodeSafe 5 image file to be loaded into the HSM(s).

optional arguments:
  -h, --help        show this help message and exit
  --timeout TIMEOUT Time to wait for service response, in seconds. Default: 30s, must be between 3s and 120s
  --esn ESN         ESN of the HSM to load CodeSafe 5 image into
  --verbose         Print verbose logs
  --json            Show result output in json format
```

On success, the `load` sub-command returns the UUID of the new SEE container, for example:

```
$ csadmin load test_app_1.cs5

FEDC-BA09-8765        SUCCESS
UUID: c2e75658-04aa-4634-a7f8-fad3692dd9d7

$ csadmin load --json test_app_1.cs5
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "uuid": "fedcba09-8765-4321-1234-567890abcdef"
```

```
            }
        }
    }
```

11.3.3. Start a CodeSafe application

> **ℹ** Although multiple SEE machines can be loaded and started at the same time, only one can talk to the `ncoreapi` service.

> **ℹ** A SEE machine will only start if a valid certificate is already loaded. Certificates have an expiry date and if the certificate has expired it will not be possible to start the SEE machine until a new certificate is loaded.

Starting an SEE machine previously loaded onto the HSM launches the container and executes the predetermined entry point. This is done using `csadmin start`:

```
$ csadmin start --help
usage: csadmin start [-h] [--timeout TIMEOUT] -u UUID [--esn ESN] [--verbose] [--json]

optional arguments:
  -h, --help            show this help message and exit
  --timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 120s
  -u UUID, --uuid UUID  The UUID of the machine being started.
  --esn ESN             ESN of the HSM hosting the SEE machine to be started.
  --verbose             Print verbose logs
  --json                Show response in json format
```

The `start` command takes as main input the UUID of the machine to be started. On success, it returns the IP(v6) address of the container, for example:

```
$ csadmin start --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765         SUCCESS
IP ADDRESS: ffff::fff:ffff:ffff:ffff

$ csadmin start --uuid fedcba09-8765-4321-1234-567890abcdef --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "ip": "ffff::fff:ffff:ffff:ffff"
        }
    }
}
```

To set the container to start automatically at startup and when the nShield 5 is reset, see Use set to enable autostart.

11.3.4. Stop a CodeSafe application

Stopping an SEE machine previously started on the HSM will gracefully terminate its pro-
grams execution and stop the container. This is done using `csadmin stop`. To display the
sub-command usage, execute it with the `--help` option as shown below:

```
$ csadmin stop --help
usage: csadmin stop [-h] [--timeout TIMEOUT] -u UUID [--esn ESN] [--verbose] [--json]

optional arguments:
  -h, --help            show this help message and exit
  --timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 120s
  -u UUID, --uuid UUID  The UUID of the machine being stopped.
  --esn ESN             ESN of the HSM hosting the SEE machine to be stopped.
  --verbose             Print verbose logs
  --json                Show response in json format
```

The `stop` command takes as main input the UUID of the machine to be stopped, for exam-
ple:

```
$ csadmin stop --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765         SUCCESS

$ csadmin stop --json --uuid fedcba09-8765-4321-1234-567890abcdef
{
    "FEDC-BA09-8765": {
        "succeeded": true
    }
}
```

11.3.5. Destroy a CodeSafe application

Destroying an SEE machine that is currently stopped on the HSM removes the application
container and any user data created during the life of the machine. This is done using `csad-
min destroy`:

```
$ csadmin destroy --help
usage: csadmin destroy [-h] [--timeout TIMEOUT] -u UUID [--esn ESN] [--verbose] [--json]

options:
  -h, --help            show this help message and exit
  --timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
  -u UUID, --uuid UUID  The UUID of the machine being destroyed.
  --esn ESN             ESN of the HSM hosting the SEE machine to be destroyed.
  --verbose             Print verbose logs
  --json                Show response in json format
```

11.3.6. List CodeSafe applications

`csadmin list` displays a list of the SEE machines currently loaded on the HSM:

```
$ csadmin list --help
usage: csadmin list [-h] [--timeout TIMEOUT] [--esn ESN] [--verbose] [--json]
```

```
optional arguments:
  -h, --help          show this help message and exit
  --timeout TIMEOUT   Time to wait for service response, in seconds. Default: 30s, must be between 3s and 120s
  --esn ESN           ESN of the HSM to retrieve SEE machines list for. Can be specified multiple times for
multiple HSMs.
  --verbose           Print verbose logs
  --json              Show response in json format
```

For example:

```
$ csadmin list
FEDC-BA09-8765
UUID                               State    Name

-------------------------------------------------------------------
12345678-90ab-cdef-fedc-ba0987654321    STOPPED   nShield Docs Server
fedcba09-8765-4321-1234-567890abcdef    RUNNING   CodeSafe App 1
09876543-21fe-dcba-abcd-ef1234567890    RUNNING   CodeSafe App 2
```

11.3.7. Configure CodeSafe 5

`csadmin config` allows the user to set and query the state of CodeSafe's auto-start and logging features:

```
$ csadmin config --help
usage: csadmin config [-h] {list,get,set} ...

positional arguments:
  {list,get,set}
    list          List CodeSafe machine configuration.
    get           Get CodeSafe machine configuration.
    set           Set CodeSafe machine configuration.

options:
  -h, --help      show this help message and exit

configuration keys:
  "autostart": "enabled" or "disabled". Automatically start the SEE machine when the module resets. Disabled by
default.
  "log": "enabled" or "disabled". Capture a container's SEE logs. Disabled by default.
```

You must use a positional argument (`list`, `get`, `set`) with the `config` command.

11.3.7.1. List machine configuration

`csadmin config list` lists the current configuration keys and values of a given machine:

```
$ csadmin config list --help
usage: csadmin config list [-h] -u UUID [--esn ESN] [--verbose] [--json] [--timeout TIMEOUT]

options:
  -h, --help             show this help message and exit
  -u UUID, --uuid UUID   The UUID of the machine.
  --esn ESN              ESN of the HSM hosting the SEE Machine.
```

```
    --verbose             Print verbose logs
    --json                Show response in json format
    --timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
```

For example:

```
$ csadmin config list --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765
        autostart             : disabled
        log                   : disabled

$ csadmin config list --json --uuid fedcba09-8765-4321-1234-567890abcdef
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "autostart": "disabled",
            "log": "disabled"
        }
    }
}
```

11.3.7.2. Get machine configuration

`csadmin config get` queries the current value of a given machine's specific configuration key:

```
$ csadmin config get --help
usage: csadmin config get [-h] -u UUID [--esn ESN] [--verbose] [--json] [--timeout TIMEOUT] key

positional arguments:
  key                   Configuration key

options:
  -h, --help            show this help message and exit
  -u UUID, --uuid UUID  The UUID of the machine.
  --esn ESN             ESN of the HSM hosting the SEE Machine.
  --verbose             Print verbose logs
  --json                Show response in json format
  --timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
```

For example:

```
$ csadmin config get log --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765         disabled

$ csadmin config get log --json --uuid fedcba09-8765-4321-1234-567890abcdef
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "log": "disabled"
        }
    }
}
```

11.3.7.3. Set machine configuration

`csadmin config set` sets a given machine's specified configuration key to the specified value:

```
$ csadmin config set --help
usage: csadmin config set [-h] -u UUID [--esn ESN] [--verbose] [--json] [--timeout TIMEOUT] key value

positional arguments:
  key                    Configuration key
  value                  Configuration value

options:
  -h, --help             show this help message and exit
  -u UUID, --uuid UUID   The UUID of the machine.
  --esn ESN              ESN of the HSM hosting the SEE Machine.
  --verbose              Print verbose logs
  --json                 Show response in json format
  --timeout TIMEOUT      Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
```

11.3.7.3.1. Use set to enable logging

```
$ csadmin config set log enabled --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765
        state                  : enabled

$ csadmin config set log enabled --json --uuid fedcba09-8765-4321-1234-567890abcdef
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "state": "enabled"
        }
    }
}
```

11.3.7.3.2. Use set to enable autostart

If Autostart is enabled the SEE machine will automatically start after a reset of the HSM.

Before starting, the system will check that a valid certificate matching the application has been loaded. If the certificate has expired the SEE machine will not start.

For this reason it is important that you check the expiry dates on your certificates and request new certificates from the CodeSafe developer for any certificates that are near their expiry date. You can check the expiry date on certificates by using `csadmin ids list` and looking for the `notAfter` field.

```
$ csadmin config set autostart enabled --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765
        state                  : enabled

$ csadmin config set autostart enabled --json --uuid fedcba09-8765-4321-1234-567890abcdef
```

```
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "state": "enabled"
        }
    }
}
```

11.3.8. Manage CodeSafe 5 logging

When logging is enabled (see Use set to enable logging), the standard output and standard error of all programs running on the container is saved into a log file that can be pulled and cleared by `csadmin log`:

```
$ csadmin log --help
usage: csadmin log [-h] {get,clear} ...

positional arguments:
  {get,clear}  manage a container's SEE logging operations (get/clear)
    get        Get a container's SEE log
    clear      Clear SEE logs on modules

options:
  -h, --help   show this help message and exit
```

You must use a positional argument (`get`, `clear`) with the `csadmin log` command.

11.3.8.1. Get CodeSafe 5 logs

To obtain CodeSafe 5 logs, use `csadmin log get`:

```
$ csadmin log get --help
usage: csadmin log get [-h] --uuid UUID [--json] [--timeout TIMEOUT] [--esn ESN] [--verbose]

options:
  -h, --help            show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to get the log from
  --json                Show response in json format
  --timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
  --esn ESN             ESN of the HSM to manage SEE logging for.
  --verbose             Print verbose logs
```

For example:

```
$ csadmin log get --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765          SUCCESS
Success: Started ipcdaemon

$ csadmin log get --json --uuid fedcba09-8765-4321-1234-567890abcdef
{
    "FEDC-BA09-8765": {
        "succeeded": true,
```

```
        "data": {
            "contents": "Success: Started ipcdaemon\n"
        }
    }
}
```

The internal log file has a maximum size of 128KB. When this size is reached, the log is cleared and a new log is started. This deletes previous data from the log. There is no rotation of the log.

11.3.8.2. Clear CodeSafe 5 logs

To clear CodeSafe 5 logs, use `csadmin log clear`:

```
$ csadmin log clear --help
usage: csadmin log clear [-h] --uuid UUID [--json] [--timeout TIMEOUT] [--esn ESN] [--verbose]

options:
  -h, --help           show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to clear/delete the log from
  --json               Show response in json format
  --timeout TIMEOUT    Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
  --esn ESN            ESN of the HSM to manage SEE logging for.
  --verbose            Print verbose logs
```

For example:

```
$ csadmin log clear --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765        SUCCESS
log: log cleared

$ csadmin log clear --json --uuid fedcba09-8765-4321-1234-567890abcdef
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "state": "log cleared"
        }
    }
}
[source]
```

11.3.9. Obtain machine-specific statistics

`csadmin stat` provides a mechanism for obtaining runtime data for CodeSafe machines on the system. It can only obtain data from machines in the `RUNNING` state:

```
$ csadmin stats --help
usage: csadmin stats [-h] [--timeout TIMEOUT] [-u UUID] [--esn ESN] [--verbose] [--json]

options:
  -h, --help           show this help message and exit
```

```
--timeout TIMEOUT     Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
-u UUID, --uuid UUID  The UUID of the machine to query stats from.
--esn ESN             ESN of the HSM hosting the SEE machine to query stats from.
--verbose             Print verbose logs
--json                Show response in json format
```

When executed without arguments, the stats command displays the runtime data of every machine currently running on the system:

```
$ csadmin stats
FEDC-BA09-8765
UUID                                     State    Name                         IP Address                 PID
CPU          Memory        KMem          Link        Tx bytes    Rx bytes    Total bytes
---------------------------------------------------------------------------------------------------------------
-----------------------------------------------------------------------------------------------
12345678-90ab-cdef-fedc-ba0987654321    RUNNING   hello            ffff::fff:ffff:ffff:eeee       1393
0.49 seconds   1.20 MiB       848.00 KiB    vethM19WNi    1.12 KiB     8.04 KiB      9.15 KiB
fedcba09-8765-4321-1234-567890abcdef    RUNNING   hello            ffff::fff:ffff:ffff:ffff       1610
0.42 seconds   1.26 MiB       836.00 KiB    vethsE1Pv4    1.12 KiB     4.98 KiB      6.10 KiB

$ csadmin stats --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "cs_stats": {
                "12345678-90ab-cdef-fedc-ba0987654321": {
                    "Name": "12345678-90ab-cdef-fedc-ba0987654321",
                    "State": "RUNNING",
                    "PID": "1393",
                    "CPU use": "0.51 seconds",
                    "Memory use": "1.24 MiB",
                    "KMem use": "888.00 KiB",
                    "Friendly Name": "hello",
                    "IP": "ffff::fff:ffff:ffff:eeee",
                    "Logging enabled": "False",
                    "Link": {
                        "name": "vethM19WNi",
                        "TX bytes": "1.12 KiB",
                        "RX bytes": "8.30 KiB",
                        "Total bytes": "9.42 KiB"
                    }
                },
                "fedcba09-8765-4321-1234-567890abcdef": {
                    "Name": "fedcba09-8765-4321-1234-567890abcdef",
                    "State": "RUNNING",
                    "PID": "1610",
                    "CPU use": "0.44 seconds",
                    "Memory use": "1.24 MiB",
                    "KMem use": "864.00 KiB",
                    "Friendly Name": "hello",
                    "IP": "ffff::fff:ffff:ffff:ffff",
                    "Logging enabled": "False",
                    "Link": {
                        "name": "vethsE1Pv4",
                        "TX bytes": "1.12 KiB",
                        "RX bytes": "5.24 KiB",
                        "Total bytes": "6.36 KiB"
                    }
                }
            }
        }
    }
}
```

When executed with the `--uuid|-u` argument, the command only displays the data of the specified machine:

```
$ csadmin stats --uuid fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765
UUID                                 State     Name                          IP Address              PID
CPU          Memory         KMem             Link        Tx bytes    Rx bytes    Total bytes
----------------------------------------------------------------------------------------------------------
------------------------------------------------------------------------------------------------
fedcba09-8765-4321-1234-567890abcdef    RUNNING    hello              ffff::fff:ffff:ffff:ffff          1610
0.48 seconds    1.21 MiB      716.00 KiB      vethsE1Pv4    1.12 KiB     5.76 KiB     6.88 KiB

$ csadmin stats --uuid fedcba09-8765-4321-1234-567890abcdef --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "cs_stats": {
                "fedcba09-8765-4321-1234-567890abcdef": {
                    "Name": "fedcba09-8765-4321-1234-567890abcdef",
                    "State": "RUNNING",
                    "PID": "1610",
                    "CPU use": "0.49 seconds",
                    "Memory use": "1.39 MiB",
                    "KMem use": "840.00 KiB",
                    "Friendly Name": "hello",
                    "IP": "ffff::fff:ffff:ffff:ffff",
                    "Logging enabled": "False",
                    "Link": {
                        "name": "vethsE1Pv4",
                        "TX bytes": "1.12 KiB",
                        "RX bytes": "6.03 KiB",
                        "Total bytes": "7.15 KiB"
                    }
                }
            }
        }
    }
}
```

11.3.10. Manage container SSHD

The container SSHD allows connections to be forwarded from the client machine to the container. Its keys and state can be queried and managed with `csadmin sshd`:

```
$ csadmin sshd --help
usage: csadmin sshd [-h] [--timeout TIMEOUT] [--esn ESN] {keys,state} ...

positional arguments:
  {keys,state}
    keys            manage a container's SSH daemon keys
    state           manage a container's SSH daemon state

optional arguments:
  -h, --help        show this help message and exit
  --timeout TIMEOUT  Time to wait for service response, in seconds. Default: 30s, must
                     be between 3s and 300s
  --esn ESN         ESN of the HSM to manage SSHD daemon for.
```

11.3.10.1. Manage SSHD keys

The public keys used to connect to the container SSHD can be managed with `csadmin sshd keys`:

```
$ csadmin sshd keys --help
usage: csadmin sshd keys [-h] {getserver,getclient,setclient} ...

positional arguments:
  {getserver,getclient,setclient}
    getserver           Obtain a container's SSH daemon keys
    getclient           Obtain a container's SSH daemon client keys
    setclient           Set a container's SSH daemon client keys

optional arguments:
  -h, --help            show this help message and exit
```

11.3.10.1.1. Retrieve keys

The SSHD's public key and the client keys set by the user can be retrieved by the `getserver` and `getclient` sub-commands respectively:

```
$ csadmin sshd keys getserver -h
usage: csadmin sshd keys getserver [-h] --uuid UUID [--json]

optional arguments:
  -h, --help            show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to retrieve SSH daemon keys for
  --json                Show response in json format
```

`getserver` returns the server public keys for the container specified by the given UUID. `getclient` returns the client public key set by the user.

For example, `getserver` produces the following output, shown in plain and json format:

```
$ csadmin sshd keys getserver -u fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765          SUCCESS
KEY_TYPE: ecdsa-sha2-nistp256
KEY_BODY: AAAA.....................w4wtS4=

$ csadmin sshd keys getserver -u fedcba09-8765-4321-1234-567890abcdef --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "server_key": {
                "key_type": "ecdsa-sha2-nistp256",
                "key_body": "AAAA.....................w4wtS4="
            }
        }
    }
}
```

11.3.10.1.2. Set client keys

The client public key that is used to connect to the SSHD server on the HSM is set by the user with `csadmin sshd keys setclient`:

```
$ csadmin sshd keys setclient --help
usage: csadmin sshd keys setclient [-h] --uuid UUID
                                   (--keyfile KEYFILE | --keytype KEYTYPE KEYBODY)

optional arguments:
  -h, --help            show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to set the SSH daemon keys for
  --keyfile KEYFILE, -f KEYFILE
                        File containing the public key body and type (a .pub file in
                        PEM format)
  --keytype KEYTYPE KEYBODY, -t KEYTYPE KEYBODY
                        Keytype of the public key (ex. ecdsa-sha2-nistp256) followed by
                        keybody
```

The key can be set using a keyfile (`--keyfile`) or by manual text entry (`--keytype`).

- Keyfile:

  When using the `--keyfile` option, the user passes in the path to the keyfile to be used. This is the public (`.pub`) key.

  Accepted public key types are `ecdsa-sha2-nistp521` and `ecdsa-sha2-nistp256`.

  For example:

  ```
  $ csadmin sshd keys setclient -u fedcba09-8765-4321-1234-567890abcdef --keyfile ~/test_key.pub
  ```

  Where `test_key.pub` is a key in OpenSSH format containing the key type and key body, for example:

  ```
  ecdsa-sha2-nistp256 AAAAE.....................5s8=
  ```

- Manual text entry:

  When using the `--keytype option`, the user manually enters the key type followed by the key body.

  For example:

  ```
  $ csadmin sshd keys setclient -u fedcba09-8765-4321-1234-567890abcdef --keytype ecdsa-sha2-nistp256
  TestKeyBody
  ```

11.3.10.2. Manage the SSHD state

The user can query and manage the container SSHD state using `csadmin sshd state`:

```
$ csadmin sshd state --help
usage: csadmin sshd state [-h] {enable,disable,get} ...

positional arguments:
  {enable,disable,get}
    enable              Enable a container's SSH daemon
    disable             Disable a container's SSH daemon
    get                 Get a container's SSH daemon state

optional arguments:
  -h, --help            show this help message and exit
```

11.3.10.2.1. Enable the container SSHD

The container SSHD will only start when the container has been enabled. By default, it is disabled. It can be enabled with `csadmin sshd state enable`:

```
$ csadmin sshd state enable --help
usage: csadmin sshd state enable [-h] --uuid UUID [--json]

optional arguments:
  -h, --help            show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to enable SSH daemon for
  --json                Show response in json format
```

If enabled, the SSHD will start when the container is started. If the container is already running when the SSHD is enabled, the SSHD will start if it was not already running.

The `enable` subcommand also returns the port and address on which the SSHD is listening, for example (output in plain text and in JSON format):

```
$ csadmin sshd state enable -u fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765          SUCCESS
SSHD PORT: 3006
LISTENING ADDRESS: aaaa::aa:aaaa:aaaa:aaaa

$ csadmin sshd state enable -u fedcba09-8765-4321-1234-567890abcdef --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "sshd_port": "3006",
            "listening_address": "aaaa::aa:aaaa:aaaa:aaaa"
        }
    }
}
```

11.3.10.2.2. Disable the container SSHD

When the container SSHD state is disabled, the container SSHD will not start. This is the default state. It can also be disabled with `csadmin sshd state disable`:

```
$ csadmin sshd state disable --help
usage: csadmin sshd state disable [-h] --uuid UUID

optional arguments:
  -h, --help           show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to disable SSH daemon for
If the container sshd is running, the disable command will also stop it.
It will not start again until is has been enabled.
```

11.3.10.2.3. Retrieve the container SSHD state

The user can also query the SSHD state with `csadmin sshd state get`. This displays the current SSHD state (enabled or disabled):

```
$ csadmin sshd state get --help
usage: csadmin sshd state get [-h] --uuid UUID [--json]

optional arguments:
  -h, --help           show this help message and exit
  --uuid UUID, -u UUID  The UUID of the machine to get the SSH daemon state for
  --json               Show response in json format
```

If the SSHD is disabled, the output only returns the state:

```
$ csadmin sshd state get -u fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765         SUCCESS
STATE: disabled

$ csadmin sshd state get -u fedcba09-8765-4321-1234-567890abcdef --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "state": "disabled"
        }
    }
}
```

If the SSHD is enabled, the output returns the port and address that the SSH server on the HSM is listening on, as well as the port on the container where connections should be forwarded to:

```
$ csadmin sshd state get -u fedcba09-8765-4321-1234-567890abcdef
FEDC-BA09-8765         SUCCESS
STATE: enabled
SSHD_LISTEN: 3001
CONTAINER_FORWARD: 9000
LISTENING ADDRESS: aaaa::aa:aaaa:aaaa:aaaa

$ csadmin sshd state get -u fedcba09-8765-4321-1234-567890abcdef --json
{
    "FEDC-BA09-8765": {
        "succeeded": true,
        "data": {
            "state": "enabled",
```

```
        "ports": {
            "sshd": "3001",
            "container": "9000"
        },
        "listening_address": "aaaa::aa:aaaa:aaaa:aaaa"
    }
  }
}
```

The container port is supplied by the user in `network-conf.json`.

11.3.10.3. Forward to the container SSHD

To establish port forwarding on the host machine, use the SSH local port forward command. Do this after the container SSHD is enabled, started, and its keys are set. Use of a secure tunnel over an open TCP connection to communicate with the container is recommended, so use local port forwarding instead dynamic port forwarding to enable SSH tunneling:

```
ssh -L local_port:destination_server_ip:remote_port ssh_server_hostname
```

The full details of the command:

```
ssh -i [PATH_TO_PRIV_KEY] -L [::1]:[LOCAL_PORT]:[CONTAINER_IP_ADDR%lxcbr0]:[CONTAINER_PORT] -f -N -p [SSHD_PORT]
launcher@[LISTENING_ADDRESS]%nshield0
```

Where:

- `PRIV_KEY` is the private key to the public client key that was set with `setclient`.
- `LOCAL_PORT` is the port on the local client where traffic to be forwarded will be sent.
- `CONTAINER_IP_ADDR` is the IP address of the container. This is the address returned when container is started.
- `CONTAINER_PORT` is the port on the container that is listening for forwarded traffic. This is set in `network-conf.json` as "ssh_tunnel" (valid range: 1024-65535).
- `SSHD_PORT` is the port the SSH daemon is listening on. This port is returned when SSHD is enabled and can also be found with `csadmin sshd state get`.
- `LISTENING_ADDRESS` is the address the host clients use to communicate with the HSM (on the HSM side). This address is returned when SSHD is enabled and can also be found with `csadmin sshd state get`.

For example:

```
ssh -i ~/test_ssh_keys/test_key -L [::1]:8888:[ffff::fff:ffff:ffff:ffff%lxcbr0]:8888 -f -C -N -p 6789
launcher@aaaa::aa:aaaa:aaaa:aaaa%nshield0
```

In preparation, there must also be some endpoint in the container that listens on the specified container port, for example, a socket. An example simple python socket script is shown below:

```python
import socket
import sys

HOST = ''
PORT = 8000 #the container port set in the network conf

soc = socket.socket(socket.AF_INET6, socket.SOCK_STREAM)

try:
    soc.bind((HOST, PORT))
except socket.error as message:
    print('Bind failed. Error code: '
        + str(message[0]) + ' Message '
        + message[1])
    sys.exit()

print('Socket binding complete')

soc.listen(9)

conn, address = soc.accept()

with conn:
    print('Connected by', address)
    while True:
        data = conn.recv(1024)
        if not data: break
        conn.send(data)

soc.close()
```

After the SSH instance is established, any traffic sent to the local port on the local client should be forwarded to the container on the container port. Because of the port forwarding restrictions set for the container SSHD in the `authorized_keys` file, attempts to forward to any port on the container other than the one specified in `network_conf.json` will be rejected.

11.3.11. Manage identities

Identities are needed in order to load and run CodeSafe applications. Developer identities (certificates) are managed with `csadmin ids`:

```
$ csadmin ids --help
usage: csadmin ids [-h] [--timeout TIMEOUT] [--esn ESN] {add,remove,get,validate,list,create} ...

positional arguments:
  {add,remove,get,validate,list,create}
    add                 Add a CodeSafe developer identity certificate to the nShield5s HSM
    remove              Remove a CodeSafe developer identity certificate from the nShield5s HSM
    get                 Get a CodeSafe developer identity certificate from the nShield5s HSM
    validate            Validate a CodeSafe developer identity certificate on the nShield5s HSM
```

```
    list            List all of the CodeSafe developer identity certificates on the nShield5s HSM
    create          Create a CSR for a named developer id key. Also create the developer id key if it does
not already exist

options:
  -h, --help        show this help message and exit
  --timeout TIMEOUT Time to wait for service response, in seconds. Default: 30s, must be between 3s and 300s
  --esn ESN         ESN of the HSM to manage developer identities certificates for.
```

The command allows end users to add new identities onto the HSM, remove existing identities from the HSM, retrieve (`get`) identities from the HSM, list identities presently installed on the HSM, and validate a specific identity present on the HSM.

11.3.11.1. Create a developer ID certificate

Only applicable to CodeSafe developers. See the *CodeSafe 5 Developer Guide*

11.3.11.2. List identity certificates

List identity certificates using `csadmin ids list`:

> 🛈 Listing certificates allows you to check their expiry dates by looking for the `notAfter` field for each certificate. You should request new certificates from the CodeSafe developer for any certificates that have expired or are near their expiry date. Entrust recommend that you remove any expired certificates using `csadmin ids remove`.

```
$ csadmin ids list --help
usage: csadmin ids list [-h] [--json] [--verbose]

options:
  -h, --help  show this help message and exit
  --json      Show response in json format
  --verbose   Print verbose logs
```

11.3.11.3. Add identities

Add identity certificates using `csadmin ids add`:

```
$ csadmin ids add --help
usage: csadmin ids add [-h] [--verbose] CERTFILE

positional arguments:
  CERTFILE    The developer identity certificate file to be added into the HSM(s).

options:
  -h, --help  show this help message and exit
  --verbose   Print verbose logs
```

11.3.11.4. Remove identities

Remove identity certificates using `csadmin ids remove`:

> **i** Entrust recommend that you remove any expired certificates so that only valid certificates remain in the system.

```
$ csadmin ids remove --help
usage: csadmin ids remove [-h] -s SERIALNO [--verbose]

options:
  -h, --help            show this help message and exit
  -s SERIALNO, --serialno SERIALNO
                        Serial number of the certificate in base ten
  --verbose             Print verbose logs
```

11.3.11.5. Retrieve identities

Retrieve identity certificates using `csadmin ids get`:

```
$ csadmin ids get --help
usage: csadmin ids get [-h] -s SERIALNO [--json] [--verbose]

options:
  -h, --help            show this help message and exit
  -s SERIALNO, --serialno SERIALNO
                        Serial number of the certificate in base ten
  --json                Show response in json format
  --verbose             Print verbose logs
```

11.3.11.6. Validate identities

Validate identity certificates using `csadmin ids validate`:

```
$ csadmin ids validate --help
usage: csadmin ids validate [-h] -s SERIALNO [--verbose]

options:
  -h, --help            show this help message and exit
  -s SERIALNO, --serialno SERIALNO
                        Serial number of the certificate in base ten
  --verbose             Print verbose logs
```

12. Remote Operator

This chapter explains:

- The concept of Remote Operator
- How to configure Remote Operator.

> **ℹ** If you wish to use the Remote Operator feature, you must have enabled it as described in Enabling optional features. The Remote Operator feature must have been ordered for, and enabled on, the nShield module that you intend to use as the remote, unattended module.

12.1. About Remote Operator

The Remote Operator feature enables the contents of a smart card inserted into the slot of one module (the *attended module*) to be securely transmitted and loaded onto another module (an *unattended module*). This is useful when you need to load an OCS-protected key onto a machine to which you do not have physical access (because, for example, it is in a secure area).

For Remote Operator to work, the modules must be in the same Security World. You insert the required cards from the OCS into a slot in the attended module. From this module, the contents of the OCS are transmitted over secure channels to the unattended module, which then loads them. You do not need physical access to the unattended module in order to load the OCS onto it.

The following limitations apply to Remote Operator:

- You cannot access non-persistent card sets remotely
- You cannot use the `createocs` command-line utility to write new cards or card sets remotely.

You can export a slot from an attended module and import a slot to any (unattended) module in the Security World. Before you can import a slot to one module, you must first export it from another module.

12.2. Configuring Remote Operator

This section explains how to configure Remote Operator.

12.2.1. Overview of configuring Remote Operator

Before you can use Remote Operator, you must perform the following initial configuration tasks:

1. Configure the HSMs for Remote Operator.

   The HSMs must be in the same Security World, and must have been initialized with remote card set reading enabled.

   Both the attended and the unattended HSM must be in operational mode before they can import or export slots. See Checking and changing the mode on the HSM for more about changing the mode.

2. Configure the HSMs for slot import and export, as appropriate.

Starting from 12.81, you can export and import dynamic slots as Remote Operator slots.

After the initial configuration is complete, to use Remote Operator you must:

1. Create a Remote OCS (that is, an OCS with the correct permissions for Remote Operator).

2. Generate keys that are protected by the Remote OCS.

3. Ensure your application is configured to use keys protected by the Remote OCS.

12.2.2. Configuring HSMs for Remote Operator

1. Ensure both HSMs are initialized into the same Security World; see Adding or restoring an HSM to the Security World.

   > By default, HSMs are initialized with remote card-set reading enabled. If you do not want an HSM to be able to read remote card sets, you can initialize it by running the `new-world` with the `-S` *MODULE* (where *MODULE* is the HSM's ID number).

2. For the unattended HSM:

   a. Check whether the Remote Operator feature is enabled by running the `enquiry` command-line utility. The output for the HSM must include `Remote Share` in its list of Features.

   b. Check whether the HSM has permission to allow loading of Remote OCSs by selecting **Security World mgmt** > **Display World info**.

   c. Check whether the correct software, with permission to receive remote shares, is present by running the `nfkminfo` command-line utility.

The output from this selection must show that `flags` are set to include **ShareTarget**, as in the following example:

```
Module #1
generation 2
state 0x2 Usable
flags 0x10000 ShareTarget
n_slots 3
esn 8851-43DF-3795
hkml 391eb12cf98c112094c1d3ca06c54bfe3c07a103
```

12.2.3. Configuring slot import and export

For information about the parameters controlled by the hardserver configuration file, see:

- slot_exports
- slot_imports
- slot_mapping

Before you can configure hardservers for Remote Operator, ensure that:

- You have configured the attended and unattended HSMs for Remote Operator as described in Configuring HSMs for Remote Operator.
- Your network firewall settings are correct. See the *Installation Guide* for more information about firewall settings.

When the HSMs have been configured, use one of the following methods to configure slot import and export:

- Use the nShield HSM front panel, see Configuring slot import and export using the nShield HSM front panel.
- Update the HSM configuration file, see Configuring hardservers for Remote Operator using the HSM configuration file.

12.2.3.1. Configuring slot import and export using the nShield HSM front panel

1. Configure the attended HSM to export a slot by following these steps:

   a. From the main menu, select **Security World mgmt** > **Set up remote slots** > **Export slot**.

      Use this option for exporting slot #0 only.

      If you need to configure the export of slots other than 0, see Configuring hardservers for Remote Operator using the HSM configuration file.

   b. Specify the HSM to which the slot is being export by supplying values for:

      ▪ The IP address of the unattended HSM

      ▪ The ESN of the unattended HSM.

2. Configure the unattended HSM to import the slot that you are exporting from the attended HSM by following these steps:

   a. From the main menu, select **Security World mgmt** > **Set up remote slots** > **Import slot**.

   b. Specify the details of the Remote Operator slot by supplying values for:

      ▪ The IP address of the HSM from which the slot is being exported

      ▪ The ESN of the HSM from which the slot is being exported

      ▪ The ID of the slot on the importing HSM

      ▪ The port to use to connect to the hardserver hosting the attended HSM.

You can check that the slot was imported successfully by, on the unattended machine, running the command:

```
slotinfo -m 1
```

If slot importation was successful, the output from this command includes the line:

```
Slot Type           Token   IC   Flags   Details
#0   Smartcard      present 3    A
#1   Software Tkn   -       0
#2   Smartcard      -       0    AR
```

The `R` in the `Flags` column indicates that slot `2` is a Remote Operator slot.

Applications running on the unattended machine can now use slot `2` to load OCSs that are presented to slot `0` on the attended machine. If any of the cards require a passphrase, the application must pass this to the unattended HSM in the usual way.

For the application to be able to load the OCS onto the unattended HSM, it must be able to read the card set files associated with the OCS from the local Key Management Data directory. If the OCS was created on a different machine, you must copy the card set files in the Key Management Data directory onto the unattended machine (either manually or by using client cooperation; for more information, see Setting up client cooperation).

The same applies for any keys that an application on an unattended HSM needs to load but that were not generated on that machine.

12.2.3.2. Configuring hardservers for Remote Operator using the HSM configuration file

1. On the attended HSM's host machine, configure the hardserver to allow slot 0 of the local HSM (with ESN AAAA-AAAA-AAAA) to be exported to a remote HSM (with ESN BBBB-BBBB-BBBB, hosted by the machine with the IP address 222.222.222.222):

   a. Create a copy of the configuration file as `config.new` in the `/opt/nfast/kmdata/hsm-ESN/config` (**Linux**) or `C:\ProgramData\nCipher\nfast\kmdata\hsm-ESN\config` (**Windows**) directory.

   b. Edit the sections related to slot export in `config.new`:

   ```
   [slot_exports]
   # Start of the slot_exports section
   # Local slots that the hardserver should allow remote modules to import. Note
   # that if a slot which has been remapped in the slot_mapping section is to be
   # exported, it must be referred to in this section by its original
   # (pre-mapping) local_slotid.
   # Each entry has the following fields:
   #
   # ESN of the local module whose slot is allowed to be exported.
   #  local_esn=ESN
   #
   # SlotID of the slot which is allowed to be exported. (default=0)
   #  local_slotid=INT
   #
   # IP address of the machine allowed to import the slot or empty to allow all
   # machines. (which is the default)
   #  remote_ip=ADDR
   #
   # ESN of the module allowed to import the slot or "" to allow all modules
   # which are permitted in the security world. (default ="")
   #  remote_esn=ESN
   ```

   ```
   [slot_mapping]
   # Start of the slot_mapping section
   # Slot remapping configuration. Notes for Remote Operator users: If a slot
   # which is remapped in this section is also exported in the slot_exports
   # section, the local_slotid field in the slot_exports section must be set to
   # the original (pre-mapping) local_slotid. When importing that slot in another
   # module, the slot_imports section must refer instead to the new
   # (post-mapping) remote_slotid.
   # Each entry has the following fields:
   #
   # ESN of the module on which slot 0 will be remapped with another.
   #  esn=ESN
   #
   # Slot to exchange with slot 0. Setting this value to 0 means do
   # nothing.(default=0)
   #  slot=INT
   ```

   c. Run the `cfg-pushnethsm` utility on the updated configuration file, specifying the updated file and the network address of the nShield HSM to load the new configuration.

   ```
   cfg-pushnethsm --address=<module_address> <path_to_config_file>
   ```

d. Check that the configuration file has been updated. This can be confirmed using the timestamp on the updated config file.

e. Clear the HSM for the changes to take effect, run the `nopclearfail` command:

2. On the unattended module's host machine, configure the hardserver to import slot 0 from the remote attended module (with ESN *AAAA-AAAA-AAAA*, hosted by the machine with the IP address *111.111.111.111*) to the local module (with ESN *BBBB-BBB-BBBB*).

a. Edit the sections related to slot import in `config.new`:

```
[slot_imports]
# Start of the slot_imports section
# Remote slots that the hardserver should import to modules on this machine.
# Note that if a remote slot which has been remapped in the slot_mapping
# section on the remote system is to be imported, it must be referred to in
# this section by its new (post-mapping) remote_slotid.
# Each entry has the following fields:
#
# ESN of the local module to import the slot to
#  local_esn=ESN
#
# SlotID to use to refer to the slot when it is imported on the local module.
# Setting this value to 0 means it will be automatically assigned to the
# lowest available value. (default=0)
#  local_slotid=INT
#
# IP address of the machine hosting the slot to import
#  remote_ip=ADDR
#
# Port to connect to on the remote machine
#  remote_port=PORT
#
# ESN of the remote module to import the slot from
#  remote_esn=ESN
#
# SlotID of the slot to import on the remote module (default=0)
#  remote_slotid=INT
```

b. Run the `cfg-pushnethsm` utility on the updated configuration file, specifying the updated file and the network address of the nShield HSM to load the new configuration.

```
cfg-pushnethsm --address=<module_address> <path_to_config_file>
```

c. Check that the configuration file has been updated. This can be confirmed using the timestamp on the updated config file.

d. Clear the HSM for the changes to take effect, run the `nopclearfail` command:

3. Check the Remote Operator slot configuration:

```
slotinfo -m 1
```

If slot import was successful, the output from this command includes the line:

```
Slot Type           Token  IC   Flags   Details
#0   Smartcard      present 3    A
#1   Software Tkn   -      0
#2   Smartcard      -      0     AR
```

The `R` in the `Flags` column indicates that slot `2` is a Remote Operator slot.

Applications running on the unattended machine can now use slot `2` to load OCSs that are presented to slot `0` on the attended machine. If any of the cards require a passphrase, the application must pass this to the unattended HSM in the usual way.

For the application to be able to load the OCS onto the unattended HSM, it must be able to read the card set files associated with the OCS from the local Key Management Data directory. If the OCS was created on a different machine, you must copy the card set files in the Key Management Data directory onto the unattended machine (either manually or by using client cooperation; for more information, see Setting up client cooperation).

The same applies for any keys that an application on an unattended HSM needs to load but that were not generated on that machine.

12.2.4. Using Remote Operator with applications requiring cards in slot 0

If you want to use Remote Operator, but have an application that expects cards to be presented in slot 0, you must configure a slot mapping for each affected HSM.

1. Do one of the following:

   a. Use the `slot_mapping` section in the module configuration file to define a Dynamic Slot to exchange with slot 0 for an HSM and push the updated configuration file to the nShield HSM.
   See HSM and client configuration files for more about module configuration file, slot_mapping for more about the `slot_mapping` section and xref:configure.adoc#AboutUserPrivileges,About user privileges>> for more about editing the module configuration file.

   Or:

   b. Use the front panel controls to navigate to **Security World mgmt > Set up dynamic slots > Slot mapping** and follow the instructions on the screen.

You can check the mapping by:

- Running the command:

```
slotinfo -m 1
```

For example, if remote slot #2 has been mapped to slot #0, the output from this command includes the lines:

```
Slot Type Token IC Flags Details
#0 Smartcard - 1 AR
#1 Software Tkn - 0
#2 Smartcard - 0 A
```

- The R in the Flags column indicates that slot #0 is now a Remote Slot

> 🛈 Slot mapping can also be configured for a dynamic remote slot, i.e. a dynamic slot in a different HSM which has been imported to the relevant HSM. The **Flags** column will contain the flags ARD.

or:

- Using the front panel controls to navigate to **Security World mgmt > Display World**.

When dynamic slots are added to an HSM after the initial configuration was done with only remote slots, the dynamic slots will take precedence over the remote slots. The slot numbers of the remote slots will therefore change. You will have to revise the slot mapping and specify the new slot number of the remote slot.

12.2.5. Using Remote Operator on Remapped Slots

If a slot has been mapped to slot #0 on the attended HSM, it is still possible to export the local slot to an unattended HSM. Further, if the mapped slot is a dynamic slot, it is possible to export it as well. To do this, do the following:

1. On the attended HSM's host machine, configure the hardserver to allow the export of the relevant slot by referring to it by its original slotID.
   a. To export the local slot, local_slotid=0.
   b. To export a dynamic slot, local_slotid=2 (or higher if the HSM is configured with multiple dynamic slots).
2. On the unattended HSM's host machine, configure the hardserver to import the relevant slot by referring to it by its new slotID.
   a. To import the exported local slot, remote_slotid=2 (or higher, same as the slotID specified in the mapping section of the attended HSM's configuration file).
   b. To import the exported dynamic slot, remote_slotid=0.

12.2.6. Configuration Example for Using Remote Administration and Remote Operator Concurrently

Below is an example of the relevant portions of a hardserver config file to achieve concurrent usage of Remote Administration and Remote Operator. It is broken up and explained per config file section.

The `dynamic_slots` section allocates exactly 1 dynamic slot to each of modules 1 and 2.

```
[dynamic_slots]
esn=BBBB-BBBB-BBBB
slotcount=1
-----
esn=AAAA-AAAA-AAAA
slotcount=1
```

The `slot_imports` section first imports module 1 slot #0 to module 2 slot #3 and then imports module 1 slot #2 to module 2 slot #4.

```
[slot_imports]
local_esn=AAAA-AAAA-AAAA
remote_ip=127.0.0.1
remote_port=9004
remote_esn=BBBB-BBBB-BBBB
remote_slotid=0
-----
local_esn=AAAA-AAAA-AAAA
remote_ip=127.0.0.1
remote_port=9004
remote_esn=BBBB-BBBB-BBBB
remote_slotid=2
```

The `slot_exports` section allows module 1 slot #0 and module 1 slot #2 to be exported by that module.

```
[slot_exports]
local_esn=BBBB-BBBB-BBBB
local_slotid=0
-----
local_esn=BBBB-BBBB-BBBB
local_slotid=2
```

The `slot_mapping` section swaps module 2 slot #0 and module 2 slot #2.

```
[slot_mapping]
esn=AAAA-AAAA-AAAA
slot=2
```

After making the changes above to the hardserver configuration file:

1. Push the hardserver configuration file to the nShield HSM by running `cfg-pushnethsm`.

2. Clear the modules by running `nopclearfail`.

This is the expected system configuration output for the relevant modules:

```
slotinfo -m1
Slot Type          Token    IC     Flags    Details
#0   Smartcard     -        0      A
#1   Software Tkn  -        0
#2   Smartcard     -        0      AD

slotinfo -m2
Slot Type          Token    IC     Flags    Details
#0   Smartcard     -        0      AD
#1   Software Tkn  -        0
#2   Smartcard     -        0      A
#3   Smartcard     -        0      AR
#4   Smartcard     -        0      ARD
```

12.2.7. Using Remote Operator with Remote Administration with Older Versions of the Software

Versions of Remote Operator older than 12.81 do not support its concurrent use with the Remote Administrator feature. In such a case, the following features are not supported:

- Exporting and importing dynamic slots

- Mapping remote slots to slot #0

- Automatic assignment of slotID when importing slots

It is possible to use some of the features when the attended HSM (exporting end) has the new version of the software (12.81+) and the unattended HSM (importing end) has an older version (pre-12.81).

A dynamic slot which has been exported by the attended HSM can be imported to the unattended HSM. Its local slotID will need to be manually specified if the unattended HSM has any dynamic slots configured. This is due to the default import slot (slot #2) being occupied by the dynamic slot. The unattended HSM can remap its dynamic slots to slot #0, but cannot remap any of its imported slots.

12.3. Creating OCSs and keys for Remote Operator

When you have configured the HSMs and slot import and export, you can create Remote OCSs and generate keys protected by them. These Remote OCSs and keys can be used by applications running on the unattended HSM.

For the most part, card sets and keys intended to be used with Remote Operator are similar

to their ordinary, non-Remote counterparts.

12.3.1. Creating OCSs for use with Remote Operator

You can generate Remote OCSs by using KeySafe or by running the `createocs` command-line utility with the `-q|--remotely_readable` option specified. The cards in a Remote OCS must be created as persistent; see Persistent Operator Card Sets.

To check whether the card in a slot is from a Remote OCS, select **Security World mgmt** > **Display World info** from the main menu or run the `nfkminfo` command-line utility. The output displays slot section information similar to the following:

```
Module #1 Slot #0 IC 1
generation        1
phystype          SmartCard
slotlistflags     0x2
state             0x5
Operator flags    0x20000 RemoteEnabled
shareno           1
shares            LTU(Remote)
error             OK
```

In this example output, the `RemoteEnabled` flag indicates the card in the slot is from a Remote OCS.

> ℹ️ If you create a Remote OCS on the attended machine, then you must copy the Key Management Data files on the attended machine to the unattended machine.

> ℹ️ Both the attended and unattended HSMs must be in the same Security World before you generate a Remote OCS. If you are not using client cooperation, the Key Management Data directories must be manually synchronized after you generate the Remote OCS.

> ℹ️ If you already have recoverable keys protected by a non-Remote OCS, you can transfer them to a new Remote OCS by using KeySafe or the `replaceocs` command-line utility.

12.3.2. Loading Remote Operator Card Sets

Once configured, the Remote Operator slots can be used by all the standard nShield libraries. A Remote Operator slot can be used to load any OCSs that have been created to allow remote loading. For more information about the applications to use with remote cards, see Application interfaces. For more information about Remote Operator slots, see

Remote Operator.

> **ℹ** After an OCS has been inserted into a Remote Operator slot, for each
> time a given card is inserted, the module only allows each share on that
> card to be read one time. If there is a second attempt to read shares
> from that card before the card is reinserted, the operation fails with a
> `UseLimitsUnavailable` error.

12.3.3. Generating keys for use with Remote Operator

After you have created a Remote OCS, to generate keys protected by it you can run
KeySafe or the `generatekey` and `preload` command-line utilities on the unattended module,
inserting cards to the slot attached to the attended module. For more information about
generating and working with keys, see Working with keys.

> **ℹ** If you generate keys protected by a Remote OCS on the attended mod-
> ule, then you must copy the files in the Key Management Data directory
> on the attended machine to the unattended module.

> **ℹ** KeySafe can list imported slots, but cannot use them.

If you already have an OCS-protected key that you want to use, but the protecting OCS is
not a Remote OCS, you can use KeySafe to protect the key under a new Remote OCS if the
key was originally generated with the key recovery option enabled.

However, if the key was not generated with key recovery enabled, you cannot protect it
under a different OCS. In such a case, you must generate a new key to be protected by a
Remote OCS.

12.3.4. Configuring the application

After you have configured the HSMs and slot import and export, created a Remote OCS,
and generated keys protected by the Remote OCS, configure the application with which
you want to use these keys as appropriate for the particular application.

After you have configured the application, start it remotely from the attended machine.
Insert cards from the OCS into the attended machine's exported slot as prompted.

12.3.5. Managing Remote Operator slots using the unit front panel

12.3.5.1. Editing Remote Operator slots

You can change the details of a Remote Operator slot. You must always update the details of both the exported slot on the local module and the imported slot on the remote module.

To update an exported a slot on the module:

1. From the main menu, select **Security World mgmt** > **Set up remote slots** > **Edit exported slot**.
2. Select the exported slot that you want to update. Slots are identified by the IP address of the remote module.
3. Update the details of the slot.

To update an imported slot on the unit:

1. From the main menu, select **Security World mgmt** > **Set up remote slots** > **Edit imported slot**.
2. Select the imported slot that you want to update. Slots are identified by the IP address of the remote module.
3. Update the details of the slot.

12.3.5.2. Deleting Remote Operator slots

You can delete Remote Operator slots.

To delete an exported slot, from the main menu, select **Security World mgmt** > **Set up remote slots** > **Delete exported slot** and select the slot you want to delete.

To delete an imported slot, from the main menu, select **Security World mgmt** > **Set up remote slots** > **Delete imported slot** and select the slot you want to delete.

13. Resource Watchdog

The Resource Watchdog monitors useful information such as CPU usage, ethernet interface states and addresses, etc. and reports it in the syslog of the Connect. The Watchdog is entirely configurable with its own bespoke configuration file.

It is possible to redirect the syslog of the Connect to either the RFS or a client, for more information, see syslog.

13.1. Enabling or disabling the Watchdog

The Resource Watchdog is disabled by default and starts automatically at Connect boot. To disable it, or to re-enable it, go in the FPUI to Menu **System** > **System Configuration** > **Watchdog Config** > **Enable watchdog** and select the appropriate setting (**ENABLE** or **DISABLE**).

> ℹ️ | Connect needs to reboot before Enable or Disable setting takes place.

You can also modify the `[nethsm_watchdog]` section in the Connect Configuration File by setting **enable_watchdog** to **yes** or **no**. Then, push it to the Connect as usual using `cfg-pushnethsm`.

> ℹ️ | See the *nShield Connect User Guide* section *Configuring the nShield Connect to use the client* for more information.

The watchdog automatically applies the new setting after a cycle (usually up to a minute).

13.2. Understanding the default settings

The default Watchdog Configuration File is as follows:

```
Network Devices:
- eth0
- eth1
- nshield0
Monitored Processes:
- hardserver
- netui
- cosmod
System Information Report:
  enabled: False
  interval: 1h
CPU Usage Monitoring:
  threshold: 180%
  interval: 60s
  frequency: 1h
```

13.2.1. Network Devices

The Watchdog monitors when the ethernet interfaces listed in its configuration file change state. By default, it monitors **eth0**, **eth1** and **nshield0** (only available on nShield 5c). You can add or remove ethernet interfaces from the configuration file.

A typical report looks like this:

```
Mar 13 10:47:50 nethsm nfwatchdog: Interface eth0 is now up.
Mar 13 10:47:50 nethsm nfwatchdog: Interface eth1 is now down.
```

The Watchdog also monitors address changes on these interfaces. A typical report looks like this:

```
Mar 13 10:47:50 nethsm nfwatchdog: The addresses of interface eth0 have been set to ["00:60:e0:87:a1:59",
"172.23.135.129"].
Mar 13 13:58:25 nethsm nfwatchdog: The addresses of interface eth0 have changed from ["00:60:e0:87:a1:59",
"172.23.135.8"] to ["00:60:e0:87:a1:59", "172.23.135.129"].
```

13.2.2. Monitored Processes

The Watchdog monitors when the processes listed in its configuration file start and stop. By default, it monitors **harderver**, **netui** and **cosmod**. You can add or remove processes from the configuration file. Please provide only the process name, do not include the path.

A typical report looks like this:

```
Mar 13 10:47:50 nethsm nfwatchdog: Process netui is not running.
Mar 13 10:47:50 nethsm nfwatchdog: Process hardserver is running: {"arguments": ["../sbin/hardserver", "-
Llogfile"], "children": [{"arguments": ["/opt/nfast/python3/bin/python3", "-uIBm", "nshield.entrypoint", "--",
"/opt/nfast/bin/hsc_servicehosts"], "children": [], "name": "python3", "pid": 607, "status": "running", "uptime":
"0:00:01"}], "name": "hardserver", "pid": 598, "status": "sleeping", "uptime": "0:00:01"}
```

If the process is running, the Watchdog reports, in JSON format, the process name, arguments, PID, uptime, status and children, if any.

13.2.3. System Information Report

The Watchdog can report System Information. By default, the report is disabled. To enable it, set **System Information Report** > **enabled** to **yes**.

By default, the System Information Report is printed every hour (if enabled). To change the interval, set **System Information Report** > **interval** to the desired interval. This field can be expressed in seconds, minutes, hours, or days. Examples: 8h, 2d8h5m20s, 2m4s. A unit is

required.

A typical report looks like this:

```
Mar 13 14:05:32 nethsm nfwatchdog: System Report for the Connect is {"system_uptime": "3:20:13.930539",
"virtual_memory": {"total": 8253513728, "free": 7962390528, "available": 7996977152, "used": 173985792},
"load_avg": [0.09, 0.04, 0.01]}
```

It contains the following information in JSON format: * system up time * virtual memory in bytes (total, free and available, as defined in psutil) * system load over the last 1, 5 and 15 minutes, as defined in psutil. * The "load" represents the processes which are in a runnable state, either using the CPU or waiting to use the CPU (for example, waiting for disk I/O).

Field	Value	Comment
enabled	True or False, default: False.	Whether the System Information Report is enabled.
interval	An interval in seconds, minutes, hours or days, default: 1h.	The interval at which the System Information Report is printed, if enabled. A unit (s, m, h or d) is required.

13.2.4. CPU Usage Monitoring

The Watchdog monitors both the global CPU usage on the Connect and the CPU usage per process. Should the total CPU usage or the CPU usage of any process (as computed with psutil's cpu_percent function) be above the threshold (default: 180%) for a sustained interval (default: 60s), then the Watchdog will report it. This check is performed every hour by default (**CPU Usage Monitoring** > **frequency**).

A typical report looks like this:

```
Mar 13 10:48:44 nethsm nfwatchdog: Global CPU usage for the Connect is 92.1%, higher than the threshold of 90%.
Mar 13 10:48:44 nethsm nfwatchdog: CPU usage for the the following processes is higher than the threshold of 90%.
Mar 13 10:48:44 nethsm nfwatchdog: Process hardserver, 90.5%: {"arguments": ["../sbin/hardserver", "-Llogfile"],
"children": [{"arguments": ["../sbin/hardserver", "--spawn-svc"], "children": [], "name": "hardserver", "pid":
620, "status": "sleeping", "uptime": "0:00:54"} {"arguments": ["/opt/nfast/bin/ncssh", "-id",
"/opt/nfast/services/client/ncoreapi/ssh/id_ecdsa", "-known-hosts", "/opt/nfast/services/module/5CA8-5A32-
80F2/ncoreapi/known_hosts", "-hosts", "/opt/nfast/services/module/hosts.txt", "-hostname", "nshield-5CA8-5A32-
80F2.local", "-port", "2203", "-user", "ncoreapi", "ncoreapi", "operational"], "children": [], "name": "ncssh",
"pid": 626, "status": "sleeping", "uptime": "0:00:54"}], "name": "hardserver", "pid": 598, "status": "sleeping",
"uptime": "0:00:55"}
```

For each process, the Watchdog reports the process name, the CPU usage as reported by psutil's cpu_percent, and in JSON format: the process name, arguments, PID, children, status and uptime.

The threshold must be a number between 0% and 100% x (number of CPUs on the Connect). The % sign is optional. The default is 180%.

The interval is the time during which the Watchdog samples the CPU usage. This interval is blocking, that is to say that the Watchdog is not able to perform any other monitoring while the sampling is happening. The default is 60s. This field can be expressed in seconds, minutes, hours, or days. Examples: 8h, 2d8h5m20s, 2m4s. A unit is required.

The frequency sets how often the Watchdog performs the sampling. By default, the Watch dog performs the CPU usage sampling for 60s every hour (frequency = 1h). This field can be expressed in seconds, minutes, hours, or days. Examples: 8h, 2d8h5m20s, 2m4s. A unit is required.

Field	Value	Comment
threshold	A percentage between 0% and 100% x number of CPU, default: 180%.	Threshold above which global CPU and process CPU usage is reported. The % sign is optional.
interval	An interval in seconds, minutes, hours or days, default: 60s.	The interval during which the Watchdog computed the CPU usage (both total and per process). A unit (s, m, h or d) is required.
frequency	A frequency in seconds, minutes, hours or days, default: 1h.	The frequency at which the Watchdog monitors the CPU usage. A unit (s, m, h or d) is required.

13.2.5. Other features

13.2.5.1. Zombie processes

On top of the aforementioned configurable features, the Watchdog also monitors zombie processes, or processes that are waiting for their parent.

A typical report looks like this:

```
Aug  7 03:47:46 nethsm nfwatchdog: Information: Process netui exited, awaiting parent.
Aug  7 03:47:51 nethsm nfwatchdog: Information: 3 processes (netui, cosmod, hardserver) exited, awaiting parent.
```

13.2.5.2. Report filtering

Should any Watchdog event occur several times in a row, the reports will be filtered out, so as not to flood the log. This will be reported by the Watchdog in the following way:

```
May  9 09:35:47 nethsm nfwatchdog: Global CPU usage for the Connect is 200.0%, higher than the threshold of 180.0%.
```

13.3. Reading or modifying the Watchdog Configuration File

1. Create a Watchdog Configuration file called `nfwatchdog.yml.new` on your client with your changes.

2. Prepare to push the new configuration file. Perform one of the following:

   ° On the Connect UI:

      i. On Menu **System** > **System Configuration** > **Watchdog Config** > **Config push mode**, set **push** to **on**.

      ii. On Menu **System** > **System Configuration** > **Watchdog Config** > **Client address**, set the address of your client.

   ° Update the Connect Configuration File directly:

      i. Open the file and locate the `[nethsm_watchdog]` section.

      ii. Set **push** to **on**.

      iii. Set **remote_ip** to the IP of your client.

      iv. (Optionally) Set **remote_keyhash** as well. For more information, see config_op.

3. Push the Connect Configuration File as usual, for example using `cfg-pushnethsm`.

4. Use the following command to send your new Watchdog Configuration File to the Connect:

```
sudo nfcp ./nfwatchdog.yml.new <connect_ip>:cfg-nfwatchdog
```

The Connect will pick up the new configuration and the Watchdog will apply it after a cycle (usually up to 1 minute). The syslog should show:

```
Feb 10 13:22:36 nethsm ../sbin/config-update: found new pushed watchdog config, attempting to install
Feb 10 13:22:36 nethsm ../sbin/config-update: successfully installed new watchdog config
```

The Connect will then push the current configuration back to `/opt/nfast/kmdata/hsm-<esn>/config` on your client, alongside the Connect Configuration File, if you have enabled `[config_op]` push.

13.3.1. Configuring via hardserver configuration file

The Watchdog can be configured using the hardserver configuration file, in the following steps.

1. Set the Watchdog entries in the hardserver configuration file.

```
[nethsm_watchdog]
# Start of the nethsm_watchdog section
# Connect Watchdog configuration. This section allows you to enable or disable
# the Connect Watchdog and set up config file push.
# Each entry has the following fields:
#
# Enable the Connect watchdog. (default=no)
#   enable_watchdog=ENUM
#
# Whether to allow a client to push new watchdog config files to the netHSM.
# If "on" then this effectively allows a client to remotely configure the
# nethsm_watchdog. (default=off)
#   push=ENUM
#
# The IP address of the client allowed to push watchdog config files. If not
# set, or set to 0.0.0.0 or ::, allows ANY IP address to push on a new config
# file.
#   remote_ip=ADDR
#
# The hash of the key that the authorised client should use to authenticate
# itself, or 40 zeros to indicate no key authentication required. (Default is
# 40 zeros).
#   remote_keyhash=KEYHASH
```

`enable_watchdog`	This section allows you to enable or disable the Connect Watchdog and set up config file push. The default is 'no'.
`push`	Whether to allow a client to push new Watchdog config files to the netHSM. If this is "on", it allows a client to remotely configure the nethsm_watchdog. The default is 'off'.
`remote_ip`	The IP address of the client allowed to push Watchdog config files. If this is not set, or set to 0.0.0.0 or ::, it allows any IP address to push a new config file.
`remote_keyhash`	The hash of the key that the authorised client should use to authenticate itself, or 40 zeros to indicate no key authentication required. The default is 40 zeros.

2. Restart the hardserver to load the updated configuration file.

3. Reboot the connect if enabling or disabling the Watchdog.

13.4. Troubleshooting

The Watchdog is designed to fall back to its default values in case the Watchdog Configuration File is missing or malformed. If the Watchdog does not behave as you intended, examine the syslog for clues.

> ℹ️ Factory stating the Connect makes the Watchdog Configuration File revert to its defaults.

Error message	Likely cause	Solution
Could not open Watchdog Configuration File. Watchdog Configuration File has been restored to its default state.	Watchdog Configuration File is missing or contains a YAML syntax error.	Push a new Watchdog Configuration File.
Watchdog Configuration File was unexpectedly empty. Watchdog Configuration File has been restored to its default state.	Watchdog Configuration File is missing or empty.	Push a new Watchdog Configuration File.
Could not read missing ... field. Using default ... instead.	The field is missing or there is a typo.	Update the Watchdog Configuration File and push it again.
Ignoring unrecognized category ... from Watchdog Configuration File. Valid categories are ...	There is a typo in the category or it is out-of-date.	Update the Watchdog Configuration File and push it again.
Ignoring unrecognized subcategory ... from Watchdog Configuration File. Valid subcategories are ...	There is a typo in the subcategory or it is out-of-date.	Update the Watchdog Configuration File and push it again.
Could not parse any time information from ... in the Watchdog Configuration File. Examples of valid strings: ... Could not use ... as .... Using default ... instead.	The string is not a valid time.	Update the Watchdog Configuration File and push it again.
Could not parse 'CPU Usage Monitoring/threshold': must be a number, not ...	The threshold provided is not a valid number.	Update the Watchdog Configuration File and push it again.
Could not parse 'CPU Usage Monitoring/threshold': must be between 0% and 200%, not ...	The threshold provided is not a number.	Update the Watchdog Configuration File and push it again.

> ℹ️ The intervals and frequencies at which the Watchdog works cannot be enforced. That is, if the Watchdog is computing CPU usage and is due to compute the System Information Report, it will do so as soon as it is available.

14. Working with keys

This chapter explains how to use the facilities we provide to work with keys. There is often more than one way of performing a particular task. The methods available for working with keys are:

- KeySafe
- `generatekey` and related utilities
- The unit front panel.

You cannot generate keys from the front panel on the unit. You can generate keys on the client using the methods described in this chapter and view them on the module.

14.1. Generating keys

Whenever possible, generate a new key instead of importing an existing key. Because existing keys have been stored in a known format on your hard disk, there is a risk that the existing key has been compromised. Key material can also persist on backup media.

> **ℹ** Some applications can generate keys directly.

When you attempt to generate keys for a Security World that complies with FIPS 140 Level 3, you are prompted to insert an Administrator Card or Operator Card.

> **❗** Use Operator Cards for FIPS authorization. You should only use the Administrator Card Set for setting up new Security Worlds or performing administrative functions.

You may need to specify to the application, the slot you are going to use to insert the card. You need to insert the card only once in a session.

> **ℹ** For softcard protected key generation, you must use an Operator Card Set.

Generating a key creates both a key and a certificate request for the following application types:

- `embed` (OpenSSL)
- `kpm`

These requests are generated in PKCS #10 format with base-64 encoding.

14.1.1. Generating keys using the command line

Keys are generated using the command line with the `generatekey` utility. The `--generate` option creates a new key on the host computer that is protected either by the module or by an Operator Card set from the Security World. No key material is stored in an unencrypted form on the host computer.

When you generate a key with `generatekey`, choose a new identifier for the key and use whichever application type is appropriate. The key identifier can only contain digits, lower-case ASCII letters, and hyphens (`-`).

> ℹ️ Any uppercase letters you enter in the key identifier are converted to lowercase when the key is generated.

You can use `generatekey` in two ways:

- In interactive mode, by issuing commands without parameters and supplying the required information when prompted by the utility
- In batch mode, by supplying some or all of the required parameters using the command line (`generatekey` prompts interactively for any missing but required parameters).

In interactive mode, you can input `abort` at any prompt to terminate the process.

Batch mode is useful for scripting. In batch mode, if any required parameters are omitted, `generatekey` does not prompt for the missing information but instead will either use available defaults or fail. If you specify one or more parameters incorrectly, an error is displayed and the command fails.

If the Security World was created with audit logging selected then you can request that the usage of a key for cryptographic operations is logged in the audit log. By default only key generation and destruction is logged. For further information see Audit Logging.

To generate a key, use the command:

```
generatekey --generate [OPTIONS] <APPNAME> [<NAME>=<VALUE> ...]
```

In this command:

- `--generate` option specifies that this instance of `generatekey` is generating a key. Other options can be specified to perform tasks such as importing or retargeting keys. To see a list of options run the command `generatekey --help`.
- the `<APPNAME>` parameter specifies the name of the application for which the key is to be generated. For details of the available application types (`APPNAME`), Key application type (APPNAME).

- The `<NAME>=<VALUE>` syntax is used to specify the properties of the key being generated. For details of the available application types (`APPNAME`), see Key properties (`NAME=VALUE`).

For details of the available application types (`APPNAME`) and parameters that control other key properties (`NAME=VALUE`), see Key generation options and parameters and parameters.

In interactive mode, `generatekey` prompts you for any required parameters or actions that have not been included in the command. When you give the command:

1. Enter parameters for the command, as requested. If you enter a parameter incorrectly, the request for that information is repeated and you can re-enter the parameter.

2. When all the parameters have been collected, `generatekey` displays the final settings. In a FIPS 140 Level 3 compliant Security World, you are prompted to insert a card for FIPS authorization if no such card is present.

3. If prompted, insert an Administrator Card or an Operator Card from the current Security World.

4. If you want to protect the key with an OCS, you are prompted to insert the relevant cards and input passphrases, as required.

14.1.1.1. Example of key generation with generatekey

To generate a simple RSA key in batch mode, protected by module protection, use the command:

```
generatekey --generate --batch simple type=rsa size=2048 plainname=keya ident=abcd certreq=yes
```

The `generatekey` utility prompts you to insert a quorum of Operator Cards from the `operatorone` OCS. After you have inserted the appropriate number of cards, `generatekey` generates the key.

Although it is not explicitly specified, the created key is recoverable by default if OCS and softcard replacement is enabled for the Security World.

14.1.2. Generating keys with KeySafe

In order to generate a key with KeySafe, follow these steps:

1. Start KeySafe. (For an introduction to KeySafe and for information on starting the software, see Using KeySafe.)

2. Click the **Keys** menu button, or select **Keys** from the **Manage** menu. KeySafe takes you

to the **Keys** panel, which shows the keys in the Security World.

3. Click the **Create** button to open the **Generate Key** panel.

4. Select an application with which you want to use your key from the list, and then click the **Next** button. KeySafe takes you to the **Key Generation Parameters** panel.

5. Select and enter your desired parameters for key generation.

   The types of information that you need to provide on the **Key Generation Parameters** panel differs slightly depending on the application you selected on the **Generate Key** panel.

6. When you have supplied your desired key generation parameters, click the **Commit** button.

> In order to generate a key protected by a FIPS 140 Level 3 compliant Security World, you need authorization from an Operator Card or Administrator Card from the current Security World. Follow the onscreen instructions.

7. If you choose to generate a key that is protected by a smart card or softcard, KeySafe takes you to a panel from which you can load the protecting card or softcard. Follow the onscreen instructions, inserting any necessary Operator Cards and supplying any passphrases as needed.

8. KeySafe displays a message indicating that the key has been successfully generated. Click the **OK** button.

9. KeySafe returns you to the **Generate Key** panel, from which you can generate another key or choose another operation.

14.1.3. Generating NVRAM-stored keys

NVRAM key storage provides a mechanism for generating keys stored in a module's non-volatile memory and hence within the physical boundary of an nShield module. You can store only a few keys in this way: the number depends on the memory capacity of the module, the size of the key and whether the key has recovery data associated with it.

> We recommend that you *do not store keys in NVRAM unless you must do so to satisfy regulatory requirements.* NVRAM key storage was introduced only for users who must store keys within the physical boundary of a module to comply with regulatory requirements. NVRAM-stored keys provide no additional security benefits and their use exposes your ACS to increased risk. Storing keys in nonvolatile memory also reduces load-balancing and recovery capabilities. Because of these factors, we

recommend you always use standard Security World keys unless explic-
itly required to use NVRAM-stored keys.

When you generate an NVRAM-stored key, you must have sufficient nonvolatile memory
available in the module or the command fails.

> You need backup and recovery procedures, which must be consistent
> with regulatory requirements, to protect your NVRAM-stored keys. Do
> *NOT* use Remote Administration to back-up keys to a smart card, as, in
> transit, the keys would not be physically protected from access by the
> host system.

> An NVRAM-stored key can only be loaded successfully by using the `pre
> load` command-line utility on the generating module. Attempts to load
> such a key on other modules that have NVRAM fail with `UnknownID`
> errors.

We provide the `nvram-backup` utility to enable the copying of files, including NVRAM-stored
keys, between a module's nonvolatile memory and a smart card.

14.2. Importing keys

Importing a key takes an unprotected key stored on the host and stores it in the Security
World in encrypted form.

> We recommend generating a new key (or retargeting a key from within
> the Security World) instead of importing an existing key whenever possi
> ble. The import operation does not delete any copies of the key material
> from the host, and because existing keys have been stored in a known
> format on your hard disk (and key material can persist on backup
> media), there is a risk that an existing key has been compromised. It is
> your responsibility to ensure any unprotected key material is deleted. If
> a key was compromised before importation, then importing it does not
> make it secure again.

The following key types can be imported by the tools we provide:

- RSA keys in PEM-encoded PKCS #1 format (from a file). The PEM key that contains the
  key to import must not require a passphrase.
- DES, DES2 and Triple DES keys (entered in hex).

  > You cannot import keys into a Security World that complies with FIPS

> 140 Level 3. Attempting to import keys into a FIPS 140 Level 3 Security
> World returns an error.

This request is a PKCS #10 format request in base-64 encoding.

14.2.1. Importing keys from the command line

You can import keys using the `generatekey` utility. To import a key, give the command:

```
generatekey --import [<OPTIONS>] <APPNAME> [<NAME>=<VALUE> ...]
```

This command uses the following options:

Option	Description
`--import`	This option specifies key importation.
`<OPTIONS>`	You can specify particular options when running `generatekey` that control details of key importation.
`<APPNAME>`	This option specifies the name of the application for which the key is to be imported. This must be an application for which `generatekey` can generate keys.
`<NAME>=<VALUE>`	This specifies a list of parameters for the application.

For RSA keys, you can include `pemreadfile=filename` in the command to specify the file name of the PEM file that contains the key. Otherwise, you are prompted for this information during import.

In interactive mode, you are prompted for any required parameters or actions that have not been included in the command:

- Enter parameters, as requested. If you enter a parameter incorrectly, the request for that information is repeated and you can re-enter the parameter.
- If you want to protect the key with an OCS, you are prompted to insert the relevant cards and input passphrases, as required.
- If prompted, insert an Administrator Card or an Operator Card from the current Security World.

14.2.1.1. Example of key importation with generatekey

To import an RSA key stored in `/opt/projects/key.pem` (**Linux**) or `C:\projects\key.pem` (**Windows**) for use with an nShield native application and protect it with the Security World,

use the command:

Linux

```
generatekey --generatekey --import simple pemreadfile=/opt/projects/key.pem plainname=importedkey ident=abc
protect=module
```

Windows

```
generatekey --generatekey --import simple pemreadfile=C:\projects\key.pem plainname=importedkey ident=abc
protect=module
```

In this example, `generatekey` requires you to input `RSA` for the key type.

Although not explicitly specified, this key is, by default, recoverable if OCS and softcard replacement is enabled for the Security World.

14.2.2. Importing keys with KeySafe

Any user who has write access to the directory that contains the Security World can import a key.

In order to import a key with KeySafe, follow these steps:

1. Start KeySafe. (For an introduction to KeySafe and for information on starting the software, see Using KeySafe.)

2. Click the **Keys** menu button, or select **Keys** from the **Manage** menu. KeySafe takes you to the **Keys** panel.

3. Click **Import** to open the **Import Key** panel.

4. Select the application associated with the key that you want to import, and then click the **Next** button. KeySafe takes you to the **Key Import Parameters** panel.

5. Select and enter the desired parameters for the key that you want to import.

   The types of information that you need to provide on the **Key Import Parameters** panel will differ slightly depending on the application you selected on the **Import Key** panel.

6. When you have supplied parameters for the key that you want to import, click the **Commit** button.

7. If you choose to import a key that is protected by a smart card, KeySafe takes you to the **Load Operator Card Set** panel. Follow the onscreen instructions, inserting the required number of Operator Cards and supplying any passphrases as needed.

8. KeySafe displays a message indicating that the key has been successfully imported.

Click the **OK** button.

9. KeySafe returns you to the **Import Key** panel, from which you can import another key or choose another operation.

14.3. Listing supported applications with generatekey

To list supported applications, use the command:

```
generatekey --list-apps
```

14.4. Retargeting keys with generatekey

The `--retarget` option to takes an existing key in the Security World and makes it available for use by another application as if it had been expressly generated for use by that application. Because no key material is exposed during retargeting, this operation is as secure as generating a new key.

> When you retarget a key, `generatekey` does not remove the original key from the Security World. If required, you can use KeySafe to discard the original key.

When you retarget a key, you cannot change its protection method. You cannot change the key from module-protected to card-protected, or from card-protected to module-protected.

To retarget a key, use the command:

```
generatekey --retarget [<OPTIONS>] <APPNAME> [from-application=<appname>]
[from-ident=<keyident>]
```

In this command:

Option	Description
`--retarget`	This option specifies key importation.
`<OPTIONS>`	This option specifies any options to include when the command is run. Run the command `generatekey --help` for details about the available options.
`<APPNAME>`	This option specifies the name of the application for which the key is to be generated. This must be an application for which `generatekey` can generate keys.

Option	Description
`from-application=<appname>`	This option specifies the name of the application with which the key is currently associated.
`from-ident=<keyident>`	This option specifies the identifier of the key to be retargeted. You can find this identifier by using the `nfkminfo` command-line utility.

If `generatekey` cannot identify the key type for retargeting, you are prompted to specify the key type. Input the key type and press Enter.

14.5. Viewing keys

You can view existing keys in the Security World using KeySafe, the unit front panel, or the `nfkminfo` command-line utility.

14.5.1. Viewing information about keys on the unit front panel

You can view keys that have been created on the client on the same computer as the RFS with SEE machines. You cannot view other keys until they are transferred to the RFS.

To view keys:

1. From the main menu, select **Security World mgmt > Keys > List keys**.
2. Select the application to which the key belongs.
3. Select a key to view its full details.
4. If you wish, select **Verify key ACLs** to verify the key's ACL.

14.5.2. Viewing keys with KeySafe

In order to view a list of keys on the client computer on which you are running KeySafe, follow these steps:

1. Start KeySafe. (For an introduction to KeySafe and for information on starting the software, see Using KeySafe.)
2. Click the **Keys** menu button, or select **Keys** from the **Manage** menu. KeySafe takes you to the **Keys** panel, which lists all the keys in the Security World on this client computer. It displays the name of the key, the application for which it was created, the protection method that was used and whether the key is stored in NVRAM.

   If you click a key's listing, KeySafe displays additional information about that key, for

example, the application with which the key is used, whether or not the key is recoverable, and the key's name, creation date, hash, instance, and copy ID.

From the **Keys** panel, you can choose to:

- Create a new key (see Generating keys with KeySafe)
- import a key (see Importing keys with KeySafe)
- discard a key from the Security World (see Discarding keys)

14.5.3. Viewing keys using the command line

The `nfkminfo` command-line utility is used to list keys. To list the keys that have been created in the current Security World, use one of the following commands:

```
nfkminfo -k [<APPNAME>[<IDENT>]]
```

```
nfkminfo -l [<APPNAME>[<APPNAME>...]]
```

The `-k|--key-list` option lists keys only. The `-l|--name-list` option lists keys and their names.

With either option, `<APPNAME>` is an optional application name. If you specify an application name, `nfkminfo` lists only the keys for that application. Commonly, `<APPNAME>` is often one of:

- `custom`
- `embed`
- `pkcs11`
- `kpm`
- `kps`
- `mscapi` (**Windows**)
- `seeconf`
- `seeinteg`
- `simple`

You can also specify your own application names for *APPNAME* as appropriate to your system.

> For example, user-defined application names can be created by using the `nfkm` library to generate arbitrary keys.

With the `--name-list` option, `<IDENT>` is the key identifier.

The command `nfkminfo --key-list` returns output of the form:

```
Key summary - 4 keys:
 AppName appname    Ident <ident> AppName <appname>
Ident <ident> AppName <appname>
Ident <ident> AppName <appname>              Ident <ident>
```

To list information about a specific key, specify the `--key-list` option with an application and key identifier:

```
nfkminfo --key-list <appname> <ident>
```

This command returns output of the form:

```
Key AppName <appname> Ident <ident> BlobKA length   752
 BlobPubKA length      316
 BlobRecoveryKA length 868
 name                 "name"
 hash                 hash recovery             Enabled
 protection           CardSet
 other flags          PublicKey +0x0
 cardset              hash_ktBlobKA
 format               6 Token
 other flags          0x0
 hkm                  hash_km hkt               hash_kt hkr       none
BlobRecoveryKA
 format               8 Indirect
 other flags          0x0
 hkm                  none
 hkt                  none
 hkr                  hash_krBlobPubKA
 format               5 Module
 other flags          0x0
 hkm                  hash_km hkt               none
 hkr                  none
 No extra entries
```

To list keys and names, specify the `--name-list` option. The command `nfkminfo --name-list` returns output of the form:

```
Key summary - 30 keys
 in format key_<appname>_<ident> '<name>')
  key_appname_ident'name '
  key_appname_ident'name '
  key_appname_ident'name '
  key_appname_ident'name '
  key_appname_ident'name '
  key_appname_ident'name '
  key_appname_ident'name '
```

14.6. Verifying Key Generation Certificates with nfkmverify

The `nfkmverify` command-line utility verifies key generation certificates. You can use `nfkmverify` to confirm how a particular Security World and key are protected. It also returns some information about the Security World and key.

The `nfkmverify` utility compares the details in the ACL of the key and those of the card set that currently protects the key.

A key that has been recovered to a different card set shows a discrepancy for every respect that the new card set differs from the old one. For example, a key recovered from a 2-of-1 card set to a 1-of-1 card set has a different card-set hash and a different number of cards, so two discrepancies are reported. The discrepancy is between the card set mentioned in the ACL of the key and the card set by which the key is currently protected (that is, the card set mentioned in the key blobs).

A key that has been transferred from another Security World shows discrepancies and fails to be verified. Entrust recommends that you verify keys in their original Security World at their time of generation.

> If you must replace your Security World or card set, Entrust recommends that you generate new keys whenever possible. If you must transfer a key, perform key verification immediately before transferring the key; it is not always possible to verify a key after transferring it to a new Security World or changing the card set that protects it.

14.6.1. Usage

To verify the key generation certificates from the command line, run the command:

```
nfkmverify [-f|--force] [-v|--verbose] [-U|--unverifiable] [-m|--module=MODULE] [appname ident [appname ident
[...]]]
```

Optionally, the command can also include the following:

Option	Description
`-h|--help`	This option displays help for `nfkmverify`.
`-V|--version`	This option displays the version number for `nfkmverify`.
`-u|--usage`	This option displays a brief usage summary for `nfkmverify`.
`-m|--module=MODULE`	This option performs checks with module `MODULE`.

Option	Description
`-f|--force`	This option forces display of an output report that might be wrong.
`-U|--unverifiable`	This option permits operations to proceed even if the Security World is unverifiable. ⓘ If you need the `-U|--unverifiable` option, there may be some serious problems with your Security World.
`-v|--verbose`	This option prints full public keys and generation parameters.
`-C|--certificate`	This option checks the original ACL for the key using the key generation certificate. This is the default.
`-L|--loaded`	These options check the ACL of a loaded key instead of the generation certificate.
`-R|--recov`	This option checks the ACL of the key loaded from the recovery blob.
`--allow-dh-unknown-sg-group`	This option allows an operation to proceed even if a Diffie-Hellman key is using an unrecognized Sophie-Germain group.

14.7. Discarding keys

Discarding a key deletes the information about the key from the host disk. This option is only available in KeySafe.

If you have backup media, you can restore the information and the key becomes usable again. Likewise, if you have copies of the Security World data on several computers, erasing the data on one computer does not remove it from any other computer.

To destroy a key permanently you must either erase the OCS that is used to protect it or erase the Security World completely. There are no other ways to destroy a key permanently.

14.8. Restoring keys

We do not supply tools for restoring a key that has been discarded. However if you have kept a backup of the host data for the Security World, you can restore a key from the backup data.

ⓘ If you have NVRAM-stored keys, you must additionally ensure you have a backup of the key data stored on the relevant modules.

15. Using KeySafe

> **!** This section describes how to use the legacy **Java-based nShield KeySafe**. For documentation on the new generation product, see the **nShield KeySafe 5** documents on https://nshielddocs.entrust.com/.

KeySafe provides a GUI based interface to perform many of the main tasks required to use an nShield Security World. This appendix describes KeySafe, the Security World management tool. It includes information about:

- Starting KeySafe
- Using the graphical user interface (GUI) for KeySafe
- Using buttons to select and run operations
- Using the keyboard to navigate KeySafe
- KeySafe error reporting.

To perform Security World management, card-set management, and key management tasks using KeySafe, see the relevant chapters of this guide.

> **i** By default, KeySafe uses the same mechanisms and supports the same features and applications as the `generatekey` utility.

15.1. Setting up KeySafe

1. You must have Java JRE/JDK 1.7, 1.8 or 11. We recommend that you install Java before you install the Security World Software. On Linux, the Java executable must be on your path.

   Java software is available from http://www.oracle.com/technetwork/java/. If your security policy does not allow the use of downloaded software, these components can be obtained on removable media from Oracle or from your operating system vendor.

   > **i** After you have set up the path, check that you are using the correct Java version by running java with the `-version` option.
   >
   > Example:
   >
   > ```
   > >>java -version
   > java version "1.8.0_05"
   > Java(TM) SE Runtime Environment (build 1.8.0_05-b13)
   > Java HotSpot(TM) 64-Bit Server VM (build 25.5-b02, mixed mode)
   > ```

2. The Security World Software must be installed.

3. In the configuration file at `opt/nfast/kmdata/config/config` (**Linux**) or `%NFAST_KM-DATA%\config\config` (**Windows**), set the following port values in the `server startup` section:

```
nonpriv_port=9000
priv_port=9001
```

You must restart the hardserver after this change.

> **ⓘ** See the Installation Guide for more about ports and firewall settings.

15.2. Starting KeySafe

To start KeySafe:

Linux

Ensure that X11 is properly configured and running before starting KeySafe.

Start KeySafe by running the `/opt/nfast/bin/ksafe` script (assuming you installed the Security World Software in the default `/opt/nfast/` directory).

Windows

Start KeySafe from the Windows Start menu: **Start > Entrust nShield Security World > KeySafe**. You may need administrator privileges to run KeySafe.

The Windows KeySafe launcher checks that the components required to run KeySafe are installed. You will be prompted to install any missing components.

15.3. About the KeySafe window

The KeySafe window is divided into two areas:

- The sidebar (on the left), subdivided into:
    - The menu buttons (at the top of the sidebar)
    - The Security World status pane (at the bottom of the sidebar)
- The main panel area (on the right).

This layout is consistent throughout the KeySafe application.

15.3.1. Sidebar

The sidebar provides access to different parts of the KeySafe application (with the menu buttons) and also displays information about both the current Security World and your module or modules (with the Module Status tree).

> ℹ️ | The options listed below are also available from the **Manage** menu.

15.3.2. Menu buttons

There are five menu buttons at the top of the sidebar:

Menu button	Description
Introduction	Clicking the **Introduction** menu button opens the introductory panel that KeySafe displays at startup.
World	Clicking the **World** menu button opens the **World Operations** panel, from which you can: • Add modules to a Security World • Remove modules from a Security World. You cannot perform these operations on a module that is not in the pre-initialization mode. Do not use the **Initialize** option. Creating a Security World from KeySafe is deprecated.
Card Sets	Clicking the **Card Sets** menu button opens the **List Operator Card Sets** panel, from which you can: • Examine or change an Operator Card Set or its passphrase • Create a new Operator Card Set • Replace an Operator or Administrator Card Set • Discard an Operator Card Set.
Softcards	Clicking the **Softcards** menu button opens the **List Softcards** panel, from which you can: • Create a new softcard • Change or recover the passphrase on a softcard • Discard a softcard

Menu button	Description
Keys	Clicking the **Keys** menu button opens the **Keys** panel, from which you can: • Create a key • Import a key • Discard a key • View details of a key.

While KeySafe is executing a command, the menu buttons are disabled. Their normal functionality returns when the command is completed.

15.3.3. Menus

Three menu options are available from the menu bar at the top of the screen:

- **File**
  - **Exit** - displays a dialog asking whether you are sure you wish to quit KeySafe. Click **Yes** (or press the **Enter** key) to close KeySafe. Click **No** to close the dialog and return to your KeySafe session.
- **Manage**
  - **Introduction** - opens the **Introduction** panel. See Introduction button.
  - **World** - opens the **World Operations** panel. See World button.
  - **Card sets** - opens the **List Operator Card Sets** panel. See Cardsets button.
  - **Softcards** - opens the **List Softcards** panel. See Soft Cardsets button.
  - **Keys** - opens the **Keys** panel. See Keys button.
- **Help**
  - **About KeySafe** - opens the **About KeySafe** panel, which displays current version numbers for KeySafe, kmjava and nfjava. You will need to quote these version numbers if you contact Support about KeySafe.

15.3.4. Module Status tree

The Module Status tree, in the lower part of the KeySafe sidebar, displays information about the current Security World and your modules in the form of a tree diagram.

```
Security world status:

□
└─ ⚷ Security World
     ├── ● Cipher suite: Unknown
     ├── ● Initialized: Yes
     ├── ● Strict FIPS 140-2 Level 3: No
     ├── ● Key Recovery: Yes
     ├── ● Passphrase Recovery: No
     └─ 📁 Advanced
          ├── ● Initialized: Yes
          ├── ● Strict FIPS 140-2 Level 3: No
          ├── ● Key Recovery: Yes
          ├── ● Passphrase Recovery: No
          ├── ● RTC Key: Yes
          ├── ● NVRAM Key: Yes
          ├── ● SEE Debug Key: Yes
          ├── ● Open SEE Debugging: No
          └── ● FTO Key: Yes
└─ Module: #1
     ├── ● State: Operational:Usable
     ├── 🖥 Card Slot  0: Empty
     ├── 🖥 Card Slot  2: Empty
     ├── 🖥 Card Slot  3: Empty
     └─ 📁 Advanced
          ├── ● ESN: 56FF-28CF-56D8
          ├── ● nShield 20000
          ├── ● Firmware version: 13.2.1
          ├── ● Has RTC: Yes
          ├── ● Has NVRAM: Yes
          ├── ● RO Compatible: Yes
          └── ● RO Permitted: Yes
     └── ● Outside Security World
```

At the top of the Module Status tree is an icon representing the computer on which the run ning copy of KeySafe is installed. The name of this computer is shown to the right of the icon.

Below the computer icon in the Module Status tree are icons and text identifiers represent- ing the current Security World and your module(s). To the left of each icon is an expand/col lapse toggle, or node. By default, when KeySafe starts, these nodes are collapsed and show a minus sign. Click the node to display expanded information about the Security World or module. Click the node again to collapse this information.

15.3.4.1. Security World information

At the top level of the Security World tree, the following information is displayed:

- **Cipher suite** — the type of key protecting the Security World
- **Initialized** — whether the Security World is initialized (Yes or No)

- **Strict FIPS 140 Level 3** — whether the Security World is operating at FIPS 140 Level 3 (Yes or No)
- **Key Recovery** — whether key recovery is enabled (Yes or No)
- **passphrase Recovery** — whether passphrase recovery is enabled (Yes or No). For more information, see passphrase replacement.

When the **Advanced** node is expanded, the following additional information is displayed:

- **RTC Key** — whether a real-time clock authorization key has been generated (Yes or No)

  This is not applicable for nShield 5c.

- **NVRAM Key** — whether a non-volatile memory authorization key has been generated (Yes or No)
- **SEE Debug Key** — whether SEE debugging has been enabled (Yes or No)
- **Open SEE Debugging** — whether Open SEE debugging has been enabled (Yes or No)
- **FTO Key** — whether a foreign token key has been generated (Yes or No)

15.3.4.2. Module information

Module information may be displayed either inside or outside the Security World. Modules that have not been incorporated into a Security World will be shown beneath the **Outside Security World** node.

At the top level of the Module tree, the following information is displayed:

- The module's state, which is one of the following:

Mode	Description
PreInitMode	The module is in pre-initialization mode.
InitMode	The module is in initialization mode.
Operational:Useable	The module is in the current Security World and useable for key operations.
Operational:Unknown	The mode of the module cannot be determined.
Operational:Uninitialized	The module key is set and the module must be initialized before use.
Operational:Factory	The module key is set to the factory default.
Operational:Foreign	The module is from an unknown Security World.

Mode	Description
Operational:AccelOnly	The module is an acceleration-only module.
Operational:Unchecked	Although the module appears to be in the current Security World, KeySafe cannot find a module initialization certificate (a **module_*ESN*** file) for this module
Failed	The module has failed.
PreMaintMode	The module is in the pre-maintenance mode.
MaintMode	The module is in the maintenance mode.

- the status of the smart card reader slot(s).

> ℹ️ For FIPS 140 Level 3 Security Worlds, a **FIPS Auth Loaded** entry shows if an Administrator Card or Operator Card has been inserted to authorize an operation that requires a FIPS key.

The Module status tree has an **Advanced** folder that shows the following details when expanded:

- **ESN** — the module's electronic serial number (ESN), which is a unique identifier. You must quote a module's ESN if you need to contact Support. Keep a record of the ESN(s) associated with your module(s).
- the **HSM type** and model number
- **Firmware version** — the version of the module's firmware
- **Has RTC** — whether the module has a Real Time Clock (RTC)
- **Has NVRAM** — whether the module has nonvolatile memory (NVRAM).
- **RO Compatible** —
- **RO Permitted** —

15.3.5. Main panel area

The KeySafe main panel area is used to display information and options pertaining to a chosen operation. For example, clicking the **World** menu button takes you to the **World Operations** panel in the main panel area.

15.3.5.1. Navigation and command buttons

On each **Operations** panel there are a number of *navigation buttons*. Clicking a navigation button does *not* commit you to an action, but instead selects an operation and loads

another panel of additional information and options related to the selected operation. From the **World Operations** panel, for example, clicking the **Erase Module** navigation button does not itself erase a module, but rather loads the **Erase Module** panel.

On the next panel, the **Commit** button executes an operation, while the **Back** button returns to the previous panel. For example, on the **Erase Module** panel, clicking the **Commit** button will erase the module, while clicking the **Back** button will return to the **World Operations** panel.

> ⊗ Clicking the **Commit** button tells KeySafe to write or delete data: it is not necessarily possible to reverse such changes even if you subsequently cancel the operation. In some cases, clicking the **Commit** button causes KeySafe to display a dialog asking you to confirm your command. Such features help prevent you from accidentally destroying your data or keys.

Some panels require that you select options by means of radio buttons or that you enter data into text fields before clicking the **Commit** button. For example, if you click the **Reprogram Module** button on the **World Operations** panel, the next panel prompts you to choose whether the module can receive remote operator card shares.

Input may be in the form of radio buttons (allowing several options, one of which — the *default* — will be already selected) or text boxes (allowing you to enter text: no default value is set). If you click the **Commit** button without having entered data into a mandatory text field, or if KeySafe detects that the information you provided is inconsistent or invalid, KeySafe displays an error message. Click the message's **OK** button, and then provide or correct the necessary information.

After you successfully issue a command by clicking the **Commit** button, the menu buttons are disabled until the requested command is completed.

15.3.5.2. Navigating with the keyboard

The **Tab** key always takes you to the next field or button. If the cursor is not currently active in a text field, pressing the space bar or the **Enter** key activates the currently selected button (as if you had clicked it). Pressing the **Shift**-**Tab** button combination takes you to the previous field (if any) or deselects an automatically selected button (if any).

15.4. Errors

If KeySafe detects an error from which it cannot recover, it may display a Fatal Error mes-

sage.

15.4.1. Unable to establish KeySafe session.

Error

```
Please ensure that the hardserver is running and accepting TCP connections.
Click OK to exit.
```

Possible causes

- The hardserver is unable to receive TCP connections. The server program communicates with clients by using named pipes or TCP sockets.
- The hardserver is not running, or is physically disconnected.

Suggested solutions

- Check the hardserver configuration file settings: see server_startup.
- To restart the hardserver:
  1. Exit KeySafe
  2. Restart the server (as described in Stopping and restarting the client hardserver)
  3. Restart KeySafe.

15.4.2. Unable to generate key

*Error reported by nShield hardware module in response to `GenerateKeyPair`:

```
nFast error: UnknownFlag
```

Possible causes

Your hardware or firmware may not be up to date.

Suggested solutions

To update your firmware:

1. Exit KeySafe
2. Update the firmware as described in Upgrading the image file and associated firmware
3. Restart KeySafe

The firmware upgrade process destroys all persistent data held in a key-management mod-

ule. If your security system requires that the persistent data held in a key-management mod ule must survive intact during the upgrade or initialization of the key-management module, a backup and recovery mechanism of your `kmdata` (**Linux**) or `_%NFAST_KMDATA%` (**Windows**) directory must be implemented.

If you receive any error message titled **Unexpected Error**, contact Support with details of what you were doing, and the exact error message.

16. Warrant Management

16.1. Warrant management for nShield Solo and nShield Edge

You must use a Windows machine to manage warrants for nShield Edge HSMs.

This appendix describes how you can ensure that a suitable warrant is available to allow an nShield Remote Administration Card to be used with nShield Solo and Edge HSMs. To be able to use an nShield Remote Administration Card you need to ensure that:

- The appropriate firmware is installed on the nShield Solo or Edge HSM. (Firmware 2.61.2 or later)
- The nShield Solo or Edge HSM has a KLF2 warrant installed in the appropriate place.

16.1.1. Warranting steps for nShield Solo and nShield Edge

> **ℹ** You need an appropriate support contract to obtain a KLF2 warrant from Entrust.

Ensure v12.xx Security World Software has been installed on your host computer (to access the `nfwarrant` tool) and the nShield Solo or Edge HSM has Firmware 2.61.2 firmware or later installed.

You then need to carry out the following steps to ensure a suitable warrant is available

1. Check if the relevant module has the appropriate firmware.
2. Check if a warrant upgrade is required, if so, follow steps 3-6.
3. Generate a Certificate Signing Request (CSR) for the warrant.
4. Send the CSR to Entrust.

> **ℹ** Ensure that the ESN contained in the upgrade request is the one that belongs to the relevant module, for example, by running the `nfkminfo` command-line utility. See Displaying information about a Security World with nfkminfo for more about viewing an ESN.

5. Validate the warrant that you receive from Entrust to ensure that it matches the sent request.
6. Install the warrant.

16.1.2. nfwarrant command-line utility

The `nfwarrant` command-line utility enables you to carry out all of the relevant warrant steps. It is used to:

- Identify modules that have the appropriate firmware and KLF2 key
- Identify modules that need their KLF2 key to be warranted by Entrust
- Generate a warrant upgrade request for a specific module, as required
- Install an upgraded warrant
- List KLF2 warrants

Usage:

```
nfwarrant [--help] [--list] [--check] [--warrant] [--csr] [--details= FILE]
[--install= FILE] [--req= MODULE] [--out= FILE] [--verbose] [--version]
```

Options:

Option	Description
-h|--help	Displays the options you can use with the utility.
--list	List ESNs of installed warrants
--check	List ESNs of known modules and their warrant state
--warrant	Perform warrant operations
--csr	Perform CSR operations
--details=<FILE>	Display the module ESN found in the CSR/warrant *<file>*
--install=<FILE>	Install the warrant from *<file>*
--req=<MODULE>	Request a warrant CSR for the given module number/ESN
--out=<FILE>	Save the new requested CSR to *<file>*
--verbose	Print extra information about CSR and warrant files
--version	Print the version number of the nfwarrant tool

16.1.2.1. Check the available hardware

```
$ nfwarrant --check
```

The following is an example output:

```
1 XXXX-XXXX-E0D2 Local, Warrant installed
2 XXXX-XXXX-CF11 Local, Warrant upgrade request possible
3 XXXX-XXXX-F1F2 Local, Warrant upgrade not supported
4 XXXX-XXXX-213B Remote, Warrant upgrade not required
```

In this example:

- (1) already has a relevant warrant installed.

- (2) is available for a warrant upgrade.

- (3) cannot be upgraded. For example, the appropriate firmware is not installed.

- (4) no warrant upgrade is required. The module is an nShield Connect.

16.1.2.2. Generate a warrant upgrade request for nShield Solo

Run the following command:

```
$ nfwarrant --csr --req <module>
```

The following is an example output, displaying the location of the resultant upgrade request for a module with ESN XXXX-XXXX-CF11:

```
CSR written to 'C:\ProgramData\nCipher\Key Management Data\warrants\csr_XXXX-XXXX-CF11'
```

Ensure that the ESN in this request file is the correct one and send the file to Entrust to be signed.

16.1.2.3. Validate the warrant you receive from Entrust

1. Run the following command:

    ```
    $ nfwarrant --warrant --details <file>
    ```

    The following is an example output:

    ```
    Warrant details: Filename: XXXX-XXXX-CF11 ESN: XXXX-XXXX-CF11 Keytype: ECDSAPublic Curve: NISTP521
    ```

2. Compare the ESN in the file received from Entrust with the one in the original request, by running the following command:

    ```
    $ nfwarrant --csr --details <file>
    ```

The following is an example output:

```
XXXX-XXXX-CF11
```

16.1.2.4. Install a warrant for nShield Solo

Run the following command:

```
$ nfwarrant --warrant --install <file>
```

<file> is the signed warrant provided by Entrust.

16.2. Warrant management for nShield Connect + and nShield Connect XC

You do not need to manage the warrants for nShield Connect HSMs. They copy the warrant back to the host or RFS on startup.

16.3. Warrant management for nShield 5s and nShield 5c

You do not need to manage the warrants for nShield 5s or nShield 5s. Entrust supplies these HSMs with the required warrants pre-installed and stored within the module. The Security World software fetches warrants from the module when they are needed.

This includes a KLF2 and a KLF3 warrant. The KLF3 warrant is currently unused and is installed in preparation for multi-tenant systems.

To view the warrants installed on a module, run `retrievewarrants`. This stores a copy of the warrants in the host file system.

17. Supplied utilities

This appendix describes the executable command-line utilities (utilities) that you can use for performing various configuration and administrative tasks related to your module.

These utilities exist in the `bin` subdirectory of your Security World Software installation. Unless noted, all utilities have the following standard help options:

- `-h|--help` displays help for the utility.
- `-v|--version` displays the version number of the utility.
- `-u|--usage` displays a brief usage summary for the utility.

17.1. Utilities for general operations

Use the utilities described in this section to:

- Check the module configuration and verify that it functions as expected.
- Obtain statistics for checking the performance of the module.

17.1.1. `enquiry`

Obtain information about the hardserver (Security World Software server) and the modules connected to it.

- Check if the software has been installed correctly
- Check the firmware version
- Check if the Remote Operator feature is enabled
- Check if the Serial Console feature is available
- Check the hardware status of internal security modules

See Testing the installation for more information.

17.1.2. `checkmod`

Check modulo exponentiations performed on the module against the test data located in `opt/nfast/testdata` (**Linux**) or `%NFAST_HOME%\testdata` (**Windows**).

17.1.3. `cfg-mkdefault`

Create a default client configuration file for the hardserver configuration sections.

17.1.4. `cfg-reread`

Load the hardserver configuration from the configuration file.

17.1.5. `fet`

- Activate features
- View the status of features
- Verify that a feature has been successfully enabled on a connected module

To view the status of features, run the tool without a smart card. If a FEM card is not present, or if any of the features are not enabled successfully, the utility prompts you to indicate what to do next.

> **ℹ** To enable features, and view the status of or verify features on an nShield HSM, use the front panel rather than the `fet` utility.

For more information, see Enabling optional features

17.1.6. `ncdate`

View, set, and update the time on a module's real-time clock.

17.1.7. `ncversions`

Obtain and verify the versions of the Security World Software components that are installed. This utility lists the following information:

- Versions of all components, irrespective of whether they are installed individually or as part of a component bundle
- Version of each component bundle

17.1.7.1. `nfdiag`

Obtain information about the module and the host on which it is installed. This diagnostic utility can save information to either a `ZIP` file or a text file.

For more information, see nfdiag: diagnostics utility.

> **ℹ** | Run this utility only if requested to do so by Support.

17.1.8. `nopclearfail`

Clear an HSM, put an HSM into the error state, retry a failed HSM, or change the HSM mode.

You must use a privileged connection to use this utility with the following parameters:

- change the mode of the HSM (`nopclearfail -I/M/O`)
- Clear the module (`nopclearfail -c`)

17.1.9. `nvram-backup`

Copy files between a module's NVRAM and a smart card, allowing files to be backed up and restored.

17.1.10. `nvram-sw`

View and modify information about NVRAM areas.

17.1.11. `pubkey-find`

Obtain information of the public key from a certificate or certificate request (in a Base-64 encoded PEM file).

17.2. `randchk`

Run a universal statistical test on random numbers returned by the module.

17.2.1. `rtc`

View and set the module's real-time clock.

The 'rtc' command is not used for obtaining and setting the RTC on nShield 5c.

By default, `rtc` reads the real-time clock of module 1.

- `--adjust`: The module uses the difference between its idea of the current time and the new time, together with how long it's been since the clock was last set, to compute

how much its clock is drifting.

- `--set-clock`: The module's clock is set to either TIME, if it is provided as a list of six integers separated by non-digit characters, or to the host's current time.

17.2.2. `slotinfo`

- Obtain information about tokens in a module
- Format a smart card

17.2.3. `snmpbulkwalk snmpget snmpgetnext snmptable snmpset snmptest snmptranslate snmpwalk`

Obtain system, module, connection and software information from the SNMP agent.

For more information, see Using the SNMP command-line utilities.

17.2.4. `stattree`

Obtain statistics gathered by the Security World Software server and modules.

For more information, see stattree: information utility.

Linux

Archive the existing hardserver log from `/opt/nfast/log/hardserver.log` and re-open as a fresh log file.

When run with no arguments, it will automatically archive the existing log to `/opt/nfast/log/archive/hardserver.DATETIME.log` (where *DATETIME* is the current date and time). The directory `/opt/nfast/log/archive/` is created if it does not already exist.

Optionally, a single argument can be provided with the full file name to archive the existing hardserver log to.

This script must be run as root.

Windows

Extract Windows event log entries and output them to the console or a text file.

As required, specify:

- `-s \| --source`: The event log source. The default is the `nCipherlog`

- `-c \| --count`: The number of records read from the event log. The default is `10000`
- `-f \| --file`: The output filename.

17.3. Hardware utilities

Use the following utilities to manage the firmware installed on an nShield HSM.

17.3.1. `fwcheck`

Verify the firmware installed on a module.

17.3.2. `nfloadmon`

Upgrade the module monitor and firmware of nShield PCIe and network-attached HSMs.

17.4. Test analysis tools

Use the following utilities to test the cryptographic operational behavior of a module.

> All the listed utilities, except the `floodtest` utility, are supported only on FIPS 140 Level 2 Security Worlds.

Utility	Enables you to...
`cryptest`	Test all defined symmetric cryptographic mechanisms.
`des_kat`	Perform DES known-answer tests. This utility indicates if any of them fail.
`floodtest`	Perform hardware speed-testing by using modular exponentiation.
`kptest`	Test the consistency of encryption and decryption, or of signature and verification, with the RSA and DSA algorithms.
`ncthread-test`	Stress test modules and test nCore API concurrent connection support.
`perfcheck`	Run various tests to measure the cryptographic performance of a module. For more information, see perfcheck: performance measurement checking tool.
`sigtest`	Measure module speed using RSA or DSA signatures or signature verifications.
`ncperftest`	Test the performance of various crypto commands using attached nShield hardware. Available since v12.10 it contains all the functionality in `sigtest` and `floodtest` as well as several new features and greater accuracy and throughput capability in performance management.

17.5. Security World utilities

Use the utilities described in this section to:

- Set up and manage Security Worlds.
- Create and manage card sets and passphrases.
- Generate keys and transfer keys between Security Worlds.

Utility	Enables you to...
`bulkerase`	Erase multiple smart cards including Administrator Cards, Operator Cards, and FEM activation cards, in the same session. ⓘ Do not use the `bulkerase` utility to erase Administrator Cards from the current Security World.
`cardpp`	Change, verify, and recover a passphrase of an Operator Card. For more information, see: • Verifying the passphrase of a card with cardpp. • Changing known passphrases with cardpp. • Changing unknown or lost passphrases.
`createocs`	Create and erase an OCS. For more information, see: • Creating an Operator Card Set from the command line. • Erasing cards from the command line.
`initunit`	Initialize an nShield module. For more information, see Erasing a module with initunit.
`generatekey`	Generate, import, or retarget keys. This utility is included in the `Core Tools` bundle, which contains all the Security World Software utilities. For more information, see: • Generating keys with the command line. • Importing keys from the command line. • Example of key generation with generatekey, for an example of key generation in batch mode. • Example of key importation with generatekey, for an example of importing an RSA key. • Listing supported applications with generatekey. • Retargeting keys with generatekey.
`kmfile-dump`	Obtain key management information from a Security World's key management data file.
`migrate-world`	Migrate existing keys to a destination Security World. For more information, see Security World migration.

Utility	Enables you to...
`mkaclx`	Generate non-standard cryptographic keys that can be used to perform specific functions, for example, to wrap keys and derive mechanisms. This utility includes options that are not available with the `generate-key` utility. ⓘ Ensure that you run the `mkaclx` utility with the options that are appropriate for your security infrastructure. If the appropriate options are not chosen, the security of existing keys might potentially be compromised.
`new-world`	Create and manage Security Worlds on nShield modules. You must use a privileged connection to use this utility with the following parameter: • Initialize the HSM (`new-world -e/i/l`)
`nfkmcheck`	Check Security World data for consistency.
`nfkminfo`	Obtain information about a Security World and its associated cards and keys. For more information, see: • Displaying information about a Security World with nfkminfo. • Viewing card sets from the command line. • Viewing softcards with nfkminfo. • Viewing keys using the command line. • nfkminfo: information utility.
`nfkmverify`	Perform Security World verification. For more information, see Verifying Key Generation Certificates with nfkmverify.
`postrocs`	Transfer PKCS #11 keys to a new card set in the new Security World. When transferring keys by using either the `key-xfer-im` utility or the `migrate-world` utility, run the `postrocs` utility if there are any PKCS #11 keys that are protected by OCSs. ⓘ PKCS #11 keys either have `keys_pkcs_um` or `key_pkcs_uc` as the prefix.

Utility	Enables you to...
ppmk	• Create and manage softcards. Use this utility to: • View details of a softcard • Create and delete a softcard • View, change, and recover the passphrase of a softcard For more information, see: • Creating a softcard with ppmk. • Erasing softcards with ppmk. • Viewing softcards with ppmk. • Verifying the passphrase of a softcard with ppmk. • Changing known softcard passphrases with ppmk. • Replacing unknown passphrases with ppmk.
preload	Load keys into a module before an application is run in another session.
racs	Create a new ACS to replace an existing ACS. For more information, see Replacing an Administrator Card Set using racs.
rocs	• Restore an OCS from a quorum of its cards • Restore softcards For more information, see: • Replacing OCSs or softcards with rocs. • Using rocs from the command line.

17.6. CodeSafe utilities

Use the following helper utilities to develop and sign SEE machines. For more information about these utilities, see the *CodeSafe Developer Guide*.

Utility	Enables you to...
elftool	Convert ELF format executables into a format suitable for loading as an SEE machine.
hsc_loadseemachine	Load an SEE machine into each module that is configured to receive one, then publishes a newly created SEE World, if appropriate.
loadsee-setup	Set up the configuration of auto-loaded SEE machines.
modstate	View the signed module state.

Utility	Enables you to...

`see-sock-serv`
 	
`see-stdioe-serv`
 	
`see-stdioesock-serv`
 	
`see-stdoe-serv`
	Activate or enable standard IO and socket connections for SEE machines using the `bsdlib` architecture.

`tct2` (Trusted Code Tool)	Sign, pack, and encrypt file archives so that they can be loaded onto an SEE-ready nShield module.

17.7. PKCS #11

> ❗ Do not use PKCS #11 to perform any task that requires an Administrator Card. Use the equivalent nShield utilities instead.

Use the following utilities to manage the interfaces between the PKCS #11 library and the module.

Utility	Enables you to...
`ckcerttool`	Import a certificate as a PKCS #11 `CKO_CERTIFICATE` object of type `CKC_X_509`, and optionally, associate it with the corresponding private key.
`ckcheckinst`	Verify the installation of the nShield PKCS #11 libraries. For more information, see Checking the installation of the nCipher PKCS #11 library.
`ckimportbackend`	Generate keys for use with PKCS #11 applications. When you run the `gener-atekey` utility to generate PKCS #11 keys, the `ckimportbackend` utility is executed in the background. > ℹ Do not run this utility unless directed to do so by Support.
`cknfkmid`	View values of attributes of PKCS #11 objects.
`ckshahmac`	Perform a PKCS #11 test for vendor-defined `SHA1_HMAC` key signing and verification capabilities.
`cksigtest`	Measure module signing or encryption speed when used with nShield PKCS #11 library calls.

The Security World software enables you to use the following additional PKCS #11 utilities. For more information about these utilities, see the *Cryptographic API Integration Guide*.

Utility	Enables you to…
`ckinfo`	View PKCS #11 library, slot, and token information. Use this utility to verify that the library is functioning correctly.
`cklist`	View details of objects on all slots. If invoked with a PIN argument, the utility lists public and private objects. If invoked with the `-n` (`--nopin`) option, the utility lists only the public objects. This utility does not output any potentially sensitive attributes, even if the object has `CKA_SENSITIVE` set to `FALSE`.
`ckmechinfo`	View details of the supported PKCS #11 mechanisms provided by the module.
`ckrsagen`	Test RSA key generation. You can use specific PKCS #11 attributes for generating RSA keys.
`cksotool`	Create a PKCS #11 Security Officer role, and manage its PIN.

17.8. Utilities for network-attached HSMs

The utilities described in this section are used with nShield network-attached HSMs only. Use these utilities to:

- Create and manage client configuration files.
- Enroll nTokens with an nShield HSM.
- Set up a Remote File System (RFS) and synchronize Security World data between an nShield HSM and the RFS.
- Administer an nShield HSM without using the front panel
- Configure NTP.

Utility	Enables you to…
`anonkneti`	View the ESN and HK_{NETI} key hash from a module identified by its IP address. For more information, see Configuring the remote file system (RFS).
`cfg-pushnethsm`	Copy a specified configuration file from a remote file system to the file system on a specified module. Some changes propagated with `cfg-pushnethsm` need further actions. For example, you have to clear the module after changing the dynamic slot configuration. For more information, see: • Remote configuration of additional clients. • About user privileges. • CodeSafe applications.

Utility	Enables you to...
`cfg-pushntp`	Configure time synchronisation on the nShield HSM, using NTP. For more information, see Configuring NTP in the nShield HSM.
`config-serverstartup`	Edit the `[server_startup]` section of the configuration file for the client's hard server to enable or disable TCP sockets. For more information, see: • After software installation. • config-serverstartup.
`nethsmadmin`	Administer an nShield HSM without using the front panel. Options include: • Check the Security World files on a specified nShield HSM. • Copy Security World files from the RFS to the nShield HSM. • Command the specified nShield HSM to reboot. This restarts the hard-server. • Command the nShield HSM to upgrade using the specified image file from its RFS. • Retrieve a list of image files available on the RFS. • Retrieve a list of feature certificates available on the RFS for a specified nShield HSM. • Command the nShield HSM to apply a specified feature certificate from the RFS. • Erase the Security World on the nShield HSM and re-initialize the HSM. • Get the date and time on the nShield HSM. • Set the date and time on the nShield HSM. • Enable dynamic features, including client licenses remotely. You must use a privileged connection to use this utility with the following parameters: • Reboot the HSM (`nethsmadmin -r`) • Erase the Security World (`nethsmadmin -e`) • Upgrade the HSM firmware (`nethsmadmin -i`) For more information, see: • Using nethsmadmin to copy a Security World to an nShield HSM and check the current version. • Upgrading the image file and firmware from a privileged client. • Remotely enabling dynamic feature certificates including nShield HSM client licenses.

Utility	Enables you to…
`nethsmenroll`	Edit the local hardserver configuration file to add the specified nShield HSM unit. As an alternative to hand-editing a client's configuration file, you can run this utility on a client to configure it to access an nShield HSM. For example: • Enroll an HSM, without needing to restart the hardserver • Unenroll an HSM (`nethsmenroll -r`), then restart the hardserver to update the information about the HSM estate For more information, see: • nethsm_imports. • Configuring the unit to use the client. • nethsmenroll.
`ntokenenroll`	Enroll a locally attached nToken with an nShield HSM. This utility installs the Electronic Serial Number (ESN) of the nToken within the client configuration file and displays the module's ESN and the hash of the key to be used in nToken authentication. For more information, see Configuring the unit to use the client.
`rfs-setup`	Create a default RFS hardserver configuration. Run this utility when you first configure the RFS. For more information, see: • Configuring the remote file system (RFS). • Setting up client cooperation.
`rfs-sync`	Synchronize the Security World data between a cooperating client and the RFS. This utility is run on the client. For more information, see: • Setting up client cooperation. • rfs-sync. ℹ️ You can use this utility with nShield modules if an nShield HSM is present.

17.9. MSCAPI utilities (Windows)

Use the following utilities to migrate from Windows registry-based CSP container storage to the new CSP formats. These utilities also enable you to manage the interfaces between the MSCAPI library and the module.

For more information about these utilities, see Utilities for the CAPI CSP.

> ℹ️ Utility names that end with 64 run only on 64-bit version of Microsoft Windows. All other utilities run on both 32-bit and 64-bit versions of Microsoft Windows.

Utility	Enables you to...
cspcheck cspcheck64	Check that CSP container files and keys in the %NFAST_KMDATA% directory are intact and uncorrupted and that the referenced key files exist.
cspimport cspimport64	Insert keys manually into existing CSP containers. For more information, see Installing the CAPI CSP.
cspmigrate cspmigrate64	Move CSP container information for an existing Security World from the registry into the Security World.
cspnvfix cspnvfix64	Regenerate the NVRAM key counter area for a specified nShield CSP key.
csptest csptest64	Test the installed Cryptographic Service Providers.
csputils csputils64	Obtain detailed information about CSP containers. ℹ️ You must have Administrator privileges to view or delete machine containers or containers that belong to other users.
keytst keytst64	Create, test, and display information about keys and CSP key containers.
configure-csp-poolmode configure-csp-poolmode64	Configure HSM Pool mode for the nShield CAPI CSP.

17.10. CNG (Windows)

Use the following helper utilities to manage keys and the interfaces between the CNG library and the HSM. For a list of utilities specific to the nShield CNG CSP, see Utilities for the CAPI CSP.

> ℹ️ Utility names that end with 64 run only on 64-bit version of Microsoft Windows. All other utilities run on both 32-bit and 64-bit versions of Microsoft Windows.

Utility	Enables you to…
cngimport	Migrate Security World, CAPI and CNG keys to the Security World Key Storage Provider. For more information, see: • Importing a Microsoft CAPI key into the Security World Key Storage Provider. • Importing a Microsoft CNG key into the Security World Key Storage Provider. • Importing a Security World key into the Security World Key Storage Provider • cngimport.
cnginstall32 cnginstall (nShield CNG provider installer utility)	Remove or reinstall the provider DLLs and associated registry entries manually. For more information, see cnginstall.
cnglist32 cnglist	View information about CNG providers. For more information, see: • Migrating keys for CNG. • Importing a Microsoft CAPI key into the Security World Key Storage Provider. • Importing a Microsoft CNG key into the Security World Key Storage Provider. • Importing a Security World key into the Security World Key Storage Provider. • cnglist.
cngregister (nShield CNG provider registration utility)	Unregister and re-register the nShield providers manually. For more information, see: • Registering the CNG CSP. • Unregistering or reregistering the CNG CSP. • Uninstalling or reinstalling the CNG CSP. • cngregister.
cngsoak cngsoak64 (nShield CNG soak tool)	Evaluate the performance of signing, key exchange, and key generation by using a user-defined number of threads. For more information, see cngsoak.

Utility	Enables you to…
`ncsvcdep` (nShield Service dependency tool)	Configure service-based applications such as Microsoft Certificate Services and IIS to use the nShield CNG CSP. Use this tool to add the nFast Server to the dependency list of such services. For more information, see: • Uninstalling or reinstalling the CNG CSP. • `ncsvcdep`.

`configure-csp-poolmode`
 	
`configure-csp-poolmode64`
	Configure HSM Pool mode for the nShield CNG CSP. For more information, see configure-csp-poolmode.

17.11. Developer-specific utilities

Use the following utilities to ensure that the HSMs are functioning as expected and to test the cryptographic functionality at the nCore level.

Utility	Enables you to…
`pollbare`	Obtain information about state changes. The functionality of this test utility depends on whether the server or an HSM supports nCore API poll commands. ℹ️ To know if your server or HSM supports nCore API poll commands, run the `enquiry` utility.
`trial`	Test the nCore API commands. You can use this utility interactively or from a script file.

17.12. Utilities that require a privileged connection

You must be a privileged user, that is, use a privileged connection to the HSM, to run certain utilities with certain parameters.

Utility	Use case
`nopclearfail -I/M/O`	Change the mode of the HSM
`nopclearfail -c`	Clear the module
`nethsmadmin -r`	Reboot the HSM
`nethsmadmin -e`	Erase the Security World
`nethsmadmin -i`	Upgrade the HSM firmware
`new-world -e/i/l`	Initialize the HSM

18. Using silent installations

This appendix describes how to use the command line for software installation and uninstallation for automation.

When you follow the standard installation instructions for Security World Software, the `setup.msi` installer runs automatically when you place the Security World Software installation media in the optical disc drive. You then follow the on-screen instructions from the installer to configure your installation.

However, if you run the `setup.msi` installer from the command line, you have the option to define the components you want to install via the Windows command prompt. This allows your installations to run 'silently', without the need for further interaction with the installer.

> ℹ️ See the Installation Guide for more information about installing Security World Software.

18.1. Installing using the silent install functionality

The Windows Installer (MSI) has the ability to install software without user interaction. This is useful for automation purposes.

Before starting, please ensure that:

- Any previously installed Security World Software is uninstalled
- If the directory `C:\Program Files\nCipher` or `C:\Program Files\nCipher` exists, it is deleted
- If the directory `C:\ProgramData\nCipher` exists, it is renamed or deleted

To install the nShield Software using the silent install functionality:

1. Log in as Administrator or as a user with local administrator rights.
2. Place the Security World Software installation media in the optical disc drive. If the installer runs automatically, quit the installer.
3. Open a Command Prompt, and run the command:

```
msiexec /i <PATH_TO_MSI> /quiet /forcerestart
```

This installs the nShield Security Software to the default installation directory, `%PROGRAM FILES%\nCipher\nfast\`, and restarts the machine.

To generate a verbose install log, add `/l*v <path-to-file.txt>` to the command after

the `/quiet` argument. For example:

```
msiexec /x E:\setup.msi /quiet /l*v C:\users\USER_NAME\installLog.txt> /forcerestart
```

This creates a log file in the specified directory.

18.2. Uninstalling using the silent install functionality

To uninstall the nShield Security Software using the silent uninstall functionality:

1. Log in as Administrator or as a user with local administrator rights.
2. Place the Security World Software installation media in the optical disc drive. If the installer runs automatically, quit the installer.
3. Open a Command Prompt, and run the command:

```
msiexec /x <PATH_TO_MSI> /quiet /forcerestart
```

This uninstalls the nShield Security Software packages and restarts the machine.

To generate a verbose uninstall log, add `/l*v <path-to-file.txt>` to the command after the `/quiet` argument. For example:

```
msiexec /x E:\setup.msi /quiet /l*v C:\users\USER_NAME\uninstallLog.txt> /forcerestart
```

This creates a log file in the specified directory.

19. Using nShield commands from PowerShell

PowerShell is a powerful console tool for scripting operations on Windows. nShield applications can be run from PowerShell, locally or remotely, interactively or non-interactively (batch mode). nShield library code implements a set of commands for reading text and passphrases.

19.1. Install and configure PowerShell

1. Install PowerShell, see https://docs.microsoft.com/en-us/powershell/.
2. Ensure that executing PowerShell scripts is enabled in the system.

   Script execution is not enabled by default on Windows clients. The default permissions usually allow script execution on Windows Server operating systems, but it may be necessary to enable this in a custom Windows Server configuration.

   Open PowerShell and set the signing property and scope of script execution.

   The support files for running nShield commands from PowerShell are Authenticode-signed, so the execution can be restricted to only signed scripts:

   ```
   Set-ExecutionPolicy -ExecutionPolicy AllSigned
   ```

   If unsigned PowerShell scripts are to be executed, you may want to relax this so that locally created scripts can be run without signing:

   ```
   Set-ExecutionPolicy -ExecutionPolicy RemoteSigned
   ```

   The above commands can be run with the additional parameter `-Scope CurrentUser` to restrict the changes to the currently logged-in user. For example to permit the current user to run locally-created scripts, run:

   ```
   Set-ExecutionPolicy -ExecutionPolicy RemoteSigned -Scope CurrentUser
   ```

3. By default, PowerShell commands are created for all executables in `$env:NFAST_HOME\bin`. If there are additional directories containing nShield executables that you wish to include in the nShield PowerShell module support, you can specify those directories with a semicolon separated list of paths in the `NC_PS_ADDITIONAL_DIRECTORIES` environment variable.

   ```
   $env:NC_PS_ADDITIONAL_DIRECTORIES="C:\Path1;C:\Path2"
   ```

4. To load support for nShield commands in PowerShell, import the `nShieldTools.psd1` module located in `$env:NFAST_HOME\bin`.

```
Import-Module 'C:\Program Files\nCipher\nfast\bin\nShieldTools.psd1'
```

The support module is installed in the active shell.

5. To load the `nShieldTools.psd1` module automatically every time PowerShell is opened, you can add it to your PowerShell profile with the `Add-nShieldToProfile` command, included in the `nShieldTools.psd1` module.

Parameter	Description
AllHosts	(Recommended) Profile is available for all PowerShell hosts, for example both ConsoleHost and ISE, rather just the currently running host.
AllUsers	Profile is available for all users on the local machine rather than just the current user.

Example:

```
Add-nShieldToProfile -AllHosts
```

19.2. Calling nShield commands at the PowerShell prompt

nShield commands can be called with their usual names from PowerShell, for example `enquiry`, `cardpp`, `generatekey`, or `new-world`. Aliases for the `.exe` variants, for example `new-world.exe`, are also registered so that the nShield executables in the PATH are only called with PowerShell support.

Do not call the executables directly (for example, do not call & 'C:\Program Files\nCipher\nfast\bin\new-world.exe') because this will not enable the PowerShell support.

Command-line parameters to nShield PowerShell commands can be provided in the same way as the corresponding command under regular Windows consoles. If you can run `cardpp --check` in a regular Windows console, you can run `cardpp --check` in PowerShell.

The global variable `$LASTEXITCODE` of PowerShell contains the exit code (0 for success) immediately after execution of the nShield command.

19.3. PowerShell modes: interactive and batch

nShield commands can be run in either interactive or batch (non-interactive) mode.

In the default **interactive mode**, the output is displayed incrementally on the screen in the PowerShell host user interface. UI prompts are displayed when the command attempts to read user input, for example passphrases or confirmations. Output is not written to a Power Shell pipeline in this mode.

Batch mode is intended for automation. Commands can be run from a script and an output can be redirected to a file. Batch mode does not prompt for input in the host UI. Input can be supplied programmatically with a PowerShell input pipeline. If input is needed but no suit able pipeline was supplied, the command fails rather than stall program execution to wait for user interaction. Standard output and standard error text printed by the underlying nShield program is written to output and error pipelines that can be redirected or piped. If the program fails, that is, it returns a non-zero exit code, the error code is also thrown as an exception that appears in the error pipeline. In batch mode, output and error texts are visible only at the very end of execution, because such texts are objects that the command returns to the pipeline instead of writing them incrementally to the host UI.

If the mode is not explicitly set, PowerShell normally defaults to interactive mode. However, if any input pipeline objects are supplied, PowerShell defaults to batch mode. You can there fore switch to batch mode without setting it explicitly simply by piping `$null`:

```
$null | enquiry > enquiry.txt
```

You can change the default mode to batch mode by setting the `NC_PS_INTERACTIVE` environ ment variable:

```
$env:NC_PS_INTERACTIVE=0
```

Mode change commands provided in the `nShieldTools.psd1` module:

Command	Notes
`Set-nShieldBatchMode`	
`Set-nShieldInteractiveMode`	
`Reset-nShieldCommandMode`	Can be used to restore the default PowerShell behavior based on presence or absence of pipeline input.

To restrict a setting to a particular PowerShell scope, you can use the PowerShell variable `$nShieldInteractiveCommandMode`, which can be set to `$True` or `$False`.

19.4. Input pipelines

In both interactive and batch mode, nShield commands support input pipelines with the PowerShell pipe ("|") syntax. The input pipeline can be used to automate the execution of nShield commands that would otherwise have to prompt for user input. For example, a passphrase check on an OCS card can be performed automatically by executing the following command:

```
Set-nShieldBatchMode
$passphrase | cardpp --check
```

`$passphrase` is a variable in the command or script, and contains the card's passphrase.

Multiple values can be supplied to provide the input to successive prompts. For example, the `generatekey` command can be automated to provide passphrases for operator cards, softcards, or administrator cards. In the following example, the passphrase variables are passed to the input pipeline, and the remaining key generation parameters are passed on the command-line:

```
PS C:\temp\test> $acs_passphrase, $softcard_passphrase | generatekey --batch pkcs11 protect=softcard
softcard=mysoftcard plainname=mykeyname nvram=yes
key generation parameters:
operation    Operation to perform            generate
application  Application                     pkcs11
module       Module to use                   1
protect      Protected by                    softcard
slot         Slot to read cards from         0
softcard     Soft card to protect key        mysoftcard
recovery     Key recovery                    yes
verify       Verify security of key          yes
type         Key type                        RSA
size         Key size                        2048
pubexp       Public exponent for RSA key (hex)
logkeyusage  Log key usage                   no
plainname    Key name                        mykeyname
nvram        Blob in NVRAM (needs ACS)       yes

Load Admin Card (for KNV):
Module 1 slot 0: Admin Card #1
Module 1 slot 0: Enter passphrase:
Module 1 slot 0:- passphrase supplied - reading card
Module #1 Slot #0: Processing ...
Card reading complete.

Please enter the passphrase for softcard 'mysoftcard':
Please wait........
Key successfully generated.
Path to key: C:\ProgramData\nCipher\Key Management Data\local\key_pkcs11_uce586891...
```

19.5. Secure strings

A passphrase or other sensitive string can be read into a variable in PowerShell using `$passphrase = Read-Host -AsSecureString`. `$passphrase` in this case is an instance of `System.Security.SecureString` and not the `System.String` type used for normal strings. The

contents of a `SecureString` cannot be read directly. If print the value of `$passphrase` in Pow
erShell, you only see the type name displayed and not the value that was entered to the
`Read-Host -AsSecureString` prompt.

nShield commands support using `SecureString` instances both for the input pipeline and as
parameters to the nShield command. This helps reduce the visibility of plaintext
passphrases or other sensitive values in scripts or in the shell. This is useful when using the
input pipeline to automate the presentation of passphrases to the prompts in card-loading
commands. It also means that nShield commands that take a passphrase as a command-
line parameter can be presented that string without the string becoming directly visible.

Example:

```
PS ppmk --new --newpp $passphrase newsoftcard
```

20. Preload Utility

20.1. Overview

The preload utility loads persistent cryptographic objects (keys/OCS/softcards) onto a chosen set of modules, then makes those objects available for use by applications. This removes the need for applications to load keys/cards themselves, and allows for easy sharing of keys/cards between multiple applications. Additionally, `preload` can manage keys, such that they are reloaded/maintained on modules to provide high availability.

Preloading is achieved via keys/cardsets being loaded, then once loaded the IDs of these objects are recorded persistently to a file (the preload file), which can be read via another application sharing the same session, and subsequently used.

Keys/cardsets must have previously been created before they can be preloaded, and all modules participating in a preload session must be in the same security world.

The preload binary can be found in `/opt/nfast/bin` (**Linux**) or `%NFAST_HOME%\bin` (**Windows**). This binary calls the `preload.py` script found in `/opt/nfast/python/scripts` (**Linux**) or `%NFAST_HOME%\python\scripts` (**Windows**).

The image below shows the relationship between preload, modules and applications:

20.2. Using Preload

20.2.1. Preload Commands

A command is needed in order to run preload. This command needs to be specified after the preload arguments.

The purpose of this command is to decide what needs to be done after preload has found and loaded all its crypto objects (OCS/softcards/keys).

```
> preload [arguments] command
```

Preload has a choice of 3 commands:

1. `pause` - continue to run the preload process forever. This is useful to load keys in one session and use them in another.

2. `exit` - exit preload gracefully. This is useful to add keys to the preload session. Not available in combination with high availability mode.

3. `subprocess` - execute this subprocess and exit once the subprocess has finished

> 🛈 | The only exception to this is the `--list-admin` option that does not require a command.

The preload session remains open, and thus the preloaded keys remain loaded, as long as at least one instance of preload continues to run. If/when the final preload instance terminates, all loaded objects will be cleaned up.

Example showing a single key, of type simple, being loaded and then an application being launched:

```
> preload -A simple -K key1 myapplication.py
```

20.2.2. Preload file location

The environment variable `NFAST_NFKM_TOKENSFILE` holds the path to the preload file. If it is not set, then the default location is used. A non-default location can also be set via the `--preload-file` option when invoking preload.

20.2.3. Preload Command Line Arguments

Argument	Effect
`--version`	show program's version number and exit

Argument	Effect
`-h|--help`	show help message and exit
`-m MODULE_NUMBER|--module=MODULE_NUM-BER`	Load on specified module (may be repeated; default = all).
`-c IDENT|--cardset=IDENT`	Load all cardsets matching `IDENT`. If `IDENT` looks like a hash it will be interpreted as that, otherwise it will be interpreted as a name. If it is definitely a name, use `--cardset-name`.
`--cardset-name=NAME`	Load cardset(s) named `NAME`.
`-s IDENT|--softcard=IDENT`	Load all softcards matching `IDENT`. If `IDENT` looks like a hash it will be interpreted as that, otherwise it will be interpreted as a name.
`--softcard-name=NAME`	Load softcard(s) named `NAME`.
`-o|--any-one`	Load a single cardset.
`-i|--interactive`	Load cardsets interactively until told to stop.
`-A APP|--appname=APP`	Choose the `appname` for subsequent `-K` options.
`-K IDENT|--key-ident=IDENT`	Load keys with ident matching `IDENT`.
`-n PATTERN|--name-pattern=PATTERN`	Load keys with name matching `PATTERN`. Use `*` for wildcard.
`--name-exact=NAME`	Load keys with name `NAME`.
`-M|--module-prot`	Load all module protected keys, in addition to any others requested.
`--no-cardset-keys`	Do not automatically load keys protected by requested cardsets. Deprecated.
`--no-token-keys`	Do not automatically load keys protected by requested tokens.
`--admin=KEYS`	Load admin keys (separate with commas, or use `all`).
`--list-admin`	List available admin key names (for `--admin`).
`-F|--require-fips`	Require FIPS-auth to be loaded.
`-N|--no-fips`	Do not record FIPS auth, even if available. (overrides `-F`).
`-H|--high-availability`	High availability mode.
`--polling-interval=POLLING_INTERVAL`	Interval (s) between polls for changes to the module list (default=60). High availability mode only.
`-f PRELOAD_FILE|--preload-file=PRE-LOAD_FILE`	Use specified preloaded objects file, instead of the default.
`-R|--reload-everything`	Reload keys and tokens that are already loaded.
`--show-key-info`	Display key information for keys as they are loaded.
`-l|--file-logging`	Log to file.

Argument	Effect
`-S|--no-stderr-logging`	Do not log to `stderr`, this is independent of file logging.
`--log-file=LOG_FILE`	The file destination for the log, defaults to `preload_%pid.log` in the `nfast` log directory.
`--log-level=LOG_LEVEL`	The log level to log, options: `DEBUG`, `INFO`, `WARNING`, `ERROR`. Default is `INFO`, if unrecognized option it will fall back to default.

20.2.4. Pattern Matching

Options to preload that use pattern matching, namely `--name-exact` and `--key-ident`, can accept the following wildcards:

Wildcard	Definition
`*`	matches everything
`?`	matches any single character
`[seq]`	matches any character in seq
`[!seq]`	matches any character not in seq

It is advised that all arguments that using wildcards are surrounded by quotations to ensure that they are passed to preload as intended. For example, to load all keys whose names start with `keyname`, the following pattern could be used:

```
> preload --name-pattern 'keyname*' exit
```

20.3. Preload File

The IDs of preloaded crypto objects are persistently stored in a preload file.

Each entry has the following format:

Element	Description
`Hash`	The `sha1 hash` of the crypto object.
`module`	The module which this object is present.
`objectid`	The id reference as a `M_KeyID`.
`generation`	This element is reserved for internal use.

Example `nfkminfo` output with preloaded crypto objects:

```
Pre-Loaded Objects ( 10): objecthash    module     objectid   generation
c29da3ac0d99a7c01477831ac31a4bebe283c4f8    1   0xac57be2e           1
c29da3ac0d99a7c01477831ac31a4bebe283c4f8    2   0xac57be2d           1
1080cca2be9588e6e47bcd870ebcbb133ea0561b    1   0xac57be2c           1
1080cca2be9588e6e47bcd870ebcbb133ea0561b    2   0xac57be13           1
```

By default the preload file location is `/tmp` (**Linux**) or the current user's temporary folder (**Windows**):

Linux

```
/tmp/nfpriv_<username>/default
```

Windows

```
<current user's temporary folder>\nfpriv_<username>\default
```

This location can be changed by using the command line option `-f PRELOAD_FILE|--preload-file=PRELOAD_FILE`.

20.4. Softcard Support

Softcards are now supported in preload, along with module protected keys and OCS cardsets.

In order to preload a softcard and the corresponding keys being protected by said softcard the `-s` or `--softcard-name` arguments can be used.

The `-s` option can be used with the softcard name or the hash of the softcard:

```
> nfkminfo -s
...
Operator logical token hash             name
3768b8efb7c7324dd8a1edbe2650c2015281c877  test
```

```
> nfkminfo -k simple aes128simplesoftcard1
...
name                "aes128simplesoftcard1"
hash                07c8110498dc0315455457f25564fc288c7da304
...
softcard            3768b8efb7c7324dd8a1edbe2650c2015281c877
```

```
> preload -s test nfkminfo
...
Pre-Loaded Objects ( 4):  objecthash    module objectid   generation
07c8110498dc0315455457f25564fc288c7da304    1 0xa411c0ab 1
```

```
07c8110498dc0315455457f25564fc288c7da304   2 0xa411c09e 1
3768b8efb7c7324dd8a1edbe2650c2015281c877   1 0xa411c09d 1
3768b8efb7c7324dd8a1edbe2650c2015281c877   2 0xa411c0a0 1
```

This shows the softcard is loaded on modules 1 and 2. It additionally shows that the key pro
tected by the softcard has been loaded on both modules.

20.4.1. No Cardset Keys

The `--no-cardset-keys` command line option can also be used for softcards.

This command line option will ensure that only the softcard is preloaded, and no keys pro-
tected by that cardset:

```
> preload -s test --no-cardset-keys
...
Pre-Loaded Objects ( 2): objecthash   module objectid   generation
3768b8efb7c7324dd8a1edbe2650c2015281c877     2 0xa9ba32a9 1
3768b8efb7c7324dd8a1edbe2650c2015281c877     1 0xa9ba32aa 1
```

20.5. FIPS Auth

FIPS Auth can be made available via preload.

The command line `-F` will ensure FIPS auth is preloaded everywhere.

The command line `-N` will ensure FIPS auth is not recorded, and will negate `-F`.

FIPS auth is also an admin key, see Admin Key section for more information.

20.6. Admin Keys

20.6.1. Listing

Admin keys can be listed using the `--list-admin` command line option.

This should be run without a command:

```
> preload --list-admin
```

Available admin keys are NSO, M, RA, P, NV, RTC, FIPS, MC, RE, DSEE, FTO.

20.6.2. Loading

Admin keys can be loaded using the `--admin=KEYS` command line option, supplying the value `--admin=ALL` to load all available admin keys. Note that admin key loading will require an ACS card being present in a slot of each module that is to be used.

Also note that the logical token of the admin key is preloaded alongside the key itself, for example, kfips and ltfips.

20.7. High Availability

Preload provides a high availability mode. When this mode is invoked Preload will load all requested keys, and will then periodically check for modules added or removed from the security world, or for keys becoming unloaded on existing modules. Should old or new modules be found to not have the specified keys/cardsets loaded, then preload will attempt to load them. This ensures that all available/usable modules have the requested keys loaded at all times, available for use by applications. Merged keyIDs are used to ensure applications can continually use these keys without interruption or changing key IDs. Preloaded keys are not only available to one application, but to any/all applications that share the preload session.

When preload is invoked with the `--high-availability` or `-H` option, it does the following differently:

1. Whenever preload loads a key onto the HSMs, it creates a Merged Key to represent the set of (HSM, key ID) pairs. Applications will then use these merged IDs to address the keys.

   ◦ As discussed below, this in itself provides failback, resilience and increased availability: the Merged Key ID remains usable even if some HSMs fail or are removed from the security world.

2. For as long as preload is running, it does the following repeatedly, once per polling interval:

   ◦ Consult the hardserver to get a list of operational HSMs which are in the relevant security world.

   ◦ For each Merged Key that was loaded by this instance of preload:

   ◦ Ensure there is a valid current entry for each usable HSM.

   ◦ To achieve this, check HSMs and load (or re-load) keys onto them as necessary, and update Merged Key contents.

   ◦ Ensure that the individual key IDs within each Merged Key are valid: Remove any that are no longer valid and usable (such as those for a removed HSM).

- ◦ Update the preload file to reflect changes, if any.
- ◦ When finished, sleep for an interval of time, then repeat.

In summary, this mode attempts to keep preloaded crypto objects present on all usable modules in a security world (or a set of modules if requested via the -m argument) for as long as preload is running, with a keyID that remains constant, so that keys are available for use by any applications sharing the preload session.

20.7.1. Prerequisites for high availability mode

Users should not mix and match instances of preload with and without the `-H` high availability option, if those instances are sharing a session.

Managing OCS cardset-protected keys requires the following:

- the OCS protecting the key(s) be a 1/N quorum
- the passphrase for each card of the OCS set be identical
- one card of the OCS set be left inserted in a slot (local or remote) for each module
- if the card is non-persistent, it must be left in a local slot.

20.7.2. Differences from legacy behaviour

When running in high availability mode, certain behaviours of may differ from those outside of high availability mode. This includes the prompts for PIN entry, error messages, etc. This is due to a necessary difference in implementation between the two modes, and is expected.

20.7.3. Conditions for Management/Reloading

As mentioned above, preload in high availability mode will (re)load keys onto modules when a module is usable. A module will be considered usable if that module is in operational mode and in the correct world (and in the case of OCS protected keys, if a card from the OCS set is inserted into the module, locally or remotely). Preload will not attempt to perform actions that involve world administration, such as world loading or client enrolment. Users are responsible for managing worlds and client enrolment, and thus for bringing modules into a usable state.

> The automatic loading/reloading of keys onto usable modules is not to be confused with forced reloading of keys provided by the -R option.

20.7.4. Merged Keys in the Preload File

When high availability mode is activated, all keys are represented in the preload file as Merged keys; cardsets and softcards are represented in the same way as non-high-availability mode.

Due to the fact that in high availability mode keys are represented as MergedKeys, which do not correspond to any one particular module, the module element of the preload file is no longer relevant for keys. However, for cardsets, the module field is still utilized.

For symmetric and private halves of asymmetric keys the module number is represented as a -1 and for public halves of asymmetric keys the module number is represented as a -2.

This is evident in the output from nfkminfo. (Note that nfkminfo ignores the 32-bit two's-complement representation, thus displaying -1 and -2 as $(2^{32} - 1)$ and $(2^{32} - 2)$ respectively: 4294967295 and 4294967294):

```
Pre-Loaded Objects ( 4): objecthash    module objectid  generation
84749a62d0f71db7f80c5df6469c11685f7f1b78    1 0xb5c0c7fa 1
84749a62d0f71db7f80c5df6469c11685f7f1b78    2 0xb5c0c7fd 1
28dcee51dfc53387f4dc4d55538d8b5253ee85d1 4294967295 0xb5c0c7f7 1
c2afe833ae6e823a37777c633a5b3a18a9e5dfbd 4294967294 0xb5c0c7f8 1
```

As shown above, cardsets/softcards are still module specific.

To make nfkminfo show the preloaded objects, run it as a subprocess as part of the preload command. See the section above on using preload.

> **ℹ** Merged Key IDs (just like single-key IDs) are shared between multiple instances of preload that are invoked by the same client (i.e. using the same ClientID). As such, applications must ensure that they perform no operations that delete or replace the merged key ID, or alter the keys that are part of that merged key ID.

20.7.5. Polling Interval

Preload manages its crypto objects by polling available modules, based on a polling interval.

Once per interval, if preload detects modules (new or existing) without the relevant crypto objects (keys/cards) present, it will attempt to load those missing objects.

This polling interval is configurable via the command line option --polling-interval=SEC-ONDS

By default the polling interval is 60 seconds.

20.7.6. Key timeouts and use limits

It is advised to not use OCSs or keys with timeouts in high availability mode, as preload will be unable to reload objects once their timeouts have expired.

In high availability mode, there are situations where OCS/keys that have previously timed out, or reached maximum use limits, may be reloaded (and thus their limits reset) without user interaction. In general within high availability mode keys that have timed out or reached their use limits will be left in place, unusable, respecting the limits. However if the module containing those keys reboots or resets then, upon the module's return, preload will notice that the keys are not loaded and will load them. This reloading of keys will necessarily reset timeouts and use limits. If the timeout on an OCS has reached its limit, any keys protected by that OCS will not be reloaded on newly-indoctrinated modules in the security world.

20.7.7. Multiple Preload instances in high availability mode

As described above, keys will be maintained by the preload instance that first introduces them, and will cease being maintained when that instance ends. (Here maintained means reloaded automatically onto relevant HSMs that lack them.)

Therefore when `preload` is invoked with `exit` (or a short-lived `subprocess` command) it will load the specified keys but then exit, leaving those keys unmaintained.

If a `preload` process is already running under high availability mode, any new preload process (with the same preload file) will gain access to the preloaded keys. As such that later instance must also be run in high availability mode (and preload will reject an attempt to run it in plain mode in this situation).

The `pause` command may be useful for setting up availability of keys for subsequent use by multiple applications:

First, a long-running preload instance to load keys and maintain them indefinitely:

```
$ preload --high-availability [...other options...] pause
```

Then run applications (possibly short-lived) that use those keys:

```
$ preload --high-availability [...other options...] app --args --for --app
```

20.7.7.1. Managing Keys

Given multiple preload processes run under high availability, the process that will manage the keys is the first process to find them, based on command line options.

For example, Security World crypto objects:

crypto object	name	protected by
Softcard	softcard1	N/a
Key	simple_softcard1	softcard1
Key	simple_module1	module

First preload process started:

```
> preload -H -s softcard1 pause
```

This would load the softcard `softcard1` on all modules as well as the key `simple_softcard1`:

```
> preload -H nfkminfo
...
Pre-Loaded Objects (  3): objecthash   module objectid  generation
29235f2a0b77fc1e18641b0820fe3c93e030a02e 4294967295 0x44313d41 1
5bccb6f540802ef1da3828f6b8b0f3fc985272e6   2 0x44313d47 1
5bccb6f540802ef1da3828f6b8b0f3fc985272e6   1 0x44313d46 1
...
> nfkminfo -k simple simplesoftcard1
...
name               "simple_softcard1"
hash               29235f2a0b77fc1e18641b0820fe3c93e030a02e
...
> nfkminfo -s
...
Operator logical token hash          name
5bccb6f540802ef1da3828f6b8b0f3fc985272e6  softcard1
```

Second preload process started:

```
> preload -H -n simple pause
```

This would load the key `simple_module` on all modules:

```
> preload -H nfkminfo
...
Pre-Loaded Objects (  4): objecthash   module objectid  generation
600bcc26336c13f2371bdbb54b1cde293ded9a15 4294967295 0x44313d29 1
29235f2a0b77fc1e18641b0820fe3c93e030a02e 4294967295 0x44313d41 1
5bccb6f540802ef1da3828f6b8b0f3fc985272e6   2 0x44313d47 1
5bccb6f540802ef1da3828f6b8b0f3fc985272e6   1 0x44313d46 1
...
> nfkminfo -k simple simplemodule1
...
name               "simple_module1"
```

```
hash                    600bcc26336c13f2371bdbb54b1cde293ded9a15
```

The evidence that the first preload process is still managing the key `simple_softcard1`, even though the second preload process could have loaded it, is in the `objectid`.

The object id for key `simple_softcard1` has not changed (**0x44313d41**).

20.7.8. FIPS Auth in High Availability mode

Fips auth can be preloaded when running preload in high availability mode. In this scenario fips auth will be loaded as a high availability key (ie, reloaded/maintained on modules, as with other preloaded keys).

To enable FIPS auth use the command line option `-F`.

However, note that fips auth is represented differently, in comparison to other high availability mode keys, within the preload file.

The FIPS auth key is represented in the preload file multiple times: once for each module it is loaded on, and one extra time with a negative module ID as with other merged IDs. However the `objectid` is still a Mergedkey so will remain the same across those entries. This duplication of entries is to maintain compatibility with legacy behaviour/applications.

The following shows the pre-loaded FIPS auth objects on an estate of 3 modules - note there are 4 entries, each with the same `objectid`:

```
Pre-Loaded Objects ( 4): objecthash   module objectid  generation
aa462d0dd9dfeaa80968aadda2610ac0f6f94352   3 0xa824b9ab 1
aa462d0dd9dfeaa80968aadda2610ac0f6f94352   2 0xa824b9ab 1
aa462d0dd9dfeaa80968aadda2610ac0f6f94352   1 0xa824b9ab 1
aa462d0dd9dfeaa80968aadda2610ac0f6f94352 4294967295 0xa824b9ab 1
...
hkfips      aa462d0dd9dfeaa80968aadda2610ac0f6f94352
```

20.7.9. PKCS #11 and JCE

Both PKCS #11 and JCE applications are compatible with the high availability mode of preload, provided the PKCS #11 or JCE library that the application uses is from the 12.60 release or later. Flags or environment variables only need to be set to enable this when PKCS #11 is used for key reloading.

20.7.9.1. Use PKCS #11 for key reloading

PKCS #11 key reloading requires preload to be run in high availability mode, with the follow-

ing options enabled:

- `--high-availability`.
- `--no-token-keys`.
- `--preload-file=PRELOAD_FILE`, where PRELOAD_FILE must match the location given to PKCS #11 with the NFAST_NFKM_TOKENSFILE environment variable.
- Either `--cardset=<IDENT>` or `--softcard=<IDENT>` (depending on whether using card set or softcard protected keys), where `<IDENT>` is the identifier of the card set or softcard, respectively.

> ℹ PKCS #11 key reloading is also supported for module-protected keys, but the PKCS #11 application must still be run under a preload application that is reloading tokens for another key.

> ℹ Using preload in high availability mode with Operator Card Sets has a set of restrictions, see Overview.

- Additionally, the following option is not required, but recommended:

  `--polling-interval=<POLLING_INTERVAL>`, where `<POLLING_INTERVAL>` also determines how often PKCS #11 will attempt to reload keys. The default is 60 seconds.

For more information, see section *PKCS _11 with key reloading* in the *Cryptographic API Integration Guide*.

20.7.10. Unsupported options

The `-H --high-availability` option may not be used in conjunction with any of the following options:

- `-o --any-one`
- `-i --interactive`
- `exit`
- `--admin`
- `--reload-everything`

20.8. Logging

By default `preload` logs to `stderr`.

Logs follow the format: `yyyy-mm-dd hh:mm:ss: [pid]: LogLevel: message`

For example:

```
2019-03-27 09:45:50: [439]: INFO: loading objects
```

Preload can also log to a file, this behaviour is separate from `stderr` logging. Therefore we can disable logging or log to `stderr` and/or a file.

To disable `stderr` logging, use the command line option `-S`. To enable file logging use the command line option `-l`.

The default file location for logs is `/opt/nfast/logs/preload_log_pid.log` (**Linux**) or `%NFAST_HOME%\logs\preload_log_pid.log` (**Windows**).

To change the file location, use the command line option `--log-file=FILE`.

As standard, preload has different log levels. These are:

- `DEBUG`
- `INFO`
- `WARNING`
- `ERROR`
- `CRITICAL`

The log level is by default: `INFO` and can be changed via the command line option `--log-level=LEVEL`.

20.9. Using preloaded objects - Worked example

In order to use preloaded objects, an application needs to create a connection that reads in the preload file:

Python:

```python
import nfkm

conn = nfkm.connection(existingobjects="") # Reads file from default location

# If no existingobjects parameter is specified,
# the connection will not attempt to read any preload file:
conn_no_preload = nfkm.connection()
```

If the `existingobjects` argument is the empty string, the connection will use the file from the default location.

Any other string should be a path to different preload file. It can then call `NFKM_GetInfo` to

---

get the security world info:

Python:

```
world_info = nfkm.getinfo(conn)
```

This results in a data structure with all the preloaded objects (this list is static and created at the time of connection creation):

Python:

```
import nfkm

conn = nfkm.connection(existingobjects="")

world_info = nfkm.getinfo(conn)

print world_info.existingobjects
```

Result:

```
[
ExistingObjectInfo
.module= 2
.hash= 84749a62 d0f71db7 f80c5df6 469c1168 5f7f1b78
.change= 1
.id= 0xffffffff88afd208,
ExistingObjectInfo
.module= 1
.hash= 84749a62 d0f71db7 f80c5df6 469c1168 5f7f1b78
.change= 1
.id= 0xffffffff88afd20b
]
```

Once an application has the M_KeyID references, it can use those cryptographic objects:

```
objid = world_info.existingobjects[0].id

cmd = nfkm.Command(["GetLogicalTokenInfo", 0, objid])

print conn.transact(cmd)
```

Result:

```
Reply.cmd= GetLogicalTokenInfo
.status= OK
.flags= 0x0
.reply.state= Present
.hkt= 84749a62 d0f71db7 f80c5df6 469c1168 5f7f1b78
.shares= empty
.sharesneeded= 0
```

21. Environment variables

This appendix describes the environmental variables used by Security World Software.

> ℹ️ When you are using these environment variables on Windows to configure nShield services such as the hardserver (nFast Server service), these must be set as System variables only; not as User Variables. Any service for which the environment variable changes are intended must be restarted for the change to take effect.

Variable	Description	Win	Lnx
KERNEL_HEADERS	This variable allows you to specify the path to kernel headers (if, for example, they are not in the default directory). It is necessary for the configuration script to be able to find the kernel headers when building the PCI driver during software installation.	n	y
NFAST_CERTDIR	This variable specifies the path to the dynamic feature enabling Feature Certificates directory. You only need to change the value of this variable if you move the Installation directory. See NFAST_HOME, NFAST_KMDATA, and NFAST_LOGDIR.	y	y
NFAST_DEBUG	This variable enables debug logging for the PKCS #11 library. You must set NFAST_DEBUG equal to a value in the range 1 – 7 for debug messages to be logged (see Logging, debugging, and diagnostics. For more information, see also Logging and debugging information for PKCS #11.	y	y
NFAST_DEBUGSYSLOG	This variable redirects debug logging to syslog. The value of the environment variable should be one of the syslog facilities to be used. Prefixing the facility name with + enables traditional logging and syslog simultaneously.	y	Y
NFAST_HOME	This variable specifies the path to the Installation directory, which is set by the Security World Software installation script. You only need to change the value of this variable if you move the Installation directory. See NFAST_KMDATA, NFAST_CERTDIR, and NFAST_LOGDIR.	y	y
NFAST_KMDATA	This variable sets the location of the Key Management Data directory. You only need to change the value of this variable if you move the Key Management Data directory. See NFAST_HOME, NFAST_CERTDIR, NFAST_LOGDIR, and NFAST_KMLOCAL.	y	y

Variable	Description	Win	Lnx
`NFAST_KMLOCAL`	This variable specifies the location of the Key Management and Security World Data directory. If this environment variable is not set, by default the module looks for the Security World data in the `local` subdirectory of the Key Management Data directory. See `NFAST_KM-DATA`.	y	y
`NFAST_LOGDIR`	This variable specifies the location of the Log Files directory. You only need to change the value of this variable if you move the Log Files directory. See `NFAST_HOME`, `NFAST_KMDATA`, and `NFAST_CERTDIR`.	y	y
`NFAST_USER_LOGDIR`	This variable specifies the location of log files that are specific to each user. In Security World versions before 12.60.3, the default is the user's home directory (Linux) or user profile folder (Windows). From 12.60.3, the default is the subdirectory `nshieldlogs` in the home directory or user profile folder.	y	y
`NFAST_NFKM_TOKENSFILE NFAST_NFKM_TOKENSSELECT`	This variable sets the default values for a file in which `ClientID` and `KeyIDs` are stored by the `preload` command-line utility.	y	y
`NFAST_SEE_MACHINEENCKEY_DEFAULT`	This variable is the name of the `SEEConf` key needed to decrypt SEE-machine images. Running the command `loadmache --encryptionkey='IDENT (or 'loadmache --unencrypted`) overrides any value set by this variable.	y	y
`NFAST_SEE_MACHINEENCKEY_<module>`	This variable is the name of the `SEEConf` key needed to decrypt the SEE-machine image targeted for the specified module. It overrides `NFAST_SEE_MACHINEENCKEY_DEFAULT` for the specified module. Running the command `loadmache --encryptionkey=<IDENT>` (or `loadmache --unencrypted`) overrides any value set by this variable.	y	y
`NFAST_SEE_MACHINEIMAGE_DEFAULT`	This variable is the path of the SEE machine image to load on to any module for which a specific image is not defined. Supplying the `machine-filename` parameter when running the `loadmache` command-line utility overrides this variable. This variable is not affected when running the `loadsee-setup` or `hsc_loadseemachine` utilities.	y	y

Variable	Description	Win	Lnx
`NFAST_SEE_MACHINEIMAGE_<mod-ule>`	This variable is the path of the SEE machine image to load on to the specified module. If set, this variable overrides the use of `NFAST_SEE_MACHINEIMAGE_DEFAULT` for the specified module. Supplying the `<machine-file name>` parameter when running the `loadmache` command-line utility overrides the `NFAST_SEE_MACHINEIM-AGE_<module>` variable. This variable is not affected when running the `loadsee-setup` or `hsc_loadseema-chine` utilities.	y	y
`NFAST_SEE_MACHINESIGHASH_DE-FAULT`	This variable is the default key hash of the vendor sign-ing key (`seeinteg`) that signs SEE machine images. This variable is only required if you are using a dynamic SEE feature with an encrypted SEE machine. Running the command `loadmache --sighash=<HASH>` any value set in this variable.	y	y
`NFAST_SEE_MACHINE-SIGHASH_<module>`	This variable is the key hash of the vendor signing key (`seeinteg`) that signs SEE machine images for the spec-ified module. It overrides `NFAST_SEE_MACHINE-SIGHASH_DEFAULT` for the specified module. This vari-able is only required if you are using a dynamic SEE fea-ture with an encrypted SEE machine. Running the com-mand `loadmache --sighash=<HASH>` any value set in this variable.	y	y

`NFAST_SERVER`
 	
`NFAST_PRIVSERVER`
	If these variables are set in the hardserver's environ-ment, the values specify:
	
	On Linux, the pathnames of the UNIX domain sockets that the hardserver uses for ordinary/privileged client connections to the hardserver.
	
	On Windows, the names of the Windows named pipes for ordinary/privileged client connections to the hard-server.
	
	These variables are available for this purpose for back-ward compatibility only; you should configure sockets in the hardserver configuration file, see server_startup. If you set these variables they override the values in hardserver configuration file.
		y	y

Variable	Description	Win	Lnx

`NFAST_SERVER_PORT`
 	
`NFAST_SERVER_PRIVPORT`
	If these variables are set in the hardserver's environment, the values specify the TCP port numbers that the nFast server uses for connections over TCP sockets.
	
	These variables are available for this purpose for backward compatibility only: you should configure ports in the hardserver configuration file, as described in server_startup. If you set these variables, they override the values in the hardserver configuration file.
		y	y

`NFLOG_CATEGORIES`	This variable is used to filter log messages by supplying a colon-separated list of allowable message categories; see Logging, debugging, and diagnostics. If no value is supplied, all message categories are logged.	y	y
`NFLOG_SEVERITY`	This variable is used to filter log messages by supplying a minimum severity level to be logged; see Logging, debugging, and diagnostics. If no value is supplied, the default severity level is `WARNING`.	y	y
`NFLOG_DETAIL`	This variable is used to filter log messages by supplying a bitmask of detail flags; see Logging, debugging, and diagnostics. The default is `time+severity+writeable`.	y	y
`NFLOG_FILE`	This variable is used to specify a filename (or file descriptor) in which log messages are to be written; see Logging, debugging, and diagnostics. The default is `stderr` (the equivalent of file descriptor `&2`).	y	y

22. Logging, debugging, and diagnostics

This appendix describes the settings and tools you can use to access the logging and debugging information generated by the Security World Software. You are also shown how to obtain system information using the `nfdiag` command-line utility.

22.1. Logging and debugging

The nShield HSM and the applications that use it generate log files. You can view log files using the front panel. Application log messages are handled on the client.

The current release of Security World Software uses controls for logging and debugging that differ from those used in previous releases. However, settings you made in previous releases to control logging and debugging are still generally supported in the current release, although in some situations the output is now formatted differently.

Some text editors, such as Notepad, can cause `NFLOG` to stop working if the `NFLOG` file is open at the same time as the hardserver is writing the logs.

> Debug logs may include sensitive data.
>
> | | |
> |---|---|
> | `CKNFAST_DEBUG` | All severities from `DL_Call` to `DL_DetailMutex` may result in sensitive data being logged |
> | `JCECSP_DEBUG` | No further guidance |
> | `NFJAVA_DEBUG` | No further guidance |
> | `NFLOG_DETAIL` | The `0x00010000` flag may result in sensitive data being logged |
> | `NFLOG_SEVERITY` | All the `DEBUGn` settings may result in sensitive data being logged |

22.1.1. Environment variables to control logging

The Security World for nShield generates logging information that is configured through a set of four environment variables:

`NFLOG_FILE`

This environment variable specifies the name of a file (or a file descriptor, if prefixed with the `&` character) to which logging information is written. The default is `stderr` (the equivalent of `&2`).

Ensure that you have permissions to write to the file specified by `NFLOG_FILE`.

NFLOG_SEVERITY

This environment variable specifies a minimum severity level for logging messages to be written (all log messages less severe than the specified level are ignored). The level can be one of (in order of greatest to least severity):

1. `FATAL`

2. `SEVERE`

3. `ERROR`

4. `WARNING`

5. `NOTIFICATION`

6. `DEBUG`N, where N can be an integer from 1 to 10 inclusive that specifies increasing levels of debugging detail, with 10 representing the greatest level of detail, although the type of output is depends on the application being debugged.

> ℹ️ The increasingly detailed information provided by different levels of `DEBUG`N is only likely to be useful during debugging, and we recommend not setting the severity level to `DEBUG`N unless you are directed to do so by Support.

The default severity level is `WARNING`.

NFLOG_DETAIL

This environment variable takes a hexadecimal value from a bitmask of detail flags as described in the following table (the `logdetail` flags are also used in the hardserver configuration file to control hardserver logging; see: server_settings):

Hexadecimal flag	Function	logdetail flags
0x00000001	This flag shows the external time (that is, the time according to your machine's local clock) with the log entry. It is on by default.	external_time
0x00000002	This flag shows the external date (that is, the date according to your machine's local clock) with the log entry.	external_date
0x00000004	This flag shows the external process ID with the log entry.	external_pid
0x00000008	This flag shows the external thread ID with the log entry.	external_tid

Hexadecimal flag	Function	logdetail flags
0x00000010	This flag shows the external `time_t` (that is, the time in machine clock ticks rather than local time) with the log entry.	external_time_t
0x00000020	This flag shows the stack backtrace with the log entry.	stack_backtrace
0x00000040	This flag shows the stack file with the log entry.	stack_file
0x00000080	This flag shows the stack line number with the log entry.	stack_line
0x00000100	This flag shows the message severity (a severity level as used by the `NFLOG_SEVERITY` environment variable) with the log entry. It is on by default.	msg_severity
0x00000200	This flag shows the message category (a category as used by the `NFLOG_CATEGORIES` environment variable) with the log entry.	msg_categories
0x00000400	This flag shows message writeables, extra information that can be written to the log entry, if any such exist. It is on by default.	msg_writeable
0x00000800	This flag shows the message file in the original library. This flag is likely to be most useful in conjunction with Security World Software-supplied example code that has been written to take advantage of this flag.	msg_file
0x00001000	This flag shows the message line number in the original library. This flag is likely to be most useful in conjunction with example code we have supplied that has been written to take advantage of this flag.	msg_line
0x00002000	This flag shows the date and time in UTC (Coordinated Universal Time) instead of local time.	options_utc
0x00004000	This flag shows the full path to the file that issued the log messages.	options_fullpath
0x00008000	This flag includes the number of microseconds in the timestamp.	options_time_us
0x00010000	This flag enables logging of potentially secret values in generic stub log output.	msg_secrets

NFLOG_CATEGORIES

This environment variable takes a colon-separated list of categories on which to filter log messages (categories may contain the wild-card characters `*` and `?`). If you do not supply any values, then all categories of messages are logged. This table lists the available categories:

Category	Description
nflog	Logs all general messages relating to nflog.
nflog-stack	Logs messages from `StackPush` and `StackPop` functions.
memory-host	Logs messages concerning host memory.
memory-module	Logs messages concerning module memory.
gs-stub	Logs general generic stub messages. (Setting this category works like using the `dbg_stub` flag with the logging functionality found in previous Security World Software releases.)
gs-stubbignum	Logs bignum printing messages. (Setting this category works like using the `dbg_stubbignum` flag with the logging functionality found in previous Security World Software releases.)
gs-stubinit	Logs generic stub initialization routines. (Setting this category works like using the `dbg_stubinit` flag with the logging functionality found in previous Security World Software releases.)
gs-dumpenv	Logs environment variable dumps. (Setting this category works like using the `dbg_dumpenv` flag with the logging functionality found in previous Security World Software releases.)
nfkm-getinfo	Logs `nfkm-getinfo` messages.
nfkm-newworld	Logs messages about world generation.
nfkm-admin	Logs operations using the Administrator Card Set.
nfkm-kmdata	Logs file operations in the `kmdata` (**Linux**) or `%NFAST_KMDATA%` (**Windows**) directory.
nfkm-general	Logs general NFKM library messages.
nfkm-keys	Logs key loading operations.
nfkm-preload	Logs `preload` operations.
nfkm-ppmk	Logs softcard operations.
serv-general	Logs general messages about the local hardserver.
serv-client	Logs messages relating to clients or remote hardservers.
serv-internal	Logs severe or fatal internal errors.

Category	Description
serv-startup	Logs fatal startup errors.
servdbg-stub	Logs all generic stub debugging messages.
servdbg-env	Logs generic stub environment variable messages.
servdbg-underlay	Logs messages from the OS-specific device driver interface
servdbg-statemach	Logs information about the server's internal state machine.
servdbg-perf	Logs messages about the server's internal queuing.
servdbg-client	Logs external messages generated by the client.
servdbg-messages	Logs server command dumps.
servdbg-sys	Logs OS-specific messages.
pkcs11-sam	Logs all security assurance messages from the PKCS #11 library.
pkcs11	Logs all other messages from the PKCS #11 library.
rqcard-core	Logs all card-loading library operations that involve standard message passing (including slot polling).
rqcard-ui	Logs all card-loading library messages from the current user interface.
rqcard-logic	Logs all card-loading library messages from specific logics.

You can set a minimum level of hardserver logging by supplying one of the values for the `NFLOG_SEVERITY` environment variable in the hardserver configuration file, and you can likewise specify one or more values for the `NFLOG_CATEGORIES` environment variable. For detailed information about the hardserver configuration file settings that control logging, see server_settings.

> ℹ️ If none of the four environment variables are set, the default behavior is to log nothing, unless this is overridden by any individual library. If any of the four variables are set, all unset variables are given default values.

22.1.2. Logging from the nShield CSP and CNG (Windows)

By default, logging is disabled for the nShield CSP and CNG.

To enable logging, use a suitable registry editor such as `regedit`.

Depending on whether the program you wish to debug is 64-bit or 32-bit based, you will have to create respective registry keys if they do not already exist.

For a 64-bit program on a 64-bit OS, create the following registry key if it does not already exist:

`HKEY_LOCAL_MACHINE\SOFTWARE\nCipher\Cryptography`

For a 32-bit program on a 64-bit OS, create the following registry key if it does not already exist:

`HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\nCipher\Cryptography`

Open the registry at the required `Cryptography` key as described above, and under the key create the following variables.

1. Create a new `string` variable named `PathName`.
2. Open the `PathName` variable and provide a value which is a suitable path to where you want the log file(s) to be placed (for example, `C:\Users\MyName\Documents`.) Do not give a log file name. The log file name(s) will be created automatically when logging starts.

   > ℹ It must be possible for the provider to write to the specified path.

3. Create a new `DWORD (32 bit)` variable named `LogLevel`.
4. Open the `LogLevel` variable and provide a suitable value (for example, `2`).

   Permitted values for `LogLevel` are:

Value	Output
0	Messages are sent to the event log.
1	Error messages are sent to the log file.
2	Function calls and error messages are sent to the log file.
3	All information, including debugging information, is sent to the log file.

> ℹ Do not set this value to 3 unless either you are asked to do so by Support or you are debugging your own code. At this level of logging, a single SSL connection generates approximately 30 kilobytes of log messages. In addition, sensitive information may appear in the log file.

> ℹ If `LogLevel` is not set, then the provider only logs messages of warning severity or greater (equivalent to setting `NFLOG_SEVERITY=warning`).

If neither `PathName` nor `LogLevel` are set for the CSP or CNG, logging remains disabled.

If logging is successfully enabled, the log file(s) should appear at the location specified in `PathName`, and will have names similar to:

- `nfdebug.txt`
- `ncspdddebug.txt`
- `nckspswdebug.txt`

22.1.3. Logging and debugging information for PKCS #11

In order to get PKCS #11 logging and debugging output, you must set the `CKNFAST_DEBUG` variable. A debug output file (with path) can also be set using the `CKNFAST_DEBUGFILE` variable. These variables can be set in the environment or the `/opt/nfast/cknfastrc` (**Linux**) or `%NFAST_HOME%\cknfastrc` (**Windows**) file. Normally settings in the environment should override the same settings (if any) in the `cknfastrc` file. You can specify any appropriate PKCS #11 categories using the `NFLOG_CATEGORIES` environment variable.

To produce PKCS #11 debug output, the `CKNFAST_DEBUG` variable can be a given value from 1 through to 11, where the greater the value the more detailed debug information is provided. A value of 7 is a reasonable compromise between too little and too much debug information. A value of 0 switches the debug output off.

This environment variable takes a colon-separated list of categories on which to filter log messages (categories may contain the wildcards characters `*` and `?`).

The following table maps PKCS #11 debug level numbers to the corresponding `NFLOG_SEVERITY` value:

PKCS #11 debug level	PKCS #11 debug meaning	NFLOG_SEVERITY value	Output in log
0	DL_None	NONE	
1	DL_EFatal	FATAL	`"Fatal error:"`
2	DL_EError	ERROR	`"Error:"`
3	DL_Fixup	WARNING	`"Fixup:"`
4	DL_Warning	WARNING	`"Warning:"`
5	DL_EApplic	ERROR	`"Application error:"`
6	DL_Assumption	NOTIFICATION	`"Unsafe assumption:"`
7	DL_Call	DEBUG2	`">> "`
8	DL_Result	DEBUG3	`"< "`
9	DL_Arg	DEBUG4	`"> "`

PKCS #11 debug level	PKCS #11 debug meaning	NFLOG_SEVERITY value	Output in log
10	DL_Detail	DEBUG5	`"D "`
11	DL_DetailMutex	DEBUG6	`"DM "`

22.1.4. Debugging information for Java

This section describes how you can specify the debugging information generated by Java.

22.1.4.1. Setting the Java debugging information level

In order to make the Java generic stub output debugging information, set the Java property `NFJAVA_DEBUG`. The debugging information for `NFJAVA`, `NFAST`, and other libraries (for example, `KMJAVA`) can all use the same log file and have their entries interleaved.

You set the debugging level as a decimal number. To determine this number:

1. Select the debugging information that you want from the following list:

   ```
   NONE = 0x00000000 (debugging off)
   MESS_NOTIFICATIONS = 0x00000001 (occasional messages including important errors)
   MESS_VERBOSE = 0x00000002 (all messages)
   MESS_RESOURCES = 0x00000004 (resource allocations)
   FUNC_TRACE = 0x00000008 (function calls)
   FUNC_VERBOSE = 0x00000010 (function calls + arguments)
   REPORT_CONTEXT = 0x00000020 (calling context e.g ThreadID and time)
   FUNC_TIMINGS = 0x00000040 (function timings)
   NFJAVA_DEBUGGING = 0x00000080 (Output NFJAVA debugging info)
   ```

2. Add together the hexadecimal value associated with each type of debugging information.

   For example, to set `NFJAVA_DEBUGGING` and `MESS_NOTIFICATIONS`, add 0x00000080 and 0x00000001 to make 0x00000081.

3. Convert the total to a decimal and specify this as the value for the variable.

   For example, to set `NFJAVA_DEBUGGING` and `MESS_NOTIFICATIONS`, include the line:

   ```
   NFJAVA_DEBUG=129
   ```

   For `NFJAVA` to produce output, `NFJAVA_DEBUG` must be set to at least `NFJAVA_DEBUGGING` + `MESS_NOTIFICATIONS`. Other typical values are:

   ◦ 255: All output

- ◦ `130`: nfjava debugging and all messages (`NFJAVA_DEBUGGING` and `MESS_VERBOSE`)
- ◦ `20`: function calls and arguments and resource allocations (`FUNC_VERBOSE` and `MESS_RESOURCES`)

22.1.4.2. Setting Java debugging with the command line

You can set the Java debug options by immediately preceding them with a `-D`. Use the `NJAVA_DEBUGFILE` property to direct output to a given file name, for example:

```
java -DNFJAVA_DEBUGFILE=myfile -DNFJAVA_DEBUG=129 -classpath .... classname
```

> ℹ️ Do not set `NFJAVA_DEBUG` or `NFJAVA_DEBUGFILE` in the environment because Java does not pick up variables from the environment.

If `NFJAVA_DEBUGFILE` is not set, the standard error stream `System.err` is used.

> ℹ️ Set these variables only when developing code or at the request of Support.

> ❗ Debug output contains all commands and replies sent to the hardserver in their entirety, including all plain texts and the corresponding cipher texts as applicable.

22.2. Diagnostics and system information

> ℹ️ Besides the diagnostic tools described in this section, we also supply a performance tool that you can use to test Web server performance both with and without an nShield HSM. This tool is supplied separately. If you require a copy, contact your Sales representative.

22.2.1. nfdiag: diagnostics utility

The `nfdiag` command-line utility is a diagnostics tool that gathers information about the system on which it is executed. It can save this information to either a `.zip` file or a text file.

Under normal operating conditions, you do not need to run `nfdiag`. You can run `nfdiag` before contacting Support and include its output file with any problem report.

22.2.1.1. Usage

> **ℹ** Run `nfdiag` with the standard `-h|--help` option to display information about the options and parameters that control the program's behavior.

If you want to supply additional diagnostic files, run:

```
nfdiag -e|--extrainfo <FILENAME>
```

You can only attach plaintext files.

The `nfdiag` command-line utility is an interactive tool. When you run it, it prompts you to supply the following information:

Option	Actions to take
`which application(s) you are using`	Identify all application software installed on the machine on which any problem with the nShield product occurs.
`what APIs you are using`	Describe any custom software, especially any interaction it has with the nShield security system.
`a description of the problem`	Include as much detail as possible, including any error messages you have seen.
`a Support ticket number` (if you have one)	When you contact Support you are supplied with a Support ticket number. Enter this number to help Support expedite the collection of any information you have sent.
`a contact email address`	Supply an email address that has as few e-mail/spam filters as possible so that any additional files that Support sends to you are not blocked. We use the e-mail address you supply here *only* for communication directly related to your problem report.
`a contact name`	Enter your name (or the name of an appropriate person for contact by Support).
`a contact telephone number`	Include the appropriate country and any region code for your contact telephone number.

> **ℹ** Except for a Support ticket number, nfdiag requires non-NULL answers to all its prompts for information.

Supplying this information helps `nfdiag` capture as much relevant information as possible for any problem report to Support. As you supply information at each prompt in turn, press **Enter** to confirm the information and continue to the next prompt. Information you supply cannot extend over multiple lines, but if you need to supply this level of information, you can include it in additional attached files, as described.

By default, `nfdiag` runs in verbose mode, providing feedback on each command that it exe-

cutes and which log files are available. If the system is unable to execute a command, the verbose output from `nfdiag` shows where commands are stalling or waiting to time out.

At any time while `nfdiag` is running, you can type **Ctrl-C** to cancel its current commands and re-run it.

22.2.1.2. Output

After you have finished supplying information for each required prompt, `nfdiag` generates a plain text output file and displays its file name.

If the file `opt/nfast/log/logfile` (**Linux**) or `%NFAST_HOME%\log\logfile` (**Windows**) exists, `nfdiag` automatically includes this file in its output. If this file does not exist, `nfdiag` warns you that it could not process this file. This warning does not affect the validity of the generated output file.

When complete, this output file contains the following:

- The information supplied interactively to nfdiag when run
- Details about the client machine
- Details about any environment variables
- Output from the following command-line utilities:
    - `enquiry`
    - `stattree`
    - `ncversions`
    - `nfkminfo`
- The contents of the following log files (if they are available):
    - `hardserver.log`
    - `keysafe.log`
    - `cmdadp.log`
    - `ncsnmpd.log`
- Any attached diagnostic files

Because the contents of the output file are plain text, they are human readable. You can choose to view the output file to ensure no sensitive information has been included.

> ℹ️ The `nfdiag` utility does not capture any passphrases in the output file.

22.2.1.3. NIC Reporting

Networking information regarding NIC (Network Interface Card) status and address details can be reported via the `stattree` output.

NIC status and address details are disabled by default. To enable or disable NIC details in the `stattree` output is done via the configuration menu, **System** > **System Configuration** > **NIC exposure config** > **Device Status** and select the appropriate setting (**Yes** or **No**). Similarly for enabling or disabling the NIC addressing which is done via the **Address Reporting** option.

> Reporting NIC addressing is only performed when device status is enabled.

A typical `stattree` report with NIC status and address details enabled looks like:

```
+#HostSysInfo:
   +SystemFans:
      +#1:
         -CurrentFanRPM       6240
      +#2:
         -CurrentFanRPM       6240
      +#3:
         -CurrentFanRPM       6240
      +#4:
         -CurrentFanRPM       6300
   +NetworkInterfaces:
      +#1:
         -InterfaceName       eth0
         -InterfaceLinkState   up
         +NetworkAddresses:
            +#1:
               -InterfaceNetworkAddress 172.23.135.127
      +#2:
         -InterfaceName       eth1
         -InterfaceLinkState   down
         +NetworkAddresses:
```

For more information, see stattree: information utility.

22.2.1.4. nShield HSM tamper log

The nShield HSM tamper log contains:

- The date and time of any tamper events
- Whether the tamper detection functionality was ever disabled or re-enabled.

> It is no longer possible to disable tamper detection.

You view the tamper log using the nShield HSM front panel, by selecting **System** > **System information** > **View tamper log**. You cannot erase the tamper log.

22.2.2. nfkminfo: information utility

The `nfkminfo` utility displays information about the Security World and the keys and card sets associated with it.

22.2.2.1. Usage

```
nfkminfo -w|--world-info [-r|--repeat] [-p|--preload-client-id]
```

```
nfkminfo -k|--key-list [<APPNAME> [<IDENT>]]
```

```
nfkminfo -l|--name-list [<APPNAME> [<APPNAME>...]]
```

```
nfkminfo [-c|--cardset-list]|[-s|--softcard-list] [<TOKENHASH>]
```

```
nfkminfo --cardset-list [<TOKENHASH>] --key-list [<APPNAME>[<APPNAME>]]|--name-list <APPNAME>[<IDENT>...]]
```

22.2.2.1.1. Security World options

`-w|--world-info`

This option specifies that you want to display general information about the Security World. These options are the default and need not be included explicitly.

`-r|--repeat`

This option displays the information repeatedly. There is a pause at the end of each set of information. The information is displayed again when you press Enter.

`-p|--preload-client-id`

This option displays the preloaded client ID value, if any.

22.2.2.1.2. Key, card set, and softcard options

`-k|--key-list [<APPNAME>[<APPNAME>]]`

This option lists keys without key names. If `<APPNAME>` is specified, only keys for these applications are listed.

`-l|--name-list [<APPNAME>[<IDENT>]]`

This option lists keys with their names. If `<APPNAME>` is specified, only keys for these applications are listed. If `<IDENT>` is listed, only the keys with the specified identifier are listed.

-c|--cardset-list [<TOKENHASH>]

If `<TOKENHASH>` is not specified, this option lists the card sets associated with the Security World. The output is similar to this:

```
Cardset list - 1 cardsets: (P)ersistent/(N)ot, (R)emoteable/(L)ocal-only
Operator logical token hash    k/n timeout name <hash>                    1/1 none-PL <name>
```

If `<TOKENHASH>` is specified, these options list the details of the card identified by *hash*. The output is similar to this:

```
Cardset
  name       "name"
  k-out-of-n 1/1
  flags      Persistent PINRecoveryForbidden(disabled) !RemoteEnabled
  timeout    none
  card names ""
  hkltu      hash
  gentime 2005-10-14 10:56:54
Keys protected by cardset hash:
 AppName app Ident keyident
 AppName app Ident keyident
 ...     ...  ...   ...
```

-s|--softcard-list TOKENHASH

This option works like the `-c|--cardset-list` option, except it lists softcards instead of card sets. If `<TOKENHASH>` is not specified, this option lists the softcards associated with the Security World.

22.2.2.2. Security World output info

If you run `nfkminfo` with the `-w|--world-info` option, it displays information similar to that shown in these examples:

```
generation 1
state     0x70000 Initialised Usable Recovery !PINRecovery
!ExistingClient !RTC !NVRAM !FTO !SEEDebug
n_modules  1
hknso      hash_knso
hkm        hash_km
hkmwk      hash_knwk
hkre       hash_kre
hkra       hash_kra
ex.client  none

...
```

```
...
Module #1
 generation    1
 state         0x1 Usable
 flags         0x10000 ShareTarget
 n_slots       2
 esn           34F3-9CB4-753B
 hkml          hash_kml
 Module #1 Slot #0 IC 11
 generation    1
 phystype      SmartCard
 slotlistflags 0x2
 state         0x4 Operator
 flags         0x20000 RemoteEnabled
 shareno       2
 shares
 error         OK
 Cardset
 name          "fred"
 k-out-of-n    1/2
 flags         NotPersistent
 timeout       none
 card names    "" ""
 hkltu         hash_kt
 Module #1 Slot #1 IC 0
 generation    1
 phystype      SmartCard
 slotlistflags 0x2 SupportsAuthentication
 state         0x4 Admin
 flags         0x10000 passphrase
 shareno       1
 shares        LTNSO(PIN) LTM(PIN) LTR(PIN) LTNV(PIN) LTRTC(PIN) LTDSEE(PIN)
 LTFTO(PIN)
 error         OK
 No Cardset

 No Pre-Loaded Objects
```

22.2.2.2.1. World

`nfkminfo` reports the following information about the Security World:

generation

This indicates the internal number.

state

This indicates the status of the current world:

Initialised	This indicates that the Security World has been initialized.
Usable	This indicates that there is at least one usable HSM in this Security World on this host.
!Usable	This indicates that there are no usable HSMs in this Security World on this host.

Recovery	This indicates that the Security World has the OCS and softcard replacement and the key recovery features enabled.
!Recovery	This indicates that the Security World has the OCS and softcard replacement and the key recovery features disabled.
AdminAuthRequired	This indicates that additional authorization is required for the following operations: • Key generation • Public key import • Operator cardset creation • Softcard creation. This authorization is supplied by presenting any operator or administration card from the same Security World. A passphrase is not required.
ExistingClient	This indicates that there is a Client ID set, for example, by preload. This Client ID is given in the ex.client output if the --preload-client-id flag was supplied.
!ExistingClient	This indicates that no Client ID is set. The ex.client output will be empty.
AlwaysUseStrongPrimes	This indicates that the Security World always generates RSA keys in a manner compliant with FIPS 186-3.
!AlwaysUseStrongPrimes	This indicates that the Security World leaves the choice of RSA key generation algorithm to individual clients.
SEEDebug	This indicates that the Security World has an SEE Debugging delegation key.
!SEEDebug	This indicates the Security World has no SEE Debugging delegation key.
SEEDebugForAll	This indicates no authorization is required for SEE Debugging.
PINRecovery	This indicates that the Security World has the passphrase replacement feature enabled.
!PINRecovery	This indicates that the Security World has the passphrase replacement feature disabled.
FTO	This indicates that the Security World has an FTO delegation key.
!FTO	This indicates that the Security World has no FTO delegation key.
NVRAM	This indicates that the Security World has an NVRAM delegation key.
!NVRAM	This indicates that the Security World has no NVRAM delegation key.
RTC	This indicates that the Security World has an RTC delegation key.
!RTC	This indicates that the Security World has no RTC delegation key.
AuditLogging	This indicates that Audit Logging is enabled for this Security World.
!AuditLogging	This indicates that Audit Logging is not enabled for this Security World.

`n_modules`

This indicates the number of nShield HSMs connected to this computer.

`hknso`

This indicates the SHA-1 hash of the Security Officer's key.

`hkm`

This indicates the SHA-1 hash of the Security World key.

`hkmwk`

This indicates the SHA-1 hash of a dummy key used to load the Administrator Card Set (the dummy key is the same on all HSMs that use Security Worlds and is not secret).

`hkre`

This indicates the SHA-1 hash of the recovery key pair.

`hkra`

This indicates the SHA-1 hash of the recovery authorization key.

`ex.client`

This indicates the `ClientID` required to use any pre-loaded keys and tokens.

`k-out-of-n`

This indicates the values of *K* and *N* for this Security World.

`other quora`

This indicates the number (quora) of Administrator Cards (*K*) required to perform certain other functions as configured for this Security World.

`ciphersuite`

This indicates the name of the Cipher suite that the Security World uses.

`Mode`

`none`	This indicates that the Security World is in an unregulated mode. The Security World can be configured to meet the needs of your security policy. This includes, but is not limited to, creating a Security World that is compliant with FIPS140 Level 2.
`fips1402level3`	This indicates that the Security World is in a mode compliant with FIPS 140 Level 3.

commoncriteriacmts	This indicates that the Security World is in a mode compliant with Common Criteria Protection Profile EN 419 221-5, for Cryptographic Modules for Trust Services.

Assigned Keys

max usage	This indicates the maximum key usage reauthorization condition for Assigned Keys. (common-criteria-cmts mode only).
max timeout	This indicates the maximum key timeout reauthorization condition for Assigned Keys (common-criteria-cmts mode only).

22.2.2.2.2. Module

For each HSM in the Security World, `nfkminfo` reports:

generation

This indicates the version of the HSM data.

state

This indicates one of the following:

PreInitMode	This indicates that the HSM is in the pre-initialization state.
InitMode	This indicates that the HSM is in the initialization state.
Unknown	This indicates that the HSM's state could not be determined.
Usable	This indicates that the HSM is programmed in the current Security World and can be used.
Uninitialized	This indicates that the HSM does not have the Security Officer's key set and that the HSM must be initialized before use.
Factory	This indicates that the HSM has module key zero only and that the Security Officer's key is set to the factory default.
Foreign	This indicates that the HSM is from an unknown Security World.
AccelOnly	This indicates that the HSM is acceleration only.
Unchecked	This indicates that, although the HSM appears to be in the current Security World, `nfkminfo` could not find a module initialization certificate (a `module_<ESN>` file) for this HSM.

Failed	This indicates that the HSM has failed. For internal security modules running firmware enquiry utility for further information about the failure reason, or look for hardware errors in the hardserver log.
MaintMode	This indicates that the HSM is in the maintenance state.

flags

This displays ShareTarget if the HSM has been initialized to allow reading of remote card sets.

n_slots

This indicates the number of slots on the HSM (there is one slot for each physical smart card reader, one slot for each soft token, one slot for each available Remote Operator slot and one slot for each associated Dynamic Slots).

esn

This indicates the electronic serial number of the HSM (if the HSM is not in the `Usable` state, the electronic serial number may not be available).

hkml

This indicates the hash of the HSM signing key (if the HSM is not in the `Usable` state, this value may not be available).

22.2.2.2.3. Slot

For each slot on the HSM, `nfkminfo` reports:

IC

This indicates the insertion count for this slot (which is 0 if there is no card in the slot).

generation

This indicates the version of the `slotinfo` structure.

phystype

This indicates the type of slot, which can be one of:

- `SmartCard`
- `SoftToken`

slotlistflags

These are flags describing the capabilities of the slot. Single letters in parentheses are the flag codes reported by the `slotinfo` utility:

`0x2`	`(A) SupportsAuthentication` This indicates that the slot supports token-level challenge-response authentication.
`0x40000`	`(R) RemoteSlot` This indicates that the slot is a Remote Operator slot that has been imported from a remote HSM.
`0x80000`	`(D) DynamicSlot` This indicates that it is a Dynamic Slot.
`0x100000`	`(a) Associated` This indicates that a Remote Administration Client has associated a card reader with this
`0x200000`	`(t) TimedOut` This indicates that no response has been received from the smartcard in this Dynamic Slot within the configured timeout.
`0x400000`	`(f) SecureChannelFailed` This indicates that the secure channel between the HSM and the smartcard in this Dynamic Slot has failed in some way.

state

This can be one or more of the following flags:

`Blank`	This indicates that the smart card in the reader is unformatted.
`Admin`	This indicates that the smart card in the reader is part of the Administrator Card Set.
`Empty`	This indicates that there is no smart card in the reader.
`Error`	This indicates that the smart card in the reader could not be read (the card may be from a different Security World).
`Operator`	This indicates that the smart card in the reader is an Operator Card.

flags

This displays `passphrase` if the smart card requires a passphrase.

shareno

This indicates the number of the card within the card set.

shares

If the card in the slot is an Operator Card, no values are displayed for `shares`.

If the card in the slot is an Administrator Card, values are displayed indicating what key shares are stored on the card. Each share is prefixed with the letters `LT` (Logical Token), and the remaining letters identify the key (for example, the value `LTNSO` indicates that a share of K_{NSO}, the Security Officer's key, is stored on the card).

error

This indicates the error status encountered if the smart card could not be read:

OK	This indicates that there were no errors.
TokenAuthFailed	This indicates that the smart card in the reader failed challenge response authentication (the card may come from a different Security World).
PhysTokenNotPresent	This indicates that there is no card in the reader.

If you purchased a developer kit, you can refer to the relevant developer documentation for a full list of error codes.

22.2.2.2.4. Card set

If there is an Operator Card in the reader, `nfkminfo` reports:

name

This indicates the name given to this card set.

k-out-of-n

This indicates the values of *K* and *N* for this card.

flags

This displays one or more of each of the following pairs of flags:

NotPersistent	This indicates that the Operator Card is not persistent.
Persistent	This indicates that the Operator Card is persistent.
NotRemoteEnabled	This indicates that the card in the slot is not from a Remote Operator Card Set.
RemoteEnabled	This indicates that the card in the slot is from a Remote Operator Card Set.
PINRecoveryForbidden(dis-abled)	This indicates that the card in the slot does not have passphrase replacement enabled. This is always true if passphrase replacement is disabled for the Security World.

`PINRecoveryRe-` `quired(enabled)`	This indicates that the card in the slot does have passphrase replacement enabled.

`timeout`

the period of time in seconds after which the HSM automatically removes the Operator Card Set. If timeout is set to `none`, the Operator Card Set does not time out.

`card`

lists the names of the cards in the set, not all software can give names to individual cards in a set.

`hkltu`

the SHA-1 hash of the secret on the card.

22.2.3. perfcheck: performance measurement checking tool

Use the `perfcheck` command-line utility to run various tests measuring the cryptographic performance of an nShield HSM.

Run `perfcheck` with the standard `-h|--help` option to display information about the options and parameters that control the program's behavior.

The available tests are grouped into suites:

- `kx` (key exchange)
- `keygen` (key generation)
- `signing` (signing)
- `verify` (verification)
- `enc` (encryption)
- `dec` (decryption)
- `misc` (miscellaneous)

To see the list of tests run by default in a particular suite, run a command of the form:

```
perfcheck --list suite
```

To see all available tests in a particular suite, including those not run by default, run a command of the form:

```
perfcheck --list --depth=full suite
```

For example, to list all the `signing` tests, run the command:

```
perfcheck --list --depth=full signing
>>> Suite `signing' -- Signing (374 tests)
>>>    signing 1 - DSA using RIPEMD160 with 1024-bit p and 160-bit q.
>>>    signing 2 - DSA using RIPEMD160 with 2048-bit p and 160-bit q.
>>>    signing 3 - DSA using RIPEMD160 with 3072-bit p and 160-bit q.
>>>    signing 4 - DSA using SHA1 with 1024-bit p and 160-bit q.
>>>    signing 5 - DSA using SHA1 with 2048-bit p and 160-bit q.
>>>    signing 6 - DSA using SHA1 with 3072-bit p and 160-bit q.
```

In the output, each listed test in the suite is identified with a number.

You can reference a test either by its number or by its name:

- by test number:

  ```
  perfcheck suite:test_number
  ```

  To use test `16` of the `signing` suite:

  ```
  perfcheck signing:16
  ```

- by test name:

  ```
  perfcheck "signing:RSA using RSAhSHA3b512pPSS with 4096-bit n."
  ```

  Example:

  ```
  perfcheck "signing:RSA using RSApPKCS1 with 2048-bit n."
  ```

The test numbers change between releases. If you want to rerun tests for comparison, reference the tests by their names.

`perfcheck` prints the results of individual tests to output as it goes along, and then prints a full report at the end. By default, `perfcheck` runs each test three times for both minimum and maximum queue sizes, and then collates the results in the final report. See `--help` for the options to adjust this behavior.

Optionally, `perfcheck` can write its output to a directory in multiple formats using the `--outputdir` option to specify a directory name. This will create a new subdirectory under the specified directory to write the output. The `--nosubdir` option can be added as well to write output to the specified directory directly, in which case that directory must not already exist. The output directory will contain `perfcheck.html`, `perfcheck.txt`, `perfcheck.csv`, and `perfcheck.json` files that contain the report in HTML, text, CSV, and JSON format respec-

tively. JSON files that contain the detailed results of individual tests will also be written to the output directory.

Output reports from test suites include the following information about each test:

Value	Description
Queue	This value is the number of outstanding jobs in the queue when the test was run. By default, most tests run both with a queue of 1, and with a fully maxed out module queue, to give an indication of both one-at-a-time performance and the bandwidth for the operation. The queue can be set differently using the `--queue` option, in which case only that queue length will be run with, except for some `misc` suite tests which set their own queue.
Rate (Units/s)	This value is a measure of throughput. It is calculated by dividing the number of repetitions by total time. If a test has been rerun to improve accuracy, as is the case by default, then this is the mean across all the runs. Some tests, for example `enc`, set the Unit to something other than an operation, for example KB, to indicate the amount of data that can be encrypted.
Min latency (ms)	This value is the time in milliseconds that the quickest individual job across all the test runs took to round-trip.
Mean latency (ms)	This value is the mean time in milliseconds that jobs took to round-trip. If a test has been rerun, this is the mean of the mean latency values from each run.
Max latency (ms)	This value is the time in milliseconds that the slowest individual job across all the test runs took to round-trip.
CV (%)	This value is the coefficient of variation expressed as a percentage of the mean latency. It gives an indication of the variability in the time it takes individual jobs to complete. If a test has been rerun, this is the mean of the CV (%) values from each run.
Min rate (tps)	This is the estimated lower bound of the throughput for this queue size in transactions per second. The value becomes more accurate if more test runs of the same test are done. When it is compared against Mean rate (tps) and Max rate (tps), Min rate (tps) gives an indication of the variability between runs.

Value	Description
Mean rate (tps)	This is a measure of throughput. Unlike Rate (Units/s), it is expressed in transactions per second, that is, as the number of jobs that round-trip per second. Mean rate (tps) is included for comparison against the Min rate (tps) and Max rate (tps) figures.
Max rate (tps)	This is the estimated upper bound of the throughput for this queue size in transactions per second. The value becomes more accurate if more test runs of the same test are done. When it is compared against Min rate (tps) and Mean rate (tps), Max rate (tps) gives an indication of the variability between runs.
Reps	This value is the number of repetitions that were actually carried out, that is, the number of jobs that were round-tripped over all tests of this operation for this queue size. If a test was rerun, this is the sum of the repetitions for each run. The target repetitions for an individual run can be set using the `--repetitions` option but note that in most cases more repetitions will be run depending on the `--accuracy` setting provided that the timeout is not reached. It is recommended to set `--accuracy` rather than `--repetitions` to control the accuracy of the test instead of adjusting the repetitions.

22.2.3.1. How perfcheck calculates statistics

When an nCore command is submitted to an HSM by a client application, it is processed as follows:

1. The command is passed to the client hardserver.
2. The client hardserver encrypts the command.
3. The encrypted command is sent to the unit hardserver over the network.
4. The unit hardserver decrypts the command and queues it.
5. When the internal security module is free, the command is submitted from the hardserver queue.
6. The command is executed by the HSM, and the reply is given to the unit hardserver.
7. The unit hardserver encrypts the command.
8. The unit hardserver sends the command back to the client hardserver over the network.
9. The client hardserver decrypts the reply and queues it.
10. When the client application is ready, the queued reply is returned to it.

Because an HSM can execute several commands at once, throughput is maximized by ensuring there is always at least one command in the unit hardserver queue (so that there are always commands available to give to the HSM).

The `perfcheck` utility sends multiple simultaneous nCore commands to keep the HSM busy. It can send more commands if a required number of repetitions has not yet been reached.

After sending some initial commands, `perfcheck` begins marking commands with the time at which are submitted. When a command comes back with a timestamp, `perfcheck` checks the amount of time needed to complete the command and updates the values for `Std dev` and `Latency`. The value of `Total time` is the amount of time from sending the first job to receiving the final one.

22.2.4. stattree: information utility

The `stattree` utility returns the statistics gathered by the hardserver and HSMs.

22.2.4.1. Usage

```
stattree [<node> [<node> [...]]]
```

22.2.4.2. Output

Running the `stattree` utility displays a snapshot of statistics currently available on the host machine. Statistics are gathered both by the hardserver (relating to the server itself, and its current clients) and by each attached HSM.

Times are listed in seconds. Other numbers are integers, which are either real numbers, IP addresses, or counters. For example, a result `-CmdCount 74897` means that there have been 74,897 commands submitted.

A typical fragment of output from `stattree` looks like this:

```
+PerModule:
  +#1:
    +ModuleObjStats:
      -ObjectCount       5
      -ObjectsCreated    5
      -ObjectsDestroyed  0
    +ModuleEnvStats:
      -MemTotal          15327232
      -MemAllocKernel    126976
      -MemAllocUser      0
    +ModuleJobStats:
      -CmdCount          169780
      -ReplyCount        169778
```

```
      -CmdBytes             3538812
      -ReplyBytes           4492764
      -HostWriteCount       169772
      -HostWriteErrors      0
      -HostReadCount        437472
      -HostReadErrors       0
      -HostReadEmpty        100128
      -HostReadDeferred     167578
      -HostReadTerminated   0
      -PFNIssued            102578
      -PFNRejected          1
      -PFNCompleted         102577
      -ANIssued             1
      -CPULoadPercent       0
   +ModuleSerialStats:
      -HostReadCount        437476
      -HostReadDeferred     167580
      -HostReadReconnect    167579
      -HostReadErrors       0
      -HostWriteCount       169774
      -HostWriteErrors      0
   +ModuleDriverStats:
       -DriverIRQs           2547906
       -DriverReadIRQs       1274069
       -DriverWriteIRQs      1276373
       -DriverWriteFails     0
       -DriverWriteBlocks    1276373
       -DriverWriteBytes     49625888
       -DriverReadFails      0
       -DriverReadBlocks     0
       -DriverReadBytes      0
       -DriverEnsureFail     0
       -DriverEnsure         1274065
```

`PerModule`, `ModuleObjStats`, and `ModuleEnvStats` are node tags that identify classes of statistics. `1` identifies an instance node.

`ObjectCount`, `MemTotal`, and the remaining items at the same level are statistics `ID`s. Each has a corresponding value.

If `<node>` is provided, `stattree` uses the value given as the starting point of the tree and displays only information at or below that node in the tree. Values for `<node>` can be numeric or textual. For example, to view the object counts for local module number 3:

```
$ stattree PerModule 3 ModuleObjStats
+#PerModule:
   +#3:
      +#ModuleObjStats:
          -ObjectCount       6
          -ObjectsCreated    334
          -ObjectsDestroyed  328
```

The value of `<node>` must be a node tag; it must identify a node in the tree and not an individual statistic. Thus, the following command does not work:

```
$ stattree PerModule 3 ModuleObjStats ObjectCount
+#PerModule:
```

```
  +#3:
    +#ModuleObjStats:
 Unable to convert 'ObjectCount' to number or tag name.
```

`ModuleDriverStats` fields:

Field	Description
DriverIRQs	Total number of interrupts
DriverReadIRQs	Read interrupts
DriverWriteIRQs	Write interrupts
DriverWriteFails	Write failures
DriverWriteBlocks	Blocks written
DriverWriteBytes	Bytes written
DriverReadFails	Read failures
DriverReadBlocks	Blocks read
DriverReadBytes	Bytes read
DriverEnsureFail	Read request failures
DriverEnsure	Read requests

22.2.4.2.1. Node tags

These hold statistics for each HSM:

Category	Contains
ModuleJobStats	This tag holds statistics for the Security World Software commands (jobs) processed by this HSM.
ModulePCIStats	This tag does not apply to nShield USB-attached HSMs.
ModuleSerialStats	This tag only applies to nShield USB-attached HSMs. It holds statistics for the serial connection between the HSM and the host computer.
ServerGlobals	Aggregate statistics for all commands processed by the hardserver since it started. The standard statistics (as described below) apply to the commands sent from the hardserver to HSMs. Commands processed internally by the server are not included here. The Uptime statistic gives the total running time of the server so far.

Category	Contains
Connections	Statistics for connections between clients and the hardserver. There is one node for each currently active connection. Each node has an instance number that matches the log message generated by the server when that client connected. For example, when the hardserver message is `Information: New client #24 connected`, the client's statistics appear under node #24 in the `stattree` output.
PerModule	Statistics kept by the HSMs. There is one instance node for each HSM, numbered using the standard HSM numbering. The statistics provided by each HSM depend on the HSM type and firmware version.
ModuleObjStats	Statistics for the HSM's Object Store, which contains keys and other resources. These statistics may be useful in debugging applications that leak key handles, for example.
ModuleEnvStats	General statistics for the HSM's operating environment.
HostEnvStats	Environmental statistics for the HSM.
HostSysInfo	Further statistics for the HSM.

22.2.4.2.2. Statistics IDs

ID	Value
Uptime	The length of time (in seconds) since an HSM was last reset, the hardserver was started, or a client connection was made.
CmdCount	The total number of commands sent for processing from a client to the server, or from the server to an HSM. Contains the number of commands currently being processed.
ReplyCount	The total number of replies returned from server to client, or from HSM to server.
CmdBytes	The total length of all the command blocks sent for processing.
ReplyBytes	The total length of all the reply blocks received after completion.
CmdMarshalErrors	The number of times a command block was not understood when it was received. A nonzero value indicates either that the parties at each end of a connection have mismatched version numbers (for example, a more recent hardserver has sent a command to a less recent HSM that the HSM does not understand), or that the data transfer mechanism is faulty.
ReplyMarshalErrors	The number of times a reply was not understood when it was received. A nonzero value indicates either that the parties at each end of a connection have mismatched version numbers (for example, a more recent hardserver has sent a command to a less recent HSM that the HSM does not understand), or that the data transfer mechanism is faulty.

ID	Value
ClientCount	The number of client connections currently made to the server. This appears in the hardserver statistics.
MaxClients	The maximum number of client connections ever in use simultaneously to the hardserver. This gives an indication of the peak load experienced so far by the server.
DeviceFails	The number of times the hardserver has declared a device to have failed. The hardserver provides a diagnostic message when this occurs.
DeviceRestarts	The number of times the hardserver has attempted to restart an HSM after it has failed. The hardserver provides a Notice message when this occurs. The message does not indicate that the attempt was successful.
QOutstanding	The number of commands waiting for an HSM to become available on the specified client connection. When an HSM accepts a command from a client, this number decreases by 1 and DevOutstanding increases by 1. Commands that are processed purely by the server are never included in this count.
DevOutstanding	The number of commands sent by the specified client that are currently executing on one or more HSMs. When an HSM accepts a command from a client, this number increases by 1 and QOutstanding decreases by 1. Commands that are processed purely by the server are never included in this count.
LongOutstanding	The number of LongJobs sent by the specified client that are currently executing on one or more HSMs. When an HSM accepts a LongJobs command from a client, this number increases by 1 and QOutstanding decreases by 1. Commands that are processed purely by the server are never included in this count.
RemoteIPAddress	The remote IP address of a client who has this connection. A local client has the address 0.0.0.0.
HostWriteCount	The number of write operations (used to submit new commands) that have been received by the HSM from the host machine. One write operation may contain more than one command block. The operation is most efficient when this is the case.
HostWriteErrors	The number of times the HSM rejected the write data from the host. A nonzero value may indicate that data is being corrupted in transfer, or that the hardserver/device driver has got out of sync with the HSM's interface.
HostWriteBadData	Not currently reported by the HSM. Attempts to write bad data to the HSM are reflected in HostWriteErrors.
HostWriteOverruns	Not currently reported by the HSM. Write overruns are reflected in HostWriteErrors.
HostWriteNoMemory	Not currently reported by the HSM. Write failures due to a lack of memory are reflected in HostWriteErrors.

ID	Value
HostReadCount	The number of times a read operation to the HSM was attempted. The HSM can defer a read if it has no replies at the time, but expects some to be available later. Typically the HSM reports HostReadCount in two places: the number under ModuleJobStats counts a deferred read twice, once when it is initially deferred, and once when it finally returns some data. The number under ModulePCIStats counts this as one operation.
HostReadErrors	The number of times a read to an HSM failed because the parameters supplied with the read were incorrect. A nonzero value here typically indicates some problem with the host interface or device driver.
HostReadEmpty	The number of times a read from the HSM returned no data because there were no commands waiting for completion. In general, this only happens infrequently during HSM startup or reset. It can also happen if PauseForNotifications is disabled.
HostReadUnderruns	Not currently reported by the HSM.
HostReadDeferred	The number of times a read operation to the HSM was suspended because it was waiting for more replies to become available. When the HSM is working at full capacity, a sizeable proportion of the total reads are likely to be deferred.
HostReadTerminated	The number of times an HSM had to cancel a read operation which has been deferred. This normally happens only if the clear key is pressed while the HSM is executing commands. Otherwise it might indicate a device driver, interface, or firmware problem.
PFNIssued	The number of PauseForNotifications commands accepted by the HSM from the hardserver. This normally increases at a rate of roughly one every two seconds. If the hardserver has this facility disabled (or a very early version), this does not occur.
PFNRejected	The number of PauseForNotifications commands rejected by the HSM when received from the hardserver. This can happen during HSM startup or reset, but not in normal use. It indicates a hardserver bug or configuration problem.
PFNCompleted	The number of PauseForNotifications commands that have been completed by the HSM. Normally, this is one less than the PFNIssued figure because there is normally one such command outstanding.
ANIssued	The number of Asynchronous Notification messages issued by the HSM to the hardserver. These messages indicate such things as the clear key being pressed and the HSM being reset. In later firmware revisions inserting or removing the smartcard or changing the non-volatile memory also generate asynchronous notifications.
ChanJobsIssued	The number of fast channel jobs issued to the HSM. The fast channel facility is unsupported on current HSMs. This number should always be 0.
ChanJobsCompleted	The number of fast channel jobs completed by the HSM. The fast channel facility is unsupported on current HSMs. This number should always be 0.

ID	Value
`CPULoadPercent`	The current utilization of the main CPU, across all cores. If you are on a firmware version earlier than 13.1, this instead reports a load average that is scaled by 100, but could be greater than 100% if there is an average of more than one runnable thread.
`HostIRQs`	On PCI HSMs, the total number of interrupts received from the host. On current HSMs, approximately equal to the total of `HostReadCount` and `HostWrite-Count`.
`ChanJobErrors`	The number of low-level (principally data transport) errors encountered while processing fast channel jobs. Should always be `0` on current HSMs.
`HostDebugIRQs`	On PCI HSMs, the number of debug interrupts received. This is used only for driver testing, and should be `0` in any production environment.
`HostUnhandledIRQs`	On PCI HSMs, the number of unidentified interrupts from the host. If this is nonzero, a driver or PCI bus problem is likely.
`HostReadReconnect`	On PCI HSMs, the number of deferred reads that have now completed. This should be the same as `HostReadDeferred`, or one less if a read is currently deferred.
`ObjectsCreated`	The number of times a new object has been put into the object store. This appears under the HSM's `ModuleObjStats` node.
`ObjectsDestroyed`	The number of items in the HSM's object store that have been deleted and their corresponding memory released.
`ObjectCount`	The current number of objects (keys, logical tokens, buffers, SEE Worlds) in the object store. This is equal to `ObjectsCreated` minus `ObjectsDestroyed`. An empty HSM contains a small number of objects that are always present.
`CurrentTempC`	The current temperature (in degrees Celsius) of the HSM main circuit board. First-generation HSMs do not have a temperature sensor and do not return temperature statistics.
`MaxTempC`	The maximum temperature recorded by the HSM's temperature sensor. This is stored in non-volatile memory, which is cleared only when the unit is initialized. First-generation HSMs do not have a temperature sensor and do not return temperature statistics.
`MinTempC`	The minimum temperature recorded by the HSM's temperature sensor. This is stored in non-volatile memory, which is cleared only when the unit is initialized. First-generation HSMs do not have a temperature sensor and do not return temperature statistics.
`MemTotal`	The total amount of RAM (both allocated and free) available to the HSM. This is the installed RAM size minus various fixed overheads.
`MemAllocKernel`	Obsolete, retained only for backwards compatibility. Shows the same value as `MemTotal`.

ID	Value
MemAllocUser	Obsolete, retained only for backwards compatibility. Shows the same value as MemTotal.
SystemFans	The fan speed (RPM) for each fan in the HSM.
NVMFreeSpace	The total amount of free space in the NVRAM of the HSM, in bytes.
NVMWearLevel	The wear level of the HSM's NVRAM, expressed as a percentage of the ratio between the erase count and the endurance.
NVMWornBlocks	The percentage of worn blocks in the NVRAM of the HSM.

22.3. How data is affected when a module loses power and restarts

nShield modules use standard RAM to store many kinds of data, and data stored in such RAM is lost in the event that a module loses power (either intentionally, because you turned off power to it, or accidentally because of a power failure).

Therefore, after restoring power to a module, you must reload any keys that had been loaded onto it before it lost power. After reloading, the KeyIDs are different.

Likewise, after restoring power to a module, you must reload any cards that were loaded onto it before it lost power.

However, data stored in NVRAM is unaffected when a module loses power.

> ℹ️ If you are using multiple nShield modules in the same Security World, and have the same key (or keys) loaded onto each module as part of a load-sharing configuration, loss of power to one module does not affect key availability (as long as at least one other module onto which the keys are loaded remains operational). However, in such a multiple-module system, after restoring power to a module, you must still reload any keys to that module before they can be available from that module.

23. HSM and client configuration files

This appendix describes the configuration files that store the current configuration of an nShield HSM or client.

The *module configuration files* are stored on the internal file system of the nShield HSM. They are updated automatically when the nShield HSM is configured from the front panel. The configuration files are also exported automatically to the remote file system, where they can be edited manually and reloaded on the nShield HSM, if required. For more information, see Client Software and module configuration.

The *client configuration files* are stored in the client's file system. The client's hardserver can also be set up using environment variables. Environment variable settings override settings in the client configuration files. For more information, see Setting environment variables.

23.1. Location of client configuration files

The client configuration files are in `opt/nfast/kmdata/config/` (**Linux**) or `%NFAST_KM-DATA%\config` (**Windows**) on the client computer's file system. This directory can contain the following files:

File	Description
config	The master configuration file. This contains the current configuration for the client. It is always present in the directory.
config-default	The default configuration file. This can be used to restore factory settings for the client. It is created by the `cfg-mkdefault` utility.

You can also save backup copies of configuration files in this directory.

23.2. Location of module configuration files

When you configure the nShield HSM from the front panel, the configuration files are exported automatically to the remote file system defined in the `rfs_client` section of the module configuration. The exported files are in `/opt/nfast/kmdata/hsm-ESN/config` (**Linux**) or `%NFAST_KMDATA%\hsm-ESN\config` (**Windows**).

In this path, *ESN* is the electronic serial number of the network nShield HSM from which the files were exported. This directory contains the master configuration file `config`, which specifies the current configuration for the nShield HSM. It is always present in the directory.

The remote file system information is also contained in the client configuration file section `remote_file_system`.

23.3. Structure of configuration files

The configuration files are text files. They must contain only characters with ASCII values between 32 and 127, and the tab, line break, and return characters.

Lines starting with one or more number sign characters (#) are comments and are ignored. Some comments that document the configuration options are generated by the configuration process. You can add your own comments, but in some cases they may later be overwritten.

A configuration file begins with a single line that specifies the version of the file syntax. This syntax-version line has the format:

```
syntax-version=n
```

In this syntax-version line example, *n* represents the version of the syntax in which the file is written. The system can process a file with a lower syntax version than the one it uses, but not one with a higher version.

After the syntax-version line, the rest of the configuration file consists of sections that can be edited to control different aspects of nShield HSM or client behavior. Each section begins with its name in square brackets, as in this example:

```
[slot_imports]
```

Module configuration files and client configuration files have some sections in common, and some specific to the type of file:

Both	Module file only	Client file only
`server_settings`	`nethsm_settings`	`module_settings`
`server_remotecomms`	`nethsm_eth`	`server_performance`
`server_remotecomms_ipv6`	`nethsm_eth_ipv6`	
`server_startup`	`nethsm_gateway`	`nethsm_imports`
	`nethsm_gateway_ipv6`	
`load_seemachine`	`nethsm_bond`	`rfs_sync_client`

Both	Module file only	Client file only
slot_imports	nethsm_route nethsm_route_ipv6	remote_file_system
slot_exports	nethsm_eth1_enable	remote_administration_service_s-tartup
dynamic_slots	nethsm_bond_enable	
slot_mapping	nethsm_enable	
dynamic_slot_timeouts	cosmod	
auditlog_settings	hs_clients	
	rfs_client	
	sys_log	
	remote_sys_log	
	config_op	
	ui_lockout	

You can update the parameters defined in most of these sections to configure the way that the hardserver handles secure transactions between the nShield HSM and applications that run on the client.

> ℹ️ Some sections are updated automatically and should not be edited manually. For more information, see the descriptions of individual sections.

In each section, the bracketed name is followed by a specified set of fields. Each field is on a separate line. Each field begins with its name, followed by an equals sign (=) and a value of the appropriate type. White space can be included at either end of the line (for example, in order to indent lines as an aid to clarity).

Some types of field are grouped into entries. An entry is a set of fields of different types that define an instance of an object (for example, a particular client as distinct from other clients). Entries in the same section are separated by a line that contains one or more hyphens (-). Blank lines and comments are allowed between the fields in an entry.

Strings are case sensitive in the section names and field names.

> ℹ️ Multiple clients can be added to one configuration file by separating each client entry from the next with a line consisting of one or more hyphens.

If a particular section is not present in the configuration file, it is assumed to have no

entries.

23.4. Sections only in module configuration files

23.4.1. nethsm_settings

The `nethsm_settings` section defines settings specific to the nShield HSM. The section contains the following fields:

Field	Description
`enable_impath_resilience`	When set to the default `yes` value, this field enables impath resilience for the module. Setting this field to `no` disables impath resilience.
`impath_resilience_timeout`	This field specifies the interval of time that must pass for a resumable impath resilience session to expire. In the event of network errors, clients can reconnect and resume use of keys and other loaded objects until the specified interval has passed (after which all previously loaded objects must be reloaded). Specify this interval either in the form *N* `t`, where *N* is an integer and `t` is `s` for seconds, `m` for minutes, `h` for hours, `d` for days, or `w` for weeks, or as `never` (in which case sessions never expire). If you specify *N* but not `t`, the seconds are assumed. The default setting is `3600` for 3,600 seconds.
`expose_nic_status`	When set to the `yes` value, this field enables NIC details in the `stattree` output. Setting this field to default `no` disables NIC details.
`expose_nic_addresses`	When set to the `yes` value, this field enables NIC addresses in the `stattree` output, only if `expose_nic_status` is enabled. Setting this field to default `no` disables NIC addresses.

23.4.2. nethsm_eth

The `nethsm_eth` section defines the Ethernet interfaces for IPv4 for the nShield HSM. Each interface is defined by an entry as follows:

Field	Description
`iface`	The identifier for the interface. This value must be `1` or `0`.
`addr`	The IP address of the nShield HSM. The default is `0.0.0.0`.
`netmask`	The net mask for the nShield HSM. The default is `0.0.0.0`.

Field	Description
gateway	This field is retained for backwards compatibility only. The IP address of the gateway is stored in the `nethsm_gateway` section and this field is set to `0.0.0.0`.
linkspeed	This field specifies the link speed setting. It can be one of `auto/1Gb` (nShield 5s only), `10BaseT`, `10BaseT-FDX`, `100BaseTX`, or `100BaseTX-FDX`. We recommend that you accept the default `auto/1Gb` or `auto` setting. On the nShield HSM, this setting can auto-negotiate 1Gb Ethernet.

23.4.3. nethsm_eth_ipv6

The `nethsm_eth_ipv6` section defines the Ethernet interfaces for IPv6 for the nShield HSM. Each interface is defined by an entry as follows:

Field	Description
iface	The identifier for the interface. This value must be `1` or `0`.
addr	The static IPv6 address for this interface. The default is `::`. If SLAAC is enabled, this address is ignored.
netmask	The subnet prefix length of the static IPv6 address for the interface. The default is `64`.

23.4.4. nethsm_gateway

The `nethsm_gateway` section defines the default IPv4 Ethernet gateway for the nShield HSM. There is one field, as follows:

Field	Description
gateway	The IPv4 address of the gateway for the nShield HSM. The default is `0.0.0.0`.

23.4.5. nethsm_gateway_ipv6

The `nethsm_gateway_ipv6` section defines the default IPv6 Ethernet gateway for the nShield HSM. There is one field, as follows:

Field	Description
gateway	The IPv6 address of the gateway for the nShield HSM. The default is `::`.

Field	Description
`linklocal_if`	The ethernet interface (`0` or `1`) to use if the IPv6 default gateway address is a link-local address. The information is not used if the IPv6 default gateway is not a link-local address (default `0`).

23.4.6. nethsm_bond

The `nethsm_bond` section defines the HSM bond interface, used for network bonding. The bond interface's address and netmask configuration are inherited from the `eth0` (`iface=0`) configuration. Each entry has the following fields:

Field	Description
`mode`	Possible values: • `802.3ad` • `active-backup` Default: `802.3ad`.
`miimon`	The MII link monitoring frequency in milliseconds. Range: `0` - `10000`. Default: `100`.
`primary`	Primary device. The specified device will always be the active slave while it is available. Possible values: • `eth0` • `eth1` Default: `eth0`. Only valid for `active-backup` mode.
`resend_igmp`	The number of IGMP membership reports to be issued after a failover event. Range: `0` - `255`. Default: `1`. A value of `0` prevents the IGMP membership report from being issued in response to the failover event. Only valid for `active-backup` mode.

Field	Description

lacp_rate	The rate in which the Ethernet interface asks the link partner to transmit LACPDU packets in 802.3ad mode.
	
	Possible values:
	
	    • slow
	    • fast
	
	Default: slow.
	
	Only valid for 802.3ad mode.

xmit_hash_policy	The transmit hash policy to use for slave selection in 802.3ad mode.
	
	Possible values:
	
	    • layer2
	    • layer2+3
	    • encap2+3
	
	Default: layer2.

The process of network bonding, including all the fields above, are described in
https://www.kernel.org/doc/Documentation/networking/bonding.txt.

23.4.7. nethsm_route

The nethsm_route section defines the static IPv4 routes for the nShield HSM. Each route is defined by an entry as follows:

Field	Description
addr	The IPv4 address of the route destination. The default is 0.0.0.0.
masklen	The length to mask for the route.
gateway	The IPv4 address of the gateway for the route. The default is 0.0.0.0.

23.4.8. nethsm_route_ipv6

The nethsm_route_ipv6 section defines the static IPv6 routes for the nShield HSM. Each route is defined by an entry as follows:

Field	Description
addr	Routable IPv6 address block. The default is ::.

Field	Description
`masklen`	The number of bits to mask for the subnet address prefix, that is, the number in the range of `1` to `128`) after the `/` of an address in CIDR format. The default is `64`.
`gateway`	Route gateway. The default is `::`.
`linklocal_if`	The ethernet interface (`0` or `1`) to use if the IPv6 default gateway address is a link-local address. The information is not used if the IPv6 default gateway is not a link-local address (default `0`).

23.4.9. nethsm_eth1_enable

The `nethsm_eth1_enable` section defines if the eth1 interface is enabled.

Field	Description
`enable`	The indicator of whether the eth1 interface is enabled or disabled (default 'no').

23.4.10. nethsm_bond_enable

The `nethsm_bond_enable section` defines if the bond interface is enabled.

Field	Description
`enable`	The indicator of whether the bond interface is enabled or disabled (default 'no').

23.4.11. nethsm_enable

The `nethsm_enable` section defines whether IPv4 and or IPv6 are enabled or disabled for the interfaces of the unit. The enable fields are:

Field	Description
`iface`	The ethernet interface (`0` or `1`) to which the following fields apply.
`enable_ipv4`	Indicator of whether the IPv4 protocol on the interface is enabled (default: `yes`).
`enable_ipv6`	Indicator of whether the IPv6 protocol on the interface is enabled (default: `no`).
`ipv6_conf_addr`	Indicator of whether the interface uses IPv6 static addresses (`static`) or SLAAC (`slaac`). The default is `static`.

23.4.12. cosmod

The `cosmod` section defines the input configuration for the nShield HSM. The configuration is defined by an entry as follows:

Field	Description
keymap	The selected layout for the keyboard connected to the nShield HSM front panel. This value is either UK or US.

23.4.13. rfs_client

The `rfs_client` section defines the remote file system where the module configuration is backed up and the master copy of the Security World data is located, as follows:

Field	Description
remote_ip	The IP address of the RFS.
remote_port	The port number on which the RFS hardserver is listening.
remote_keyhash	Software or module KNETI hash used to authenticate the RFS, or 40 zeroes to indicate no authentication required (default is 40 zeroes).
remote_esn	ESN of the remote module used to authenticate the RFS, or empty when using software KNETI authentication or no authentication required (default is empty).
push	Whether to allow the RFS to push configuration files to the nShield HSM: • ON: This effectively allows the RFS to remotely configure the nShield HSM, and also to act as an additional privileged cryptographic client above the client licence limit. • OFF: This does not allow the RFS to remotely configure the nShield HSM. • AUTO (default): If the remote_keyhash field is set to a non-zero hash, behaviour shall be as though ON was set. If the remote_keyhash field is the all-zeroes hash, behaviour shall be as though OFF was set.

23.4.14. sys_log

The `sys_log` section defines how the nShield HSM logging works:

Field	Description
`behaviour`	The way the log is written. How the log is written is decided by setting either of the following: • `log` • `push` If `log` is set, the log is written only to the file system of the nShield HSM. It is lost if the nShield HSM is rebooted. Logging stops when the file system is full. If `push` is set, the log is automatically appended to the log file in the RFS or remote syslog server at the interval specified in `push_interval`.
`push_interval`	The number of minutes between the log being appended when `behaviour` is set to push. The default is `60`. The minimum is 1.

23.4.15. remote_sys_log

The `remote_sys_log` section defines a remote syslog server to send the nShield HSM logging (system.log and hardserver.log) to. It can be an IPv4 or an IPv6 address.

Field	Description
`remote_syslog_ip`	The IP address of the remote syslog server to send logs to. The default is `0.0.0.0`.
`remote_sys_log_port`	The UDP port of the remote syslog server to send logs to. The default is `514`.

23.4.16. config_op

The `config_op` section defines whether a client computer is allowed to update the module configuration automatically (to push a configuration) from files on the client:

Field	Description
`push`	Whether the client is allowed to push configuration files to the nShield HSM. This is decided by setting one of the following: • `ON` • `OFF` If `on`, the client specified in the `remote_ip` section is allowed to update the configuration of the nShield HSM remotely.
`remote_ip`	The IP address of the client that is allowed to push config files. If no value is set, or if it is set to `0.0.0.0` or `::`, any IP address can push on a new config file.

Field	Description
remote_keyhash	The hash of the key with which the authorized client is to authenticate itself, or (as the default) 40 zeros to indicate no key authentication required.

> ⓘ The default value for remote_keyhash (40 zeros) specifies that no authentication should occur. We recommend specifying a key hash in place of this default.

23.4.17. ui_lockout

The ui_lockout section defines whether you can lock the nShield HSM using login settings.

To be compatible with UI lockout, the OCS card(s):

- Must be persistent.
- Must not be remoteable.
- Do not need a passphrase, but if a passphrase is configured, it has to be used.

Field	Description

lockout_mode	Controls front panel access to the nShield HSM. Set to:
	
	 • locked
	
	   Enables UI lockout without requiring a logical token.
	
	 • locked_lt
	
	   Enables UI lockout with a logical token (OCS) (requires a valid ltui_hash to be set).
	
	 • unlocked
	
	   No UI lockout (default).

ltui_hash	The hash of the logical token (LTUI) required to authorize access to the nShield HSM menu structure when lockout_mode is set to locked_lt.
panel_poweroff	This controls the function of the Power button on the nShield HSM front panel when it is in operational mode. The default setting of yes enables the Power button. When set to no, the Power button is disabled.

23.5. Sections in both module and client configuration files

23.5.1. server_settings

The `server_settings` section defines the settings for the client hardserver you can modify while the hardserver is running.

> ℹ️ These flags are used by the `NFLOG_DETAIL` environment variable (see Environment variables to control logging).

The section contains the following fields:

Field	Description
`loglevel`	This field specifies the level of logging performed by the hardserver. See hardserver loglevel and Logging, debugging, and diagnostics.
`logdetail`	This field specifies the level of detail logged by the hardserver. You can supply one or more flags in a space-separated list. For more information about the flags, see the table below.
`connect_retry`	This field specifies the number of seconds to wait before retrying a remote connection to a client hardserver. The default is 10.
`connect_maxqueue`	This field specifies the maximum number of jobs which can be queued on the hardserver. The default is 4096: this is also the maximum value. Setting `connect_maxqueue` to a high value allows high throughput, but may cause long latency if the hardserver goes down.
`connect_broken`	This field specifies the number of seconds of inactivity allowed before a connection to a client hardserver is declared broken. The default is 90.
`connect_keepalive`	This field specifies the number of seconds between `keepalive` packets for remote connections to a client hardserver. The default is 10.
`accept_keepidle`	This field specifies the number of seconds before the first keepalive packet for remote incoming connections. The default is 30. Ideally, `accept_keepalive` should be at least twice the value of the `connect_keepalive` setting on the unattended machines.
`accept_keepalive`	This field specifies the number of seconds between keepalive packets for remote incoming connections. The socket will be closed after up to ten consecutive probe failures. The default is 10. Ideally, `accept_keepalive` should be a value such that (10 * `accept_keepalive`) > `connect_broken` on the unattended machine. Using the default values for both these fields will fulfil this requirement.
`connect_command_block`	When the nShield HSM has failed, this field specifies the number of seconds the hardserver should wait before failing commands directed to that HSM with a `NetworkError` message. For commands to have a chance of succeeding after the nShield HSM has failed this value should be greater than that of `connect_retry`. If it is set to `0`, commands to an nShield HSM are failed with `NetworkError` immediately. The default is 35.

Field	Description
max_pci_if_vers	This field specifies the maximum PCI interface version number. If max_pci_if_vers is set to 0 (the default), there is no limit.
enable_remote_mode	If this field is set to yes (the default value) in the module configuration file, nShield HSM mode changing using the nopclearfail utility is enabled. If set to no, mode changing using the nopclearfail utility is disabled. ℹ Do not set enable_remote_mode in the client configuration file.
enable_remote_reboot	If this field is set to yes (the default value) in the module configuration file, the nShield HSM remote reboot using the nethsmadmin utility is enabled. If set to no, remote reboot using the nethsmadmin utility is disabled. Run cfg-push-nethsm to push the new config file to the module.
enable_remote_upgrade	If this field is set to yes (the default value) in the module configuration file, the nShield HSM remote upgrade using the nethsmadmin utility is enabled. If set to no, remote upgrade using the nethsmadmin utility is disabled. Run cfg-push-nethsm to push the new config file to the module.

These flags are those used by the NFLOG_DETAIL environment variable (see Environment variables to control logging).

You can supply a number of flags with the logdetail field, which specifies the level of detail logged by the hardserver (see the table above). Supply the flags in a space separated list:

Flag	Description
external_time	This flag specifies the external time (that is, the time according to your machine's local clock) with the log entry.
external_date	This flag specifies the external date (that is, the date according to your machine's local clock) with the log entry.
external_tid	This flag specifies the external thread ID with the log entry.
external_time_t	This flag specifies the external time_ti (that is, the time in machine clock ticks rather than local time) with the log entry.
stack_backtrace	This flag specifies the stack backtrace with the log entry.
stack_file	This flag specifies the stack file with the log entry.
stack_line	This flag specifies the message line number in the original library. This flag is likely to be most useful in conjunction with example code we have supplied that has been written to take advantage of this flag.
msg_severity	This flag specifies the message severity (a severity level as used by the NFLOG_SEVERITY environment variable) with the log entry.

Flag	Description
msg_categories	This flag specifies the message category (a category as used by the NFLOG_CATEGORIES environment variable) with the log entry.
msg_writeable	This flag specifies message writeables, and extra information that can be written to the log entry, if any such exist.
msg_file	This flag specifies the message file in the original library. This flag is likely to be most useful in conjunction with example code we have supplied that has been written to take advantage of this flag.
msg_line	This flag specifies the message line number in the original library. This flag is likely to be most useful in conjunction with example code we have supplied that has been written to take advantage of this flag.
options_utc	This flag showing the date and time in UTC (Coordinated Universal Time) instead of local time.
unix_file_descriptor_max	This field sets the number of file descriptors the hardserver must be capable of having open concurrently on Linux. The value must be an integer. If unix_file_descriptor_max is set to **0** (the default), the value will be ignored by the hardserver. If it is set to a positive value, the hardserver will refuse to start if the file descriptor hard limit on the system is less than that value. This configuration entry can be used to make sure that the hardserver is capable of satisfying the maximum number of hardserver connections that applications may make use of.

23.5.2. hardserver loglevel

The table in this section describes the loglevels in increasing order of severity. If you set a custom [server_settings]/loglevel, you will get that level and all the more severe levels.

If the legacy NFAST_SERVERLOGLEVEL debug environment variable is set, it overrides any loglevel value set in the configuration file.

[server_settings]/loglevel **value**	Appears in the hardserver log as	Description
info	Information	Report about the hardserver start-up configuration, connections that have been established or closed. General information for nShield Support for debugging.
notice	Notice	Report about certain start-up events, some non-fatal and routine errors that the hardserver can handle internally.
client	Detected error in client behaviour	Malformed or invalid messages are received from a client, typically from a local client.

`[server_set-tings]/loglevel` value	**Appears in the hardserver log as**	**Description**
`remoteserver`	Remote server error	Malformed messages or protocol errors are received while communicating with remote peers over the nCipher Secure Transport/Impath protocol.
`error`	Nonfatal error	Unexpected but handled errors from system calls, for example for device or TCP I/O.
`serious`	Serious error, trying to continue	Unexpected errors from system calls, but they are more serious and likely to indicate an issue somewhere in the system.
`internal`	Serious internal error, trying to continue	Possible bug in the hardserver, but it might also be an issue in the environment the hardserver is running in.
`startup`	Fatal error during startup	The hardserver could not start, for example because of an invalid configuration file or because it cannot bind to a TCP socket which is already in use. The hardserver will abort.
`fatal`	Fatal runtime error	Fatal error usually referring to a non-ignorable error that has occurred after startup such as out of memory errors in certain contexts. Rarely used. The hardserver will abort.
`fatalinternal`	Fatal internal error	A non-recoverable failure occurred, for example certain internal self-consistency checks to detect program logic errors. The hardserver will abort.

23.5.3. server_remotecomms

The `server_remotecomms` section defines the remote IPv4 communication settings for the client hardserver. These are read only at hardserver start-up.

Any changes made to these fields in the HSM config file should be followed by a reboot of the nShield HSM.

It is not recommended that the port number be changed. If it is changed, the port number must match in the configuration of peers. For example, if the port number is changed in the nShield HSM configuration, the remote_port field of the nethsm_imports section of the client-side hardserver config file must be updated to match. The port number can also be specified with the `-P` parameter when running `nethsmenroll`.

This section contains the following fields:

Field	Description
`impath_port`	This field specifies the port on which the hardserver listens for incoming impath connections. The default is 9004. Setting this field to 0 specifies that the hardserver does not listen for incoming connections (do not set to 0 on an nShield HSM). Make sure that firewall settings are consistent with port settings. If you change this field you will need to re-enroll your client with the new port value, see Enrolling the client from the command line.
`impath_addr`	This field specifies the IPv4 address at which the hardserver listens for incoming impath connections. If this field is set to `inaddr_any` (which is the default), the hardserver listens on all IP addresses.
`impath_interface`	This field specifies a string representing the name of the Ethernet interface on which the hardserver listens. This field is only examined if `impath_addr` is set to `inaddr_any`. By default, the `impath_interface` field is empty, which means that the hardserver listens on all interfaces. If the string is not recognized as the name of one of the interfaces of the nShield HSM, the hardserver does not listen.

23.5.4. server_remotecomms_ipv6

The `server_remotecomms_ipv6` section defines the remote IPv6 communication settings for the client hardserver. These are read only at hardserver start-up.

Any changes made to these fields in the HSM config file should be followed by a reboot of the nShield HSM.

It is not recommended that the port number be changed. If it is changed, the port number must match in the configuration of peers. For example, if the port number is changed in the nShield HSM configuration, the remote_port field of the nethsm_imports section of the client-side hardserver config file must be updated to match. The port number can also be specified with the `-P` parameter when running `nethsmenroll`.

This section contains the following fields:

Field	Description
`impath_port`	This field specifies the port on which the hardserver listens for incoming impath connections. The default is 9004. Setting this field to 0 specifies that the hardserver does not listen for incoming connections (do not set to 0 on an nShield HSM). Make sure that firewall settings are consistent with port settings. If you change this field you will need to re-enroll your client with the new port value, see Enrolling the client from the command line.
`impath_addr`	This field specifies the IPv6 address at which the hardserver listens for incoming impath connections. If this field is set to `::` (which is the default), the hardserver listens on all IP addresses.
`impath_interface`	This field specifies a string representing the name of the Ethernet interface on which the hardserver listens. This field is only examined if `impath_addr` is set to `::`. By default, the `impath_interface` field is empty, which means that the hardserver listens on all interfaces. If the string is not recognized as the name of one of the interfaces of the nShield HSM, the hardserver does not listen. Setting this field to `0` will disable IPv6 in the hardserver.

23.5.5. server_startup

The `server_startup` section defines the settings for the hardserver that are loaded at start-up. Any changes you make to the settings in this section do not take effect until after you restart the hardserver. For more information, see Stopping and restarting the client hardserver.

The section contains the following fields:

Field	Description
`unix_socket_name`	This field is not used on Windows. On Linux, this field specifies the name of the socket to use for non-privileged connections. The default is `/dev/nfast/nserver`. If the `NFAST_SERVER` environment variable is set, it overrides any value set for `unix_socket_name` in the hardserver configuration file. For more information about environment variables, see Environment variables.

Field	Description
`unix_privsocket_name`	This field is not used on Windows. On Linux, this field specifies the name of the socket to use for privileged connections. The default is `/dev/nfast/privnserver`. If the `NFAST_PRIVSERVER` environment variable is set, it overrides any value set for `unix_privsocket_name` in the hardserver configuration file.
`nt_pipe_name`	This field specifies the name of the pipe to use for non-privileged connections on Windows. An empty string specifies none. The default is `\\.\pipe\crypto`. If the `NFAST_SERVER` environment variable is set, it overrides any value set for `nt_pipe_name` in the hardserver configuration file. This field is not used on Linux.
`nt_pipe_users`	This field specifies the name of the user or group who is allowed to issue non-privileged connections on Windows. If this field is empty (which is the default), any user can make non-privileged connections. User or group names must be specified unqualified (for example, *bob* not *MYDOMAIN\bob* or *bob@MYDOMAIN*). This field is not used on Linux.
`nt_privpipe_name`	This field specifies the name of the pipe to use for privileged connections on Windows. An empty string specifies none. The default is `\\.\pipe\privcrypto`. If the `NFAST_PRIVSERVER` environment variable is set, it overrides any value set for `nt_privpipe_name` in the hardserver configuration file. This field is not used on Linux.
`nt_privpipe_users`	This field specifies the name of the user or group who is allowed to make privileged connections on Windows. If this field is empty (which is the default), only processes running with local administrator privilege can make privileged connections. User or group names must be specified unqualified (for example, *bob* not *MYDOMAIN\bob* or *bob@MYDOMAIN*). This field is not used on Linux.
`nonpriv_port`	This field specifies the port on which the hardserver listens for local non-privileged TCP connections. The value `0` (which is the default) specifies none. Make sure that your network firewall settings are correct. See the Installation Guide for more about firewall settings. If the `NFAST_SERVER_PORT` environment variable is set, it overrides any value set for `nonpriv_port` in the hardserver configuration file.

Field	Description
`priv_port`	This field specifies the port on which the hardserver listens for local privileged TCP connections. The value `0` (which is the default) specifies none. Java clients default to connecting to port 9001.
	If the `NFAST_SERVER_PRIVPORT` environment variable is set, it overrides any value set for `priv_port` in the hardserver configuration file.

23.5.6. load_seemachine

The `load_seemachine` section of the hardserver configuration file defines SEE machines that the nShield HSMs connected to the client should load and, if required, start for use by other clients. Each SEE machine is defined by the following fields:

Field	Description
`module`	This field specifies the nShield HSM on to which to load the SEE machine. The value must be an integer. A nShield HSM with this ID must be configured on the client computer.
`machine_file`	This field specifies the file name of the SEE machine.
`userdata`	This field specifies the `userdata` file name to pass to the SEE machine on start-up. If this field is blank (`" "`), the SEE machine is loaded but not started. By default, this field is blank.
`worldid_pubname`	This field specifies the `PublishedObject` name to use for publishing the `KeyID` of the started SEE machine. If this field is blank (`" "`), the `KeyID` is not published. This field is ignored if the value of the `userdata` field is blank.
`postload_prog`	This field specifies the program to run after loading the SEE machine in order to perform any initialization required by the SEE machine or its clients. The specified program must accept an argument of the form `-m module#`.
	To run `see-sock-serv` directly on the nShield HSM, set this field to `sockserv`.
`postload_args`	This field specifies arguments to pass to the program specified by the `postload_prog` field. The argument `-m module#` is automatically passed as the first argument. The `postload_args` field is ignored if `postload_prog` is not specified or is blank.
	To run `see-sock-serv` directly on the nShield HSM, set this field to `-p pubname`.

Field	Description
pull_rfs	This field specifies whether the SEE machine name and userdata should be pulled from the RFS. The default is no: set to yes to pull the SEE machine and userdata from the RFS before loading on the remote nShield HSM. This field will be ignored if set on client machine configurations.

> **i** This field will not be added to existing configuration files if you are upgrading an image. If you require the new functionality enabled by this field, you can add the field to the load_seemachine section of your existing configuration file.

23.5.7. slot_imports

The slot_imports section defines slots from remote nShield HSMs that will be available to the client. Each slot is defined by the following fields:

Field	Description
local_esn	This field specifies the ESN of the local nShield HSM importing the slot.
local_slotid	This field specifies the SlotID to use to refer to the slot when it is imported on the local nShield HSM. The default is 2.
remote_ip	This field specifies the IP address of the machine that hosts the slot to import.
remote_port	This field specifies the port for connecting to the nShield HSM.
remote_esn	This field specifies the ESN of the remote nShield HSM from which to import the slot.
remote_slotid	This field specifies the SlotID of the slot to import on the remote nShield HSM. The value of this field must be an integer. The default is 0.

23.5.8. slot_exports

The slot_exports section defines the slots on the HSMs connected directly to the client that the client hardserver should export for other network HSMs to import. Each local slot has an entry for each network nShield HSM that can import it, consisting of the following fields:

Field	Description
local_esn	This field specifies the ESN of the local nShield HSM whose slot can be imported by a network nShield HSM.

Field	Description
local_slotid	This field specifies the SlotID of the slot that is allowed to be exported. The value must be an integer. The default is 0.
remote_ip	This field specifies the IP address of the nShield HSM that is allowed to import the slot. Keep this field blank to allow all machines. The default is blank.
remote_esn	This field specifies the ESN of the nShield HSM allowed to import the slot. Leave the value blank to allow all permitted nShield HSMs in the Security World. The default is blank.

23.5.9. dynamic_slots

The dynamic_slots section defines the number of Dynamic Slots that each HSM supports.

Field	Description
esn	ESN of the HSM to be configured with Dynamic Slots.
slotcount	The number of Dynamic Slots that the HSM is to support. If set to 0 (default) the HSM does not support Remote Administration.

23.5.10. slot_mapping

The slot_mapping section defines, for each specified HSM, a slot that is exchanged with slot 0 of the HSM. Slot 0 becomes a Dynamic Slot and the local slot becomes the specified slot number. This enables applications and utilities that only support slot 0 to use Remote Administration.

Field	Description
esn	ESN of the HSM to which the mapping is applied.
slot	The slot number to be swapped with slot 0, so that: • Slot 0 refers to a Dynamic Slot • The specified slot number refers to the local slot of the HSM. If slot is set to 0 (default) there is no slot mapping.

23.5.11. dynamic_slot_timeouts

The dynamic_slot_timeouts section defines timeout values that are used to specify expected smartcard responsiveness for all HSMs associated with the relevant host or client, when using the Remote Administration.

Field	Description
`round_trip_time_limit`	`round_trip_time_limit` > 5s + network latency time Round trip (HSM to smartcard and back) time limit in seconds. The card is regarded as removed, if no response has been received within the allowed time. Expected network delays need to be taken into account when setting this. The default is ten seconds.
`card_remove_detect_-time_limit`	`card_remove_detect_time_limit` >= `round_trip_time_limit` *2 Maximum number of seconds that can pass without a response from the smart card, before it is regarded as removed and all the keys that it protects are unloaded. Lower values increase network traffic. The default is 30 seconds.

23.6. Sections only in client configuration files

23.6.1. module_settings

The `module_settings` section defines the settings for the nShield HSM that can be changed while the hardserver is running. The section contains the following fields:

Field	Description
`esn`	This field specifies the electronic serial number of the nShield HSM.
`priority`	This field specifies the priority of the nShield HSM. The value for this field can be an integer from 1 (highest) to 100 (lowest). The default is 100.

23.6.2. server_performance

The `server_performance` section defines the performance settings for the client hardserver. These are read only at hardserver start-up. This section contains the following fields:

Field	Description
`enable_scaling`	This field determines whether multi-threaded performance scaling is enabled or not. If this field is set to **auto** (or not set), the hardserver automatically chooses the best option for the available hardware (enabled when using an nShield network-attached HSM, for which scaling is currently optimized, and disabled if using an nShield PCIe or USB-attached HSM). It can explicitly be enabled by setting to **yes**, and explicitly disabled by setting to **no**.

Field	Description
target_concurrency	This field allows the level of concurrency to be tuned. The value must be an integer and will only come into effect when multi-threaded performance scaling is enabled. If target_concurrency is set to **0** (the default), the value will be automatically configured by the hardserver based on the available number of physical CPU cores. The target concurrency configured is written to the hardserver log.

23.6.3. nethsm_imports

The `nethsm_imports` section defines the network nShield HSMs that the client imports. It can also be set up by the `nethsmenroll` utility. Each nShield HSM is defined by the following fields:

Field	Description
local_module	This field specifies the ModuleID to assign to the imported nShield HSM. The value must be an integer. An nShield HSM with this ID must not be already configured on the client computer.
remote_ip	This field specifies the IP address of the nShield HSM to import.
remote_port	This field specifies the port for connecting to the nShield HSM.
remote_esn	This field specifies the ESN of the imported nShield HSM.
keyhash	This field specifies the hash of the key that the nShield HSM should use to authenticate itself.
privileged	The value in this field specifies whether the client can make a privileged connection to the nShield HSM. The default is **0**, which specifies no privileged connections. Any other value specifies privileged connections.
ntoken_esn	This field specifies the ESN of this client's nToken, if an nToken is installed.

> ! The default value for `remote_keyhash` (40 zeros) specifies that no authentication should occur. We recommend that you set a specific key hash in place of this default.

23.6.4. rfs_sync_client

This section defines which remote file system the client should use to synchronize its key management data:

Field	Description
remote_ip	The IP address of the RFS against which to synchronize.
remote_port	This field specifies the port for connecting to the RFS.
use_kneti	Setting this option to yes to use a local module KNETI instead of the default hardserver's software KNETI to authenticate this client to the RFS.
local_esn	This is only required if use_kneti is set to yes. It is the ESN of the local module used for authentication.
remote_keyhash	Software or module KNETI hash used to authenticate the RFS, or 40 zeroes to indicate no authentication required (default is 40 zeroes).
remote_esn	ESN of the remote module used to authenticate the RFS, or empty when using software KNETI authentication or no authentication required (default is empty).

23.6.5. remote_file_system

> ❗ This section is updated automatically when the rfs-setup utility is run. Do not edit it manually.

The remote_file_system section defines a remote file system on the client by listing the nShield HSMs allowed to access the file system on this client. Each nShield HSM is defined by an entry consisting of the following fields:

Field	Description
remote_ip	This field specifies the IP address of the remote nShield HSM that is allowed to access the file system on this client.
remote_esn	This field specifies the ESN of the remote nShield HSM allowed to access the file system on this client.
keyhash	This field specifies the hash of the key with which the client must authenticate itself to the nShield HSM. The default is 40 zeros, which means that no key authentication is required.
native_path	This field specifies the local file name for the volume to which this entry corresponds.
volume	This field specifies the volume that the remote host would access to use this entry.
allow_read	If this field is set to yes, it means that a remote server is allowed to read the contents of the file. The default is no.
allow_write	If this field is set to yes, it means that a remote server is allowed to write to the file. The default is no.

Field	Description
`allow_list`	If this field is set to `yes`, it means that a remote server is allowed to list the contents of the file. The default is `no`.
`is_directory`	If this field is set to `yes`, it means that this entry represents a directory. The default is `no`.
`is_text`	If this field is set to `yes`, it means that line endings should be converted to and from the Linux convention for transfers.

> **ℹ** If you upgrade from an earlier software version to v12 and are using Remote Administration, you need to manually add the following sections to your configuration file.

23.6.6. remote_administration_slot_server_startup

The `remote_administration_slot_server_startup` section defines the communication settings that are applied at start-up to the Remote Administration.

Field	Description
`port`	Which port to use to connect to the Remote Administration. The default is 9005.

23.6.7. hs_clients

The `hs_clients` section defines the clients that are allowed to connect to the nShield HSM. Each client is defined by an entry as follows:

Field	Description
`addr`	Either the IP address of the client or `0.0.0.0`, `::`, or blank if the HSM is to accept clients identified by their keyhash instead of their IP address.
	If no value is set (the field is blank), or if it is set to `0.0.0.0` or `::`, only `HKNETI` identification is allowed.
	The default is blank.
	`0.0.0.0` or `::`, and blank result in the same behavior. You cannot enter these values in the front-panel user interface. You can only use them in the configuration file and you will have to use the correct key hash for identification if no IP address is configured.

Field	Description
`clientperm`	The type of connection permitted from the client. This is one of the following: • `unpriv` (non-privileged connections) • `priv` (privileged connections) • `priv_lowport` (privileged connections on ports lower than 1024) The default is `unpriv`.
`keyhash`	The hash of the KNETI key that the client should present to authenticate itself. Both software based authentication and nToken authentication are supported, For nToken authentication, a value must also be provided for the `esn` field. The default is 40 zeros, which means that no key authentication is required.
`esn`	The ESN of the client's nToken. (Only applicable if nToken authentication is used.)

24. Cryptographic algorithms

24.1. Symmetric algorithms

Symmetric Algorithms				
Algorithm	FIPS approved in a v1 or v2 Security World	FIPS approved in a v3 Security World	Key type	Supported by generatekey
AES	Y	Y	AES or Rijndael	Y
Arcfour	N	N	Arcfour	N
ARIA	N	N	Aria	N
Camellia	N	N	Camellia	N
CAST 256	N	N	CAST256	N
DES	N	N	DES	N
DES2	Y	N	DES2	Y
Triple DES	Y	N [1]	Triple DES	Y
MD5 HMAC	N	N	HMACMD5	N
RIPEMD160 HMAC	N	N	HMACRIPEMD160	N
SEED	N	N	SEED	N
SHA-1 HMAC	Y	Y	HMACSHA1	Y
SHA-224 HMAC	Y	Y	HMACSHA224	N
SHA-256 HMAC	Y	Y	HMACSHA256	Y
SHA-384 HMAC	Y	Y	HMACSHA384	Y
SHA-512 HMAC	Y	Y	HMACSHA512	Y
Tiger HMAC	N	N	HMACTiger	N

[1] Not FIPS 140 approved for encryption operations, but available for decryption operations.

24.2. Asymmetric algorithms

Asymmetric Algorithms

Algorithm	FIPS approved in a v1 or v2 Security World	FIPS approved in a v3 Security World [1]	Key type	Supported by generatekey
Diffie-Hellman	Y	Y	DH or DHEx	Y
DSA	Y	Y	DSA	Y
ECDH	Y [2]	Y [2]	ECDH or EC [3]	Y
ECDSA	Y [4]	Y [4]	ECDSA or EC	Y
ECIES	N	N	ECDH or EC	N
Ed25519	N	N	Ed25519	N
ElGamal	Y	Y	DH	Y
KCDSA	N	N	KCDSA	N
RSA	Y	Y	RSA	Y
X25519	N	N	X25519	N

[1] Some insecure key sizes are non-FIPS 140 approved.

[2] FIPS 140 approval is only for use with ECDH keys, not with EC keys, but see [3] for exception.

[3] FIPS 140 allows an EC key to be used as an ECDH key for a single use-case. In this use case, an ECDH key is allowed to perform a single signing of a Certificate Signing Request (CSR), so that a certificate for the ECDH key can be generated.

[4] FIPS 140 approval is only for use with ECDSA keys, not with EC keys.

24.3. FIPS information

In a FIPS 140 Level 3 Security World, the nShield HSM only supports FIPS-approved algorithms and key sizes.

- If you have a FIPS 140 Level 3 Security World and have any protocols that use algorithms not approved by FIPS, you have the following options:
  - If you need to use these non-approved algorithms, you can migrate to a non-FIPS Security World but continue to use hardware and firmware validated for FIPS 140 Level 3.
  - If you have strict FIPS 140 Level 3 requirements, you must replace your protocols to use approved algorithms.

- If you have a FIPS 140 Level 3 Security World and have existing long-term keys for unapproved algorithms, you have the following options:

  - Migrate to a non-FIPS Security World but continue to use hardware and firmware validated for FIPS 140 Level 3.

  - Replace the keys with approved keys before upgrading to the current firmware. Keys for unapproved algorithms are incompatible with this Security World.

To obtain more details on the specific algorithms that are FIPS approved for use in the HSM, refer to the nShield Security Policy for the particular FIPS CMVP certified nShield product that you are using.

For the FIPS CMVP certificates for nShield products, see https://csrc.nist.gov/projects/cryptographic-module-validation-program/validated-modules/search. The FIPS CMVP certificate links to the Security Policy.

24.4. Compatibility of Security World versions with FIPS

To comply with the latest FIPS cryptographic transitions, Security World v3 was introduced in firmware version 12.50. If an nShield HSM is upgraded to use firmware version 12.50 or later, any v2 Security Worlds using the HSM that were compliant with FIPS 140 Level 3 will no longer be compliant.

You can create a v3 Security World that is compliant with FIPS 140 Level 3 from a host server if you meet the following criteria:

- The host server is running Security World host-side software version 12.50 or later.
- The HSM is running firmware version 12.50 or later.

Your solution is only FIPS 140 compliant if you are running the exact firmware version that has been FIPS 140 certified.

25. Audit Logging

Audit Logging on nShield HSMs provides the following features:

- Logs generated and signed on the nShield HSM
- Tamper detection
- Deletion Detection
- Administrative operations are logged
- Key lifetime events are logged
- Per key usage events are optionally logged
- Optional key usage logging
- Public key verification of audit logs
- Compatibility with syslog and Security Information and Event Management (SIEM).

25.1. Configuring Audit Logging

Audit Logging is enabled per Security World and is configured on creation of the Security World.

Prerequisites

- If the audit logs are to be sent to a non-default location, the configuration file must be set up before the Security World is created.
- The Real Time Clock (RTC) on the HSM must be set and the setting confirmed after cre ating the Security World or indoctrinating an HSM into the Security World. The RTC on the HSM is used to timestamp audit log entries.

25.1.1. Configure audit log transport through syslog

The Audit Logging entries are delivered over syslog using UDP transport. This transport must be configured before Audit Logging is enabled in order to collect the initial messages.

1. Check the syslog transport before creating an Audit Logging enabled Security World.

   Send a log message to the configured host and port using a command, for example **log ger**, that can send messages to a syslog server over UDP.

2. Set the Audit Log entries in the hardserver configuration file.

   ```
   [auditlog_settings]
   # Start of the auditlog_settings section
   ```

```
# Hardserver settings for audit logging.
# Each entry has the following fields:
#
# The port number Auditlogging server listens to .
auditlog_port=514
#
# IP Address of the Auditlogging server
auditlog_addr=WWW.XXX.YYY.ZZZ
#
# Copy auditlog to hardserver log. (default=no)
#  auditlog_copy_hslog=ENUM
```

<table>
<tr>
<td><code>auditlog_port</code></td>
<td>This is the UDP destination port for Audit Logging syslog messages. The default is 514.</td>
</tr>
<tr>
<td><code>auditlog_addr</code></td>
<td>This is the IP address of the host to which the Audit Logging syslog messages should be delivered. The default is 0.0.0.0.</td>
</tr>
<tr>
<td><code>auditlog_copy_hslog</code></td>
<td>This indicates that the syslog messages from Audit logging should be copied to the hardserver's log file. This provides some degree of redundancy but means that the hardserver's log file may grow significantly. The default is no. It is not recommended on nShield network-attached HSMs.</td>
</tr>
</table>

3. Restart the hardserver to load the updated configuration file.

4. Push the configuration to the nShield HSMs in the Security World and restart the hardserver on each HSM. See Utilities for network-attached HSMs.

25.1.2. Create a Security World with Audit Logging enabled

For the overall procedure, see Creating a Security World using new-world.

Key considerations:

- A Security World is created with Audit Logging enabled if either the `--audit-logging` or `-G` options are passed to the `new-world` command or the Security World is created in the `common-criteria-cmts` mode. This requires that the HSM is capable of audit logging and Security World creation will fail if the HSM does not support Audit Logging.

- Additional HSMs are indoctrinated into an Audit Logging enabled Security World using the `new-world` command with the `--program` or `-l` options.

- The HSM must be capable of Audit Logging. If it is not capable the indoctrination will fail. Therefore all HSMs in an Audit Logging Security World are set to Audit Logging.

When you configure an Audit Logging Security World:

1. Audit Logging is set as enabled for this Security World and is recorded in the world file.

2. The HSM is initialized and:

   ° A flag indicating the Audit Logging status is recorded in the HSM's EEPROM

   ° A 3072bit DSA HSM specific Audit Logging Signing Key (KAL) is created and persisted in the HSM's EEPROM

   ° A Certifier Block containing the public half of KAL is generated and sent to the log receiver via the hardserver.

When you indoctrinate a new HSM into an Audit Logging Security World:

- The world file specifies that this is an Audit Logging Security World

- The same actions as for initializing an HSM when the Audit Logging Security World was created are performed. This means that:

   ° All HSMs in an Audit Logging Security World have Audit Logging enabled

   ° Each HSM has a distinct Audit Logging Signing Key.

25.1.3. Confirm the Audit Logging configuration

Check for `AuditLogging` on the `state` line in the output of `nfkminfo`.

Enabled	`AuditLogging`
Disabled	`!AuditLogging`

```
>nfkminfo
...
 state       0x37270009 Initialised Usable Recovery !PINRecovery !ExistingClient RTC NVRAM FTO
AlwaysUseStrongPrimes !DisablePKCS1Padding !PpStrengthCheck AuditLogging SEEDebug
...
```

Check `AuditLogging` on the `active modes` line in the output of `enquiry`.

Enabled	`AuditLogging`
Disabled	(AuditLogging is not listed)

```
>enquiry
...
mode            operational
...
active modes    AuditLogging UseFIPSApprovedInternalMechanisms AlwaysUseStrongPrimes
hardware status OK
```

25.1.4. Disable Audit Logging

Audit Logging is set for the lifetime of the Security World.

To disable Audit Logging on an HSM:

1. Remove that HSM from the Security World.

2. Reinitialize the HSM using `initunit`.

25.2. Audit Logging architecture

Audit Logging is implemented on the nShield HSM. The Audit Logging capability is on the embedded nShield 5s card. The following image displays the nShield HSM implementation.

The audit log entries are generated on the module, the hardserver acts as a relay to a syslog infrastructure. The logging is at the level of nCore commands as processed by the HSM.
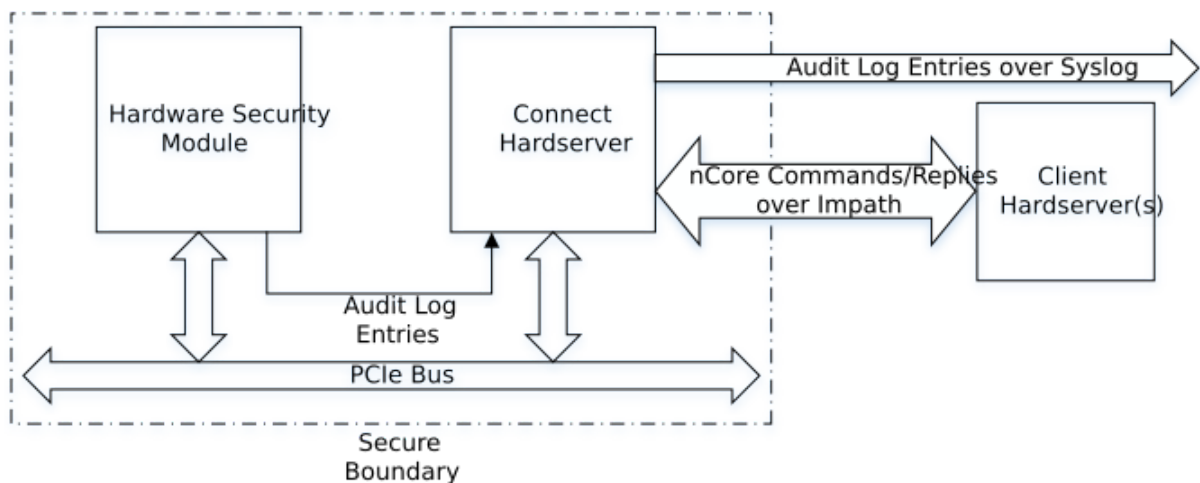
The hardserver layer implements a higher-level abstraction which is presented to application clients. The information used to manage this such as Client Identifiers etc. is not available to the HSM and therefore cannot be logged.

25.2.1. Audit Logging implementation

The Audit logging functionality is based on that described in RFC-5848 (https://tools.ietf.org/html/rfc5848). This describes a mechanism also known as syslog-sign that adds the following capabilities to syslog:

- Origin authentication

- Public verification

- Message integrity

- Replay resistance

- Message sequencing
- Detection of missing messages.

It is implemented on top of a standard syslog data stream and does not use any additional protocol. The syslog-sign logging scheme adds a number of control messages to the log entries that are to be audited. These are also implemented as syslog messages. The following sections outline the audit log entries that are present in the syslog data stream for nShield Audit Logging. The signing mechanism used is DSA with a 3072 bit key and SHA-256 as the hashing mechanism.

Audit log entry

This is the log message from the entity being audited. It is in a standard format and includes operation specific data required to provide an auditing capability. As each log message is generated on the HSM, a digest operation is performed on it and the digest buffered in the HSM.

Signature Block

When sufficient digests have been accumulated (N), a Signature Block is generated as a standard log entry containing the following:

- Digests of the previous N messages
- Information to allow the digests to be matched with their respective log entries
- A signature across the digests and other information.

The number of messages is dictated by the transport medium, the size of the digests, the size of the signature and the size of other data contained in the message. There is a limit to the size of messages that can be transported over syslog. The signature is performed using a log signing key. This key is generated and the private half is held securely in the HSM.

Certifier Block

Verification of the Signature Blocks requires that you, or the application performing the verification, has access to the public half of the log signing key. The Certifier Block provides a mechanism for the log verifier to get access to this key. The key is packaged in a form allowing the source of the audit logs to be verified. As the size of this information may be too large for the syslog transport medium it can be broken down into Certifier Block Fragments which are compatible with the syslog transport mechanism. When all of these fragments are received by the log verifier, it can reconstruct the public half of the log signing key and perform any consistency checks and origin verification that is needed.

25.2.2. Audit Log Verification process

Given the public half of the log signing key, a Signature block and its corresponding log
entries, the verifier can check the signature on the Signature Block. When this is verified,
the log entry digests in the Signature Block are implicitly verified. The integrity of the corre-
sponding log entries can be verified by performing a digest on received log entries and com
paring them to the corresponding verified digests in the Signature block. The image below
shows the basics of this approach. For more information, see Audit Log Verification.

Audit Log Generation   Audit Log Verification
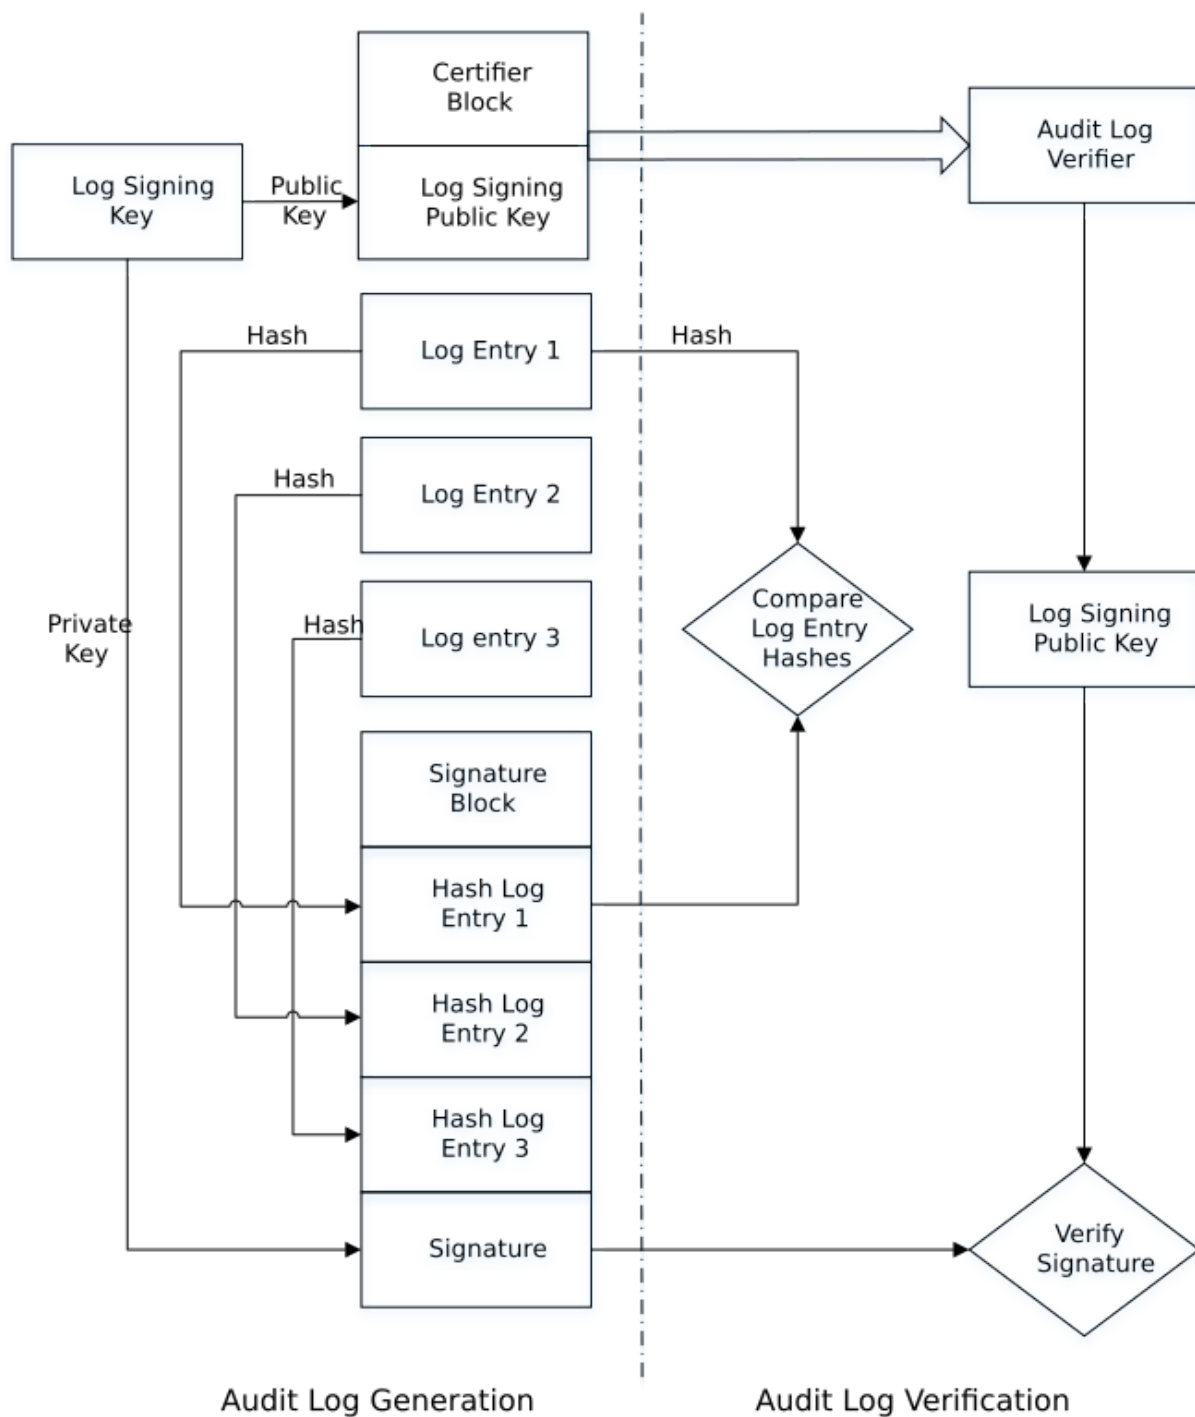
25.2.3. Log distribution

The nShield Audit Logging capability uses the RFC-3164 (https://tools.ietf.org/html/rfc3164) standard for distributing audit log messages. All audit log messages will have the following header prepended. This header is applied by the hardserver before sending the message and does not form part of the signed audit log messages. The signed portion of the audit log message starts at the `CEF:0` CEF identifier and continues to the end of the message.

```
<134>MMM DD HH:MM:SS hostname CEF:0......
```

The PRI value of this header `<134>` indicates the facility `local0` and a severity of `info`.

The syslog infrastructure used for Log distribution is out of the scope of this guide and your responsibility to implement. Log distribution for Audit Logging uses syslog as the transport medium which allows a standard protocol and message format to be used for the Audit Logging messages. This transport is compatible with SIEM systems and the wider syslog infrastructure in an organization can be used to further distribute or process the log stream. There are many possible configurations of syslog deployment, as shown below.

25.3. Configuring audit log distribution

The actual implementation of the syslog infrastructure is at your discretion. Verification of the log messages requires that the verifying application has access to the audit logs from the HSMs in the Security World. The example verifier for the nShield Audit logging facility described in Audit Log Verification processes a file containing the audit log messages. It can process audit log messages from a specific HSM identified by its ESN or will use the

first ESN found in the file.

It is recommended that logs from the nShield Audit Logging facility are separated from those from other applications. This can be accomplished by using the information in the audit log messages described in the section on Log Format. There are a number of entries that can be used to separate out the messages from the nShield Audit Logging facility. These include:

- Identifying elements in the CEF header:
    - Device Vendor
    - Device Product
- Identifying elements in the syslog header
    - Hostname or IP address of the machine hosting the hardserver which is distributing the audit log messages
- Using a distinct port for nShield Audit Logging see configuring syslog.

The log messages can be further split into those from specific HSMs using the ESN in the audit log messages.

As an example, the following `rsyslog` configuration will direct all messages with the string *nCipher Security* to a specific log file:

```
:msg, contains, "nCipher Security" /var/log/hsmauditlog
```

A similar strategy can be used with `syslog-ng`:

```
destination d_auditlog { file("/var/log/hsmauditlog"); };
filter f_auditlog { match("nCipher Security" value("MESSAGE")); };
log { source(s_log); filter(f_auditlog); destination(d_auditlog); }
```

Adjust the example paths for the platform running your syslog server as appropriate. `s_log` is the source receiving the audit logging messages.

Refer to the documentation for your syslog implementation for information on processing and distributing log messages.

25.4. Configuring the syslog message infrastructure

It is important that the syslog infrastructure does not attempt to rewrite the log messages as this will affect the ability of the Audit Logging process to verify log messages. For example, the `rsyslog` default RFC-3164 parser will rewrite log messages interpreting the **CEF:0** as a tag and will write **CEF: 0** to the log file. This means that an Audit Logging message per-

sisted by the default RFC-3164 parser cannot be verified as Audit Logging signs the log message starting at **CEF:**. You must configure your syslog infrastructure to preclude the signed part of the audit log message.

25.4.1. rsyslog

rsyslog can be configured to not reformat messages using the following approach:

1. Define a formatting template as shown below in the `/etc/rsyslog.conf` file. This should be added in the `##### MODULES #####` section of the rsyslog configuration file.

   ```
   $template myFormat,"%rawmsg%\n"
   ```

2. Apply this formatting template to the processing of Audit Logging messages. For this example it is assumed that messages containing *nCipher Security* will be persisted in the `/var/log/hsmauditlog` file. You can use any other selection mechanism such as storing messages for a particular HSM as identified by its ESN in separate files.

   Linux

   ```
   :msg, contains, "nCipher Security"          /var/log/hsmauditlog;myFormat
   ```

   Windows

   ```
   :msg, contains, "nCipher Security"          \var\log\hsmauditlog;myFormat
   ```

3. If the rsyslog server is going to be used as a relay, then the format needs to be applied to any relay statements in the rsyslog configuration file and to any receivers of the syslog message.

25.4.2. syslog-ng

`syslog-ng` does not appear to rewrite messages in the same way as `rsyslog`. Refer to the *syslog-ng* documentation for information on formatting.

25.5. Audit log format

25.5.1. CEF format

The audit log entries are emitted from the HSM in CEF format. This provides both human

readable log messages and compatibility with SIEM applications. As indicated in the previous section the audit log entries are distributed using the syslog transport mechanism.

A CEF format log message is shown below:

```
CEF:Version|Device Vendor|Device Product|DeviceVersion|Device Event Class ID|Name|Severity|[Extension]
```

Parameter	Description
`CEF:Version`	This is mandatory and Version is currently **0.**
`Device Vendor`	This is nCipher Security
`Device Product`	This identifies the family of nShield HSMs: • nShield Solo • nShield Solo XC • nShield Edge • nShield 5s
`Device Version`	This is the firmware version, for example *12.80.0*.
`Device Event Class ID`	This is an identifier for the type of message: <table><tr><th>Class ID</th><th>Description</th></tr><tr><td>1</td><td>nCore Commands</td></tr><tr><td>2</td><td>Internal HSM events: • Periodic heartbeat • Secure channel establishment</td></tr><tr><td>3</td><td>Audit logging control messages: • Signature Blocks • Certifier Blocks</td></tr><tr><td>4</td><td>Reserved</td></tr><tr><td>5</td><td>Shutdown messages</td></tr><tr><td>6</td><td>Reserved</td></tr></table>
`Name`	This is the event being logged. For Audit Logging, it is either the nCore command that is being logged, `Cmd_Destroy` for example, a description of the event such as `heartbeat` or one of `ssign` or `ssign-cert` which identifies either a Signature Block or a Certifier block.

Parameter	Description
Severity	This is an indication of the importance of the message.

Severity	Description
1	nCore Commands
2	Internal HSM events: • Reboot events • Secure channel establishment
3	nCore Commands that force a Signature Block flushing buffered message hashes.
4	Periodic Heartbeat messages
5	Audit Logging control messages: • Signature Blocks • Certifier Blocks
6	Shutdown messages
10	HSM Error messages

25.5.2. CEF extensions

The rest of the log message is made up of CEF extensions. These are name/value pairs that are used to convey specific information for the log message. The name-value pairs can be processed by SIEM applications such as Arcsight and can be displayed in tabular reports of the messages received. They can be used for filtering and further processing within the SIEM application. The following table specifies the meaning and format of the extensions used by the Audit Logging facility.

Extension Name	Description
esn	Electronic Serial Number (ESN) of the HSM in the following format: XXXX-XXXX-XXXX
rtc	Time-stamp from the HSM's Real Time Clock (RTC) as ms since the epoch (1970 Jan 01 00:00:00 UTC).
outcome	Outcome of the operation - **success** or **failure**
hkey	Identifying nCore key hash for the main key of the command being logged as a 40 character hex string

Extension Name	Description
hbase	Identifying nCore key hash for the base key of a `Cmd_DeriveKey` command being logged as a hex string
hwrap	Identifying nCore key hash for the wrap key of a `Cmd_DeriveKey` command being logged or the logical token hash for key blobbing operations as a hex string
hin3-5	Identifying nCore key hashes for the remaining keys of a `Cmd_DeriveKey` command being logged as hex strings
hknso	Identifying nCore key hash of Security Officer's key as a hex string
htok	Identifying nCore Logical Token hash as a hex string
shareindex	Index of share being operated on by Logical Token functions as a decimal number
sharesleft	Number of Logical Token shares left to read or write as a decimal number
tokenslot	Slot number for Logical Token operations as a decimal number
sharesneeded	Quorum required to reconstruct a Logical Token as a decimal number
sharestotal	Total number of shares for a Logical Token as a decimal number
timelimit	How many seconds after reassembly the Logical Token is usable for
shorthash	Short Logical Token hash used in `Cmd_EraseShare` and `Cmd_ChangeSharePIN`
hkm	Identifying nCore hash of module key KM
mode	Mode that a channel is opened in. One of **encrypt**, **decrypt**, **sign** or **verify**
source	Source of command. One of **host**, **SEE** or **internal**
flags	Flags supplied to `Cmd_SetNSOPerms` and `Cmd_InitialiseUnitEx`. `Cmd_SetNSOPerms` - AlwaysUseStrongPrimes - DisablePKCS1Padding - FIPSLevel3Enforcedv2 - CommonCriteriaCMTSRestrictions `Cmd_InitialiseUnitEx` - AuditLogging - UseFIPSApprovedInternalMechanisms
slotcount	Count of Dynamic Slots to be configured
slotid	Dynamic Slot to create association for

Extension Name	Description
prevrtc	The previous value of the HSM's RTC as ms since the epoch. Used to indicate previous value of the RTC before a `Cmd_SetRTC` timestamp or an event occurring before a restart
smartcardesn	ESN of smartcard used for Dynamic Slot operations
kmltype	Type of the Module Per-Initialization Signing Key (KML) set by `Cmd_Initialise Unit(Ex)`
sos	Indication of the sos code

25.5.3. Infrastructure extensions

The Audit Logging Implementation requires a number of infrastructure CEF extensions to provide data necessary for the RFC-5848 based signed syslog approach used. Please refer to RFC-5848 for further details on these infrastructure extensions. These CEF extensions replace the RFC-5424 Structured Data used in the original scheme but have the same meaning.

25.5.4. Message and reboot counters

There are two counters that are sent with all Audit Logging command log messages. The Reboot Session ID is also sent with Certifier Block and Signature Block messages.

Counter	Description
seqNo	Log message sequence number as a decimal number. This a counter that has a range of 1 to 9999999999. When seqNo reaches 9999999999 it is reset to 1. It is incremented for every command log message sent. This is not part of RFC-5848 and has been added to provide a direct mechanism for detecting deleted or missing log messages. The sequence number is unique in the context of a reboot session. When `rsid` is incremented `seqNo` is reset to **1**.
rsid	Reboot Session ID as a decimal number. This a counter that has a range of 1 to 9999999999. When rsid reaches 9999999999 it is reset to 1. It is incremented every time the HSM is restarted and whenever the Global Block Counter (gbc) reaches its limit and is reset.

25.5.5. Client ID Session Extension

The `session` extension is a unique identifier that refers to an HSM's client. This allows audit logs to be traced to the client machine that initiated a command for the HSM.

When a client is enrolled into the HSM, the following takes place:

- A session value is calculated from both the client's IP address and its KNETI hash.
- The HSM automatically generates an audit log with the `Cmd_SessionCreate` command that contains the IP address of the client, the client's KNETI hash and the corresponding session ID.

  For example, for an nShield 5c HSM:

```
CEF:0|nCipher Security|nShield5|13.4.0|1|Cmd_SessionCreate|1|esn=F4A9-050A-EBF8 rsid=5 rtc=1678208481435
seqNo=215 source=host outcome=success
description=IP:"172.23.136.106:512";KNETI:"76c5516f834376eb34e7345fa3870fde764c962d"
session=75c9cd2600000001
```

Subsequently, whenever the client executes a loggable command on the HSM, the audit log that's generated will contain the `session` field with the client's unique identifier.

If this session is terminated (for example, if the Impath Resilience session expires after a client disconnects and doesn't resume the session before the timeout), then the HSM will generate an audit log with the `Cmd_SessionDestroy` command to indicate this.

For the same HSM:

```
CEF:0|nCipher Security|nShield5|13.4.0|1|Cmd_SessionDestroy|1|esn=F4A9-050A-EBF8 rsid=5 rtc=1678208481437
seqNo=216 source=host outcome=success session=75c9cd2600000001
```

Name	Description
session	A unique identifier for the client that executed a command on an HSM. This field is only available on audit logs generated inside an HSM for commands initiated by the HSM's clients.

25.5.6. Certifier Block extensions

The following extensions are used in the Certifier Block where the CEF header name element is `ssign-cert`. The Certifier Block is used to distribute the log signing public key. It is sent by the HSM when logging is enabled and every time the HSM is restarted. This provides for redundancy as any Certifier Block from an HSM that has been configured for Audit Logging will contain the same log signing public key. The Certifier Block can extend over multiple syslog messages. The extensions identified here allow the data of a Certifier Block to be rebuilt from multiple fragments. Sufficient fragments are sent in separate `ssign-cert` messages to rebuild the payload block. See Certifier Block example for the details of the data included in the certifier block.

Name	Description
tpbl	Total Payload Block Length. This is the length of all Certifier Block fragments.
findex	Index of this fragment (1 based) as a decimal number.
flen	Length of the fragment as a decimal number.
frag	Base 64 encoded Certifier Block fragment.
sign	DSA signature using KAL of the data in each fragment up to the sign extension. The DSA signature is DER encoded and then base64 encoded. It is present here to support consistency checking.

25.5.7. Signature Block extensions

The following extensions are used in the Signature Block where the CEF header name element is ssign. The Signature Block supports the verification of Audit Logging messages. The Signature Block is sized to fit within a syslog message which dictates the number of audit log messages it covers. This release supports a maximum of 10 audit log messages per Signature Block. The main data for the Signature Block is the SHA-256 hashes of the log messages covered by the block.

Name	Description
gbc	Global Block Counter as a decimal number. Count of signature blocks sent in this Reboot Session prior to this Signature Block. This is a decimal number that has a range of 1 to 9999999999. When gbc reaches 9999999999 it is reset to 0. At the same time the rsid is incremented and seqNo reset to 1. It is incremented after every Signature Block is sent.
fmn	First Message Number as a decimal number. First log message seqNo in this Signature Block.
hcnt	Count of log messages included in this Signature Block as a decimal number.
hb	The log message SHA-256 hashes base64 encoded and separated by the & character. The & character does not occur in base64 encoding and avoids SIEM issues with embedded spaces.
sign	DSA signature using KAL of the data in the Signature Block up to the sign extension. The DSA signature is DER encoded and then base64 encoded.

25.5.8. Example Audit Logging messages

This section shows example Certifier Block, Signature Block and Audit Logging messages and shows how the CEF extensions are used together.

25.5.8.1. Certifier Block example

This is an example Certifier Block produced after a reboot of the HSM. The log messages have been reformatted for display as each one can be up to 1024 characters long. The Reboot Session ID (rsid) is 8. There are five fragments in this example. The first four are 450 characters and the final 340 long for a total length of the payload of 2140 characters. The Event Class Id is 3 and the severity is 5 identifying these as infrastructure messages.

```
<14>Aug 10 18:22:18 nethsm CEF:0|nCipher Security|nShield Solo XC|12.60.9|3|ssign-cert|5|esn=7109-02E0-D947
rsid=9 rtc=1628616065737 tpbl=2140 findex=1 flen=450
frag=uwAAAEQAAAArPjRMy6QnKObhNt0va3OSSdLSTSEkO2Ydxj2sRmsTvmDJ5tqXeNE0p7GkmFfgO/yU2iTl9VlSJa7Opq+BKNTw3wEAAEQAAACI
WksBX/JuSf8Cyv3WRfnUWvvCAHID+OGaoYvHqqlVtW7GzRMWMg/fPeO3UV32w2i+GkORU87i6dJKybxqhwVVQwAAAKwFAAAEAAAAAAAAAMAAAACA
AAADwAAADcxMDktMDJFMC1EOTQ3AAAANAAAAa0OXfBuCMHaS7g0dVUE5vnVx43guAAAABgAAAAAAAABEAAAA4P0MEe9CDHohvBeOvBwLykF99mopVY
1Rdam8554PDJWwhwID6+tUn8AJGOAscYtfrOxabeJ1iC+btrmJZbrlrXQBAABEAAAAdlxXLbwppzAjT8/H+off7ikbnLNQdozuNJeNaqb4+IZwCxc
NwA

<14>Aug 10 18:22:18 nethsm CEF:0|nCipher Security|nShield Solo XC|12.60.9|3|ssign-cert|5|esn=7109-02E0-D947
rsid=9 rtc=1628616065739 tpbl=2140 findex=2 flen=450
frag=s3z1d9CeG3/LQpv6zXv9mC2HoFDes/dP/x3eUAAAC7AAAADgAAAPKa01J+3cWja4OICn9Bog4a4IlRAwAAAIABAADjwRLJlKVAir+HlVAUCW
ojKksMqGyWGhwhMoqYP8ldIy7bb3UVQBp6M+fxVpSFFrz3bfDgJQNh/13YCAY1+r1JYvEner7cnGatDIjnMgNqQPN6alqM787pMz3/eIq0L0xI8rV
y99F/foV6aFcJVCvxsjL9wIQOd4AhjIgTfPTiAEC4UTl5Eg9YkKnjZXizpTxhReSZVMjIM8Fu2sjcvzh1Q8POqYcEuU5sZhQbLVjUvRpou2HpgOTw
hcXW+X4gWpMlXsVkwV7F24j4Ax6eiyaSp1HCx4savMxcyA3cxwp7dTUJnFPVr8npfqp2H3ai3khSIefMc7d8gyajJnOLlJQMzf6O5HrlVeuixd6hB
dwd

<14>Aug 10 18:22:18 nethsm CEF:0|nCipher Security|nShield Solo XC|12.60.9|3|ssign-cert|5|esn=7109-02E0-D947
rsid=9 rtc=1628616065740 tpbl=2140 findex=3 flen=450
frag=B1Ku9rirlixkgEd+73tMVJ1FQz85aCWuRqJl04YB1YwFvZgvRXhHvzqLFeJZAUerKlLgIaZwDq1twoXzvHq88QcJdbr0i4+87VorPKkEjKtS
SGHOVkkHhoBC8uNgYXnTBxqcqCqpZl4whuiEBmJQLcwgAAAAg8rgckmo3ArobecQooPxQ9AjYbCmAoKOUTRi7grTzPyAAQAA3Bvuz+tQ1uh5LvuKM
LTtGDTTplG7ks6ZkL8b+F2UW37jfN3lap27oAZq1otU4FOP4EVvoMmNSdI4uzCPi7VgcI3AcIkdjZIwbpYf9XQwvFwMxYvdBPGhPtc/t8Lslgs97r
MkES4ZciNI/NwjKp0fW4kCiSBSUQUAUcp6vgqg2vVL9naqRHhXNRJuweaRtOO60z0mBkTgCnAvscdr2ymErrWDZArHosYXJZrXghjNmXvu+rS8GvT
vTc

<14>Aug 10 18:22:18 nethsm CEF:0|nCipher Security|nShield Solo XC|12.60.9|3|ssign-cert|5|esn=7109-02E0-D947
rsid=9 rtc=1628616065742 tpbl=2140 findex=4 flen=450
frag=189gfRjMpL5aBYYAkl1XqWDGHhFcltxTSzCMgWalxMOeOQxaZLbzYDtl2/udXIo0bn/PTgaOkPYTypvzFwsQM4axpDzVYCE064YVoyjUpgWB
U4kMlC4JuH5ytJAM+uA67xu36Iqx3j/mjMozTR1rGJuH+b314zgHRjvfr8AA0juiXn8tdxkFFRQlzYC9Ulw4g6fOSa06ecBIOA5Q0ylvdVjqmcX2+
+J+snC0wTxHV+vLKWh8m7/DDjExTXKpHo5EqyQB24tHCogAEAAJRhMgdRJ9i6dxLTUzQFo7ZzcQpZFdbFF4TN5+8Z/TER2/toZg+oNQkD2Eshd/T4
4p3WQr1XkQFjFe3/syQu8pS4Q0inmZzx6leFN99pXUhxXPnNIIpuc8UHJWips7nQnlsa3ER6evViNQanQSBTYEmuibXRITin8NM3h0RTJekDGw1J2
wj0

<14>Aug 10 18:22:18 nethsm CEF:0|nCipher Security|nShield Solo XC|12.60.9|3|ssign-cert|5|esn=7109-02E0-D947
rsid=9 rtc=1628616065743 tpbl=2140 findex=5 flen=340
frag=DjWEMgrKNb0X6hYqOPg/Ozid2ytohqdRx0Hr2zel6Ha17HgFfAkXlYaPXxdIGnrsuNwub43HU9iVByMdSe8kjOigEW/jXzTZQMFPy/Iy1+GC
V3P8GW+liYE/DM5auWH1a9NWbfjGTWQK4A820Eqy9zSUzOBnbn/gCLVUUMbSAGKY41Dtx7pUNU0TQpEMydBmfaoQhyx07fj070t8Zc2Ut2e5a2s/u
xwb0nSQNYJ25SBDb8UjHnhtNJVHiANHU3Qu1JcUnEKUwQW00dLeLOaMtd2Ldy+QaHHjsk7WYhboWCCSsWiHa4z3nVQsr3BrrkxBUlilWtXQlVVmEa
oAAAA=
```

25.5.8.2. Log Messages and Signature Block

This example shows a sequence of audit log messages with a Signature Block after 10 messages. These are in the same `rsid` as the previous example. The log sequence number for this excerpt starts at 31 and the last log message before the Signature Block is sequence number 40. The name element identifies the command being executed by the HSM. Each of the example commands operates on an nCore Key and this is identified by the nCore key hash of the relevant key.

The Signature Block has name element `ssign` identifying it as a Signature Block. The `gbc` is 6 meaning this is the 7th Signature Block in this Reboot Session ID. The `fmn` is 31 and `hcnt` is 10 meaning that this Signature Block covers messages 31 to 40. As Audit Logs are generated this sequence will be repeated. Once this Signature Block has been received and with the log signing public key available the signature on this Signature Block can be verified and then the hashes of the individual log messages can be calculated and compared with the hashes recorded in the Signature Block for the corresponding log message to allow the detection of tampering.

```
<134>May 15 12:42:09 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Destroy|1|esn=1111-2222-3333 rsid=8
rtc=1526388047690 seqNo=31 source=host outcome=success
hkey=c4ab637985a542e7eb3eb4838f57872d5422bbb4

<134>May 15 12:47:09 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Destroy|1|esn=1111-2222-3333 rsid=8
rtc=1526388347977 seqNo=32 source=host outcome=success
hkey=c4ab637985a542e7eb3eb4838f57872d5422bbb4

<134>May 15 12:52:10 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Destroy|1|esn=1111-2222-3333 rsid=8
rtc=1526388648265 seqNo=33 source=host outcome=success
hkey=c4ab637985a542e7eb3eb4838f57872d5422bbb4

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.20|1|Cmd_GenerateKeyPair|1|esn=1111-2222-
3333 rsid=8 rtc=1526388719548 seqNo=34 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Export|1|esn=1111-2222-3333 rsid=8
rtc=1526388719549 seqNo=35 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Export|1|esn=1111-2222-3333 rsid=8
rtc=1526388719549 seqNo=36 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Destroy|1|esn=1111-2222-3333 rsid=8
rtc=1526388719550 seqNo=37 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Import|1|esn=1111-2222-3333 rsid=8
rtc=1526388719550 seqNo=38 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Destroy|1|esn=1111-2222-3333 rsid=8
rtc=1526388719551 seqNo=39 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|1|Cmd_Import|1|esn=1111-2222-3333 rsid=8
rtc=1526388719552 seqNo=40 source=host outcome=success
hkey=cec7b0b1ef47d4141d65fdde9f9d23e854391dea

<134>May 15 12:53:21 myhost CEF:0|nCipher Security|nShield Solo|12.60.2|3|ssign|5|esn=1111-2222-3333 rsid=8
rtc=1526388719552 gbc=6 fmn=31
hcnt=10hb=8ogF/vsd9SwQ+qWEnealDuczRE9XbE9Rf3k3dc51SXo=&Xq5dbTEtg0l6pHQ7n6G16X+muB6C6VzN4FKbuqZqNUQ=&tOdfZk5Uvs6W9
E8mJyEU4kkNHAImwYft0v+7mHSL7VY=&ct+wqWsptfc+asw9ppvYVNbmkpqU3/WbXm5nHJPri9E=&riegp83m5c3jYt2vymNf61ov+Jf8JqCeLSyh
iSDRXfA=&WGKibqrL7AEKXU3Wu2IG1VcctIIESxyXluLcAFR+qvE=&TDrCHjkuA2fS5ZBgVu4Wspta+MbdhkyrGXeHHWn9Xck=&EwflNUbAIIVVdj
06PlEvjlljRoIXiprw56cqGGLbp9w=&Y1j1x6GEOofqG9XdAqk91zFhA8bczBocsttvvNND3tI=&p5YnBsabKWL4F5+2WNmIJ9DmKuSeUz5v3qogO
ilztRY=sign=MEUCIDa7RJlKmMRJ4KrDWxDYYok1t9ptQXqH1nPE5xLihcgoAiEAr4JWBTxsF/XrN2kJVfiVpCRicbvgTNjjjnouohS6QTM=
```

25.6. Commands Audited

The Audit Logging facility generates log entries on the module for a set of nCore com-
mands and module operations. The commands and information logged for each command
are described in the following sections.

25.6.1. Key usage logging

By default the nShield Audit Logging Facility does not log usage of keys for cryptographic
operations such as sign, verify, encrypt and decrypt or their usage in channels for these pur
posed. Creation, Deletion and a number of other key operations are unconditionally logged
by default. The Audit Logging feature provides the capability to optionally log these opera-
tions. This is determined on a per-key basis by the LogKeyUsage permission group flag on
the ACL group authorizing the operation for which logging is desired. See the *nCore Devel-
oper Tutorial* for further information on ACLs.

The `generatekey` utility (see Key generation options and parameters) provides the ability to
set this permission group flag when a key is generated by either:

- Specifying `logkeyusage=yes` as an option on the command line
- Answering **yes** to the `logkeyusage` question if the command is being used interactively.

When `generatekey` is used this flag is applied to all permission groups but is only checked
by the HSM on the group authorizing the desired action.

The following example shows this set on permission group **0** of a key's ACL.

```
groups[    0].flags= LogKeyUsage
                .n_limits= 0
                .n_actions= 2
                .actions[    0].type= OpPermissions
                                    .details.oppermissions.perms= DuplicateHandle
ExportAsPlain GetAppData SetAppData
    ReduceACL ExpandACL Encrypt Verify UseAsBlobKey GetACL
```

In the following sections, the tables will indicate if this mechanism is required to generate a
log message for a specific command or key.

25.6.2. Commands generating Audit Log messages

The following tables list the nCore commands that generate Audit Logging messages. For
each command they identify command specific data that is contained in the log message
and the CEF extension used to identify it.

nCore Key and Logical Token hashes are the standard nCore identifying hashes. They are
used to identify a key or logical token as it is an invariant for the key or logical token. These

hashes are logged as lower-case hex encoding. In some cases a short hash may be presented. This is the first 10 bytes of the hash in a lower-case hex encoding.

For each command logged the command is specified by the `name` element of the CEF header. The other elements of the CEF header are filled as detailed in the previous section. All commands being logged will also include the following CEF extensions:

Extension	Description
`esn`	ESN of the HSM
`rsid`	Reboot Session ID
`rtc`	Timestamp as milliseconds since the epoch derived from the HSM's Real Time Clock
`seqNo`	Sequence number of the Audit Log message
`outcome`	Success or failure

> ℹ️ Identifying Log Messages: As Audit Logging will potentially be running for a long time, the identification of a log message from an HSM based on the `rsid` and `seqNo` will not hold if the HSM is not restarted before the `seqNo` is reset when it reaches 9999999999. In this case account can be taken of the `gbc` as this will increment at a slower rate than the `seqNo`. Therefore messages in the same `rsid` with the same `seqNo` will have significantly different values of `gbc` (the mapping between `seqNo` and `gbc` is determined by the Signature Block containing the message in question). When the `gbc` is reset the `rsid` is incremented so counting begins in a new Reboot Session. Account can also be taken of the value of the `rtc`.

25.6.3. Key commands

Commands with **Yes** in the **Requires logkeyusage ACL** column will only be logged if the Key's ACL contains the `LogKeyUsage` Flag in the permission group authorizing the operation.

Command	Command Specific Information Logged	Extension	Requires logkeyusage ACL
`Cmd_Sign`	nCore key hash	hkey	Yes
`Cmd_Encrypt`	nCore key hash	hkey	Yes
`Cmd_Decrypt`	nCore key hash	hkey	Yes
`Cmd_Verify`	nCore key hash	hkey	Yes

Command	Command Specific Information Logged	Extension	Requires logkeyusage ACL
Cmd_ChannelOpen	nCore key hash	hkey	Yes
	channel mode	mode	-
	ℹ️ Mode is one of encrypt, decrypt, sign, verify		
Cmd_Import	nCore key hash	hkey	No
Cmd_Export	nCore key hash	hkey	No
Cmd_Duplicate	nCore key hash	hkey	No
Cmd_GenerateKey	nCore key hash	hkey	No
Cmd_GenerateKeyPair	nCore key hash	hkey	No
	ℹ️ nCore Key Hashes of private and public key halves are identical		
Cmd_SetAppData	nCore key hash	hkey	No
Cmd_SetACL	nCore key hash	hkey	No
Cmd_Destroy	nCore key hash	hkey	No
	Also used for Logical Tokens		
Cmd_DeriveKey	nCore Key Hash of derived key	hkey	No
	nCore Key Hash of base key	hbase	Yes
	nCore Key Hash of wrap key	hwrsp	Yes
	nCore Key Hash of third input key	hin3	Yes
	nCore Key Hash of fourth input key	hin4	Yes
	nCore Key Hash of fifth input key	hin5	Yes
	ℹ️ The nCore Key Hashes for the input keys will only be included in the audit log if the Permission Group for the DeriveKey action has the LogKeyUsage flag.		
Cmd_MakeBlob	nCore Key Hash	hkey	No
	nCore LT Hash	hwrap	
Cmd_LoadBlob	nCore Key Hash	hkey	No
	nCore LT Hash	hwrap	
Cmd_SetKM	nCore key hash	hkey	No

Command	Command Specific Information Logged	Extension	Requires logkeyusage ACL
Cmd_RemoveKM	nCore key hash	hkey	No

25.6.4. Logical Token and Share Commands

These commands do not use the logkeyusage ACL mechanism and log unconditionally.

Command	Command Specific Information Logged	CEF Extension
Cmd_ChangeSharePIN	KM nCore Key Hash	hkm
	Short LT Hash	shorthash
	Share Index	shareindex
	Slot	tokenslot
Cmd_Destroy	LT Hash	hkey
	ℹ️ Cmd_Destroy is used for Logical Tokens as well as Keys	
Cmd_EraseShare	Short LT Hash	shorthash
	Share Index	shareindex
	Slot	tokenslot
Cmd_GenerateLogicalToken	KM nCore Key Hash	hkm
	nCore LT Hash	htok
	Token Shares Needed	sharesneeded
	Token Total Shares	sharestotal
	Token time-limit	timelimit
Cmd_LoadLogicalToken	KM Hash	hkm
	LT Hash	htok
	Token Shares Needed	sharesneeded
	Token Total Shares	sharestotal
	Token time-limit	timelimit

Command	Command Specific Information Logged	CEF Extension
Cmd_ReadShare	LT Hash	htok
	Share Index	shareindex
	SlotId	tokenslot
	Share Left	sharesleft
	ℹ️ This is the remaining number of shares required to reconstruct the Logical Token. It reduces to 0 when a quorum of the Shares have been read.	
Cmd_WriteShare	LT Hash	htok
	Share Index	shareindex
	SlotId	tokenslot

25.6.5. Administrative Commands

These commands are logged unconditionally.

Command	Command Specific Information Logged	CEF Extension
Cmd_InitialiseUnit	Type of KML key	kmltype
	ℹ️ DSAp3072s256	
Cmd_InitialiseUnitEx	Type of KML key	kmltype
	InitialiseUnitEx Flags	flags
	ℹ️ DSAp3072s256 Combination of AuditLogging and UseFIPSApprovedInternalMechanisms	
Cmd_SetNSOPerms	nCore Key Hash of Security Officers Key	hknso
	SetNSOPermsFlags	flags
	ℹ️ Combination of AlwaysUseStrongPrimes, DisablePKCS1Padding, FIPSLevel3Enforcedv2 and CommonCriteriaCMTSRestrictions	
Cmd_CreateSeeWorld		
Cmd_SetSEEMachine		

Command	Command Specific Information Logged	CEF Extension
Cmd_SetRTC	Previous RTC	prevrtc
	This command is not applicable for nShield 5c. 🛈 ┃ New RTC value will be shown in the rtc extension	

25.6.6. Dynamic Slot Commands

These commands are logged unconditionally.

Command	Command Specific Information Logged	CEF Extension
Cmd_DynamicSlotsConfigure	Count of Dynamic Slots to be Configured	slotcount
Cmd_DynamicSlotCreateAssociation	Slot Id for Association	slotid
EstablishSecureChannel	ESN of smartcard	smartcardesn
	🛈 ┃ This internal event is logged with name element EstablishSecureChannel, a Severity of 2 and a Device Event Class Id of 2 in the CEF header and with source=internal in the CEF extensions.	

25.6.7. Heartbeat

The heartbeat is a periodic audit log message sent every 15 minutes. This audit log message indicates that the HSM is still active. After a heartbeat event is logged a Signature Block is generated including the heartbeat log message and any outstanding audit log messages. Waiting until the heartbeat is logged before restarting the HSM will ensure outstanding log messages can be verified.

Command	Command Specific Information Logged	CEF Extension
heartbeat	nCore Key Hash of Security Officers Key	hknso
	🛈 ┃ The heartbeat is logged in the CEF header with name element heartbeat, Severity 4, and Device Event Class Id 2. In the CEF extensions it's logged with source=internal.	

25.6.8. Post Reboot Logging

The nShield HSM has a number of commands and errors that cannot be logged directly when they occur. This applies primarily to errors detected during processing or self test and the reboot command `Cmd_ClearUnit`. The strategy adopted for these is to persist sufficient information and replay them as log entries after a successful reboot of the HSM. These reboot event messages occur after the Certifier Block has been emitted.

Each of these messages are emitted with `rsid` and `seqNo` relating to the current session and will have a prevrtc CEF element recording the RTC at the time of the event. The name element will identify the event. If the event is associated with a nCore SOS code this will be indicated by a sos CEF extension and an appropriate code. The Device Event Class Id is set to 5 and Severity will be set to 10 for errors or 6 for shutdown events. The source CEF extension will be internal. The following table lists the events replayed in a post reboot log. The available events depend on the type of HSM.

Event Id	Event	SOS Code
`Cmd_ClearUnit`	Cmd_ClearUnit	
`Cmd_Fail`	Cmd_Fail	D
`Environment_SensorFail`		HV
`Temperature_OutofRange`		T
`RNG_PeriodicTestFail`		HRTP
`SOS`	Starting up crypto offload	HF
`SOS`	cache keygen failed	HR
`Voltage_Tamper`		V
`Battery_Tamper`		B
`Unknown_Tamper`		TAMPER
`SelfTestFail`	POST test timed out	HC0TTO
	POST test failed: lock failure detected	HC0LC
	POST test failed: TEST_STARTED	HC0TS
	POST test failed: PROCESS_STARTED	HC0PS
	POST test failed: CPUID_CHECK	HC0CC
	POST test failed: SRAM_ALLOC	HC0SA
	POST test failed: SRAM_WRITE	HC0SW
	POST test failed: SRAM_READ	HC0SR

Event Id	Event	SOS Code
	POST test failed: SRAM_FREE	HC0SF
	POST test failed: CRAM_ALLOC	HC0CA
	POST test failed: CRAM_GETCACHED	HC0CG
	POST test failed: CRAM_WRITE	HC0CW
	POST test failed: CRAM_READ	HC0CR
	POST test failed: CRAM_FREE	HC0CF
	POST test failed: LOCK_CHECK	HC0LC
	POST test failed: RTC_CHECK	HC0RT
	POST test failed: KAT_DSA	HC0KS
	POST test failed: KAT_ECDSA	HC0KC
	POST test failed: KAT_DES	HC0KE
	POST test failed: KAT_DES3	HC0KF
	POST test failed: KAT_DES3CBCMAC	HC0KO
	POST test failed: KAT_AES	HC0KA
	POST test failed: KAT_AESCMAC	HC0KB
	POST test failed: KAT_AESCBCMAC	HC0KD
	POST test failed: KAT_SHA1	HC0KH
	POST test failed: KAT_SHA1HMAC	HC0KM
	POST test failed: KAT_SHA224HMAC	HC0KN
	POST test failed: KAT_SHA256HMAC	HC0KJ
	POST test failed: KAT_SHA384HMAC	HC0KP
	POST test failed: KAT_SHA512HMAC	HC0KI
	POST test failed: KAT_RSA	HC0KRH
	POST test failed: KAT_NISTKDF	HC0KDF
	POST test failed: KAT_HASHDRBG	HC0HD
	POST test failed: KAT_RSAOAEP	HC0KZ
	POST test failed: KAT_25519	HC0KX
	POST test failed:unknown	HC0H

As an example, the following shows a post reboot log of `Cmd_ClearUnit`. In this excerpt, it can be seen after the last fragment of the Certifier Block. A Signature Block is generated after the reboot log entries.

```
....
<134>May 16 15:08:45 myhost2 CEF:0|nCipher Security|nShield Solo XC|12.60.2|3|ssign-cert|5|esn=1111-2222-4444
rsid=2 rtc=1524140117693 tpbl=2140 findex=5
flen=340frag=3/ITRJT4T/qgd2ZEJufIzCR+nR9lngOrmogj+5JM7VMFLsWGDxUqxmFlpqs52T2zWuYIeFHGQfx9WS9PUhf2eLMyF/7onn+hFUs5
7/GSZlGbCnxWybfPN27oyXjHE7pfyOrWRVKlIw8UULHVezVsxeIsZuuNEsZa5gUQ++DkoTu5M2BoPr4A+6dVL2eDhOF1m2zKATfk2moW93GkA3AO7
lNPV5xU76ujo2tT7Mttvg+vyddiF2UWe6n75U0FMFjlM9WnhpFAhNk9mJPrNZ5smf4i9JuNKZat+5tq5w2b/a8Sy01EVEKtJI5SSjahtp5z77RseQ
8H8ytsw6oAAAA=sign=MEUCIHgrF1m7t9X5xsl/gXwlju0bPfFPjJeIeIiH8TKSN7prAiEAs3lPS62zX3TE940/Dw9/1gVradNi62wrQI+WlSI4IY
U=
<134>May 16 15:08:45 myhost2 CEF:0|nCipher Security|nShield Solo XC|12.60.2|5|Cmd_ClearUnit|6|esn=1111-2222-4444
rsid=2 rtc=1524140117693 seqNo=1 source=internal prevrtc=1524140108693

<134>May 16 15:08:45 myhost2  CEF:0|nCipher Security|nShield Solo XC|12.60.2|3|ssign|5|esn=1111-2222-4444 rsid=2
rtc=1524140117693 gbc=0 fmn=1 hcnt=1
hb=nwtggjmPYA1TR07KhdOHoyytxLb7RDvg7Wpw6FfAiC4=sign=MEYCIQDxIIJZRfKsXpMMoQ3GDEkTZ/+DTuEdNLKwHQzllflMUQIhAPipdSPrB
SUnarrtjMslYS4k3RPCXcNoO16xEhg/907z
```

25.6.9. Tracing Key Usage

With the information logged as detailed in the preceding sections it is possible to trace back from a Key Command to the loading of the Key, then to loading the Logical Token and reading the Shares that constitute the Logical Token.

The following example shows the notional traceback from a `Cmd_Encrypt` operation. This command logs the nCore Key Hash **KKKKKKK**. Prior to this the Key was loaded onto the HSM using `Cmd_LoadBlob` which correlates the nCore Key Hash with the ncore Hash of the Logical Token that authorized loading the Key. Tracing further back we can identify the shares used to reconstruct that Logical Token. In this example two shares are required identified by share indices S1 and S2. The share index identifies a specific card in an OCS cardset.

Command	Key Hash	Logical Token Hash	Share Index
`Cmd_Encrypt`	KKKKKKK		
`Cmd_LoadBlob`	KKKKKKK	LLLLLLL	
`Cmd_Read Share`		LLLLLLL	S2
`Cmd_Read Share`		LLLLLLL	S1
`Cmd_LoadLogicalToken`		LLLLLLL	

25.7. Audit Log Verification

The audit logs produced when AuditLogging feature is active can be verified using the infor

mation contained in the audit logging metadata. Every HSM enrolled into a Security World with AuditLogging enabled generates an HSM-specific log signing private key (KAL) that is maintained in the HSM's non-volatile memory until the module is re-initialized. The public key corresponding to this private key is sent as a Certifier Block by the HSM when Audit Logging is configured either by Security World creation or by indoctrination into an existing Audit Logging Security World. Every Signature Block sent by the HSM is generated using the log signing private key. The Audit Log can be verified as follows:

- Extract the KAL public key from the Certifier block
- Verify the Signature Blocks
- Verify the log message hashes in the Signature Block against hashes of the received logs to determine if any messages have been tampered
- Identify any missing log messages.

The basics of the verification approach is shown on the Audit Log Verification diagram.

To support Audit Log verification, Entrust provide an example verification program written in Python to serve as an example for developing a more comprehensive verification solution.

25.7.1. Running the example verification program

The example verification program can be found in `/opt/nfast/python/examples/audit-log-verifier.py` (**Linux**) or `C:\Program Files\nCipher\nfast\python\examples\audit-log-verifier.py` (**Windows**).

This program requires the use of the nShield Python interpreter. This is necessary to provide support for the nShield specific marshalling functions used to export the log signing public key. The example verification program also requires the presence of an nShield HSM accessible to the machine on which the verification is to be performed. This is required to perform the cryptographic operations necessary to verify the log signing public key and the Signature Blocks. This HSM does not need to be the same HSM on which the logs were generated, nor does it need to be in a Security World.

The Audit Log verifier program is run with a command of the form:

```
python audit-log-verifier.py [-h] [-e ESN] SYSLOG
```

Where:

Parameter	Function
`-h|--help`	Displays the help message
`-e ESN|--esn ESN`	ESN of the logevents to be verified
`SYSLOG`	Location of the syslog file to be verified

> **ℹ** Make sure that you use the nShield Python.

25.7.1.1. Results

Results from the Audit Log Verifier are written to several different files and saved in a sub-directory called LogResult. See example below for more detail.

25.7.1.2. Example

Running a command of the form:

```
python audit-log-verifier.py AuditLogInputFile.txt
```

Should produce a screen output similar to the following:

```
FIRST LOG INSTANCE-1 for ESN:9204-02E0-D947 @ Line:1 rsid:8 ######

Verifying certifier block...
Verification of CERTIFICATE Success
Verifying a cert fragment...Line:1
Verifying a cert fragment...Line:2
Verifying a cert fragment...Line:3
Verifying a cert fragment...Line:4
Verifying a cert fragment...Line:5
Verifying SB....Lineno:7:InstanceNo:1
Verifying SB....Lineno:18:InstanceNo:1
Valid Hash list from SBs written to ValidHashes_fromSB_forInst1.txt
No entry in SB for event @ Line:19 : Seq No:12 --
No entry in SB for event @ Line:20 : Seq No:13 --

Verified cert blocks written to./LogResult/CBs.txt
Sig blocks written to./LogResult/SBs.txt
Log messages written to./LogResult/AllEvents.txt
Anything that did not match (incl.
invalid cert blockfragments)  written to :./LogResult/Inconsistent.txt
```

The screen output indicates the contents of the main results files which are stored in the *LogResult* sub-directory. The contents of the folder will vary slightly depending on the log contents and whether there were any failures, but should be similar to the following:

```
AllEvents.txt               # log entries relating to events
CBs.txt                     # certificate blocks
Inconsistent.txt            # inconsistent log entries - should be empty []
```

```
                                (assuming no inconsistencies)
Instance.txt                    # esn number and other info relating to the log
InvalidHashes_fromSB_forInst1   # only exists if verification failed due to signature
block (forInst1 refers to

                                first logging instance/world found in the
                                log file)
SBs.txt                         # signature blocks
Tampered_logs.txt               # contains log messages that did not verify - e.g. due
                                to a corrupt signature block.
This file only exists
                                if verification failed.
Unverified_logs.txt             # unverified log entries - e.g. any trailing entries
                                from the end of the log file that lack an accompanying
                                signature block
ValidHashes_fromSB_forInst1.txt # valid hashes from the signature blocks
                                (forInst1 refers to first logging
                                instance/world found in the log file)
Verified_logs.txt               # verified log messages
```

Use a text editor to examine the files as required to check the verification. Note that `Inst1` in the filenames refers to the first `logging world` instance in the log, see Program Architecture. If the log contains messages relating to more than one `logging world`, files relating to subsequent instances will be tagged with `Inst2`, `Inst3` etc.

If the verification fails, screen output should indicate the source of the failure. For example, output for a log where a log message was missing would look something like this:

```
FIRST LOG INSTANCE-1 for ESN:9204-02E0-D947 @ Line:1 rsid:8 ######

Verifying certifier block...
Verification of CERTIFICATE Success
Verifying a cert fragment...Line:1
Verifying a cert fragment...Line:2
Verifying a cert fragment...Line:3
Verifying a cert fragment...Line:4
Verifying a cert fragment...Line:5
Verifying SB....Lineno:7:InstanceNo:1
Verifying SB....Lineno:17:InstanceNo:1
Verifying SB....Lineno:28:InstanceNo:1
Valid Hash list from SBs written to ValidHashes_fromSB_forInst1.txt
No entry in SB for event @ Line:29 : Seq No:22 --
No entry in SB for event @ Line:30 : Seq No:23 --
@@@@@@@@ Some Log events present in the SB but missing in Log file           # indicates
missing log message
Verified cert blocks written to./LogResult/CBs.txt
Sig blocks written to./LogResult/SBs.txt
Log messages written to./LogResult/AllEvents.txt
Anything that did not match (incl.
invalid cert blockfragments)  written to :./LogResult/Inconsistent.txt
```

Output for a log where a log message had been tampered with or is otherwise corrupt might look like this:

```
FIRST LOG INSTANCE-1 for ESN:9204-02E0-D947 @ Line:1 rsid:8 ######

Verifying certifier block...
Verification of CERTIFICATE Success
Verifying a cert fragment...Line:1
```

```
Verifying a cert fragment...Line:2
Verifying a cert fragment...Line:3
Verifying a cert fragment...Line:4
Verifying a cert fragment...Line:5
Verifying SB....Lineno:7:InstanceNo:1
Verifying SB....Lineno:18:InstanceNo:1
Verifying SB....Lineno:29:InstanceNo:1
Valid Hash list from SBs written to ValidHashes_fromSB_forInst1.txt
Validating Log @ Line No:10 SeqNo:4 is Failed ----                            # indicates
tampered log entry
***** Hash Mismatch No entry in SB for event @ Line:30 : Seq No:22 --
No entry in SB for event @ Line:31 : Seq No:23 --
Verified cert blocks written to./LogResult/CBs.txt
Sig blocks written to./LogResult/SBs.txt
Log messages written to./LogResult/AllEvents.txt
Anything that did not match (incl.
invalid cert blockfragments)  written to :./LogResult/Inconsistent.txt
```

The tampered log line(s) will be listed in output file `Tampered_logs.txt`.

Output for a log where the signature block is corrupt will look something like this:

```
FIRST LOG INSTANCE-1 for ESN:9204-02E0-D947 @ Line:1 rsid:8 ######

Verifying certifier block...
Verification of CERTIFICATE Success
Verifying a cert fragment...Line:1
Verifying a cert fragment...Line:2
Verifying a cert fragment...Line:3
Verifying a cert fragment...Line:4
Verifying a cert fragment...Line:5
Verifying SB....Lineno:7:InstanceNo:1
Verifying SB....Lineno:18:InstanceNo:1
Signature Tampered B64 decode

//k=&s0QG1C08QBi34gaTU2+rUzp/dwtAXi9Hv0IjDvDL/yg=&Im5nW+OX0gbdlLnrFLxsZtR4meDSEXG5JXtkMmltTZU=&LGAXS1nvgHElvXhk8R
VT2lCK2NMtXyD9OYTecVOaaBk=&MbJAk706yU2+QykWmtfnCV0lxn/enber8aJK3cZyxLg=&y2qxF5VGm/X/h6ZcZ5iOes7ZAFpqM/6ND8nAXzCM/
bY=&kWjEaGIclJv494A1ZcUgGHJko7AeKvUUqVimhfExioU=  Length: 577
Verifying SB....Lineno:29:InstanceNo:1
Valid Hash list from SBs written to ValidHashes_fromSB_forInst1.txt
In-Valid Hash list from SBs written to InvalidHashes_fromSB_forInst1.txt
        Log entry found in Tampered SB. Line no:8 SeqNo:2
        Log entry found in Tampered SB. Line no:9 SeqNo:3
        Log entry found in Tampered SB. Line no:10 SeqNo:4
        Log entry found in Tampered SB. Line no:11 SeqNo:5
        Log entry found in Tampered SB. Line no:12 SeqNo:6
        Log entry found in Tampered SB. Line no:13 SeqNo:7
        Log entry found in Tampered SB. Line no:14 SeqNo:8
        Log entry found in Tampered SB. Line no:15 SeqNo:9
        Log entry found in Tampered SB. Line no:16 SeqNo:10
        Log entry found in Tampered SB. Line no:17 SeqNo:11
No entry in SB for event @ Line:30 : Seq No:22 --
No entry in SB for event @ Line:31 : Seq No:23 --
Verified cert blocks written to./LogResult/CBs.txt
Sig blocks written to./LogResult/SBs.txt
Log messages written to./LogResult/AllEvents.txt
Anything that did not match (incl.
invalid cert blockfragments)  written to :./LogResult/Inconsistent.txt
```

The failed log messages should be reported in `Tampered_logs.txt` in the `LogResult` folder.

If the certificate block is corrupt, output will be similar to that shown below. In this case, the

`CBs.txt` file may be empty and the cert block fragments will be written to `Inconsistent.txt`.

```
FIRST LOG INSTANCE-1 for ESN:9204-02E0-D947 @ Line:1 rsid:8 ######

Verifying certifier block...
Verification of CERTIFICATE Success
Verifying a cert fragment...Line:1
Verifying a cert fragment...Line:2
Verifying a cert fragment...Line:3
Signature Tampered B64 decode
('Failed fragment:', 3, '<134>May  2 16:10:38 exampleCB1.myexample.com CEF:0|nCipher Security|nShield Solo
XC|12.60.2|3|ssign-cert|5|esn=9204-02E0-D947 rsid=8rtc=4294967386734000 tpbl=2140 findex=3 flen=450
frag=B1Ku9rirlixkgEd+73tMVJ1FQz85aCWuRqJl04YB1YwFvZgvRXhHvzqLFeJZAUerKlLgIaZwDq1twoXzvHq88QcJdbr0i4+87VorPKkEjKtS
SGHOVkkHhoBC8uNgYXnTBxqcqCqpZl4whuiEBmJQLcwgAAAAg8rgckmo3ArobecQooPxQ9AjYbCmAoKOUTRi7grTzPyAAQAA3Bvuz+tQ1uh5LvuKM
LTtGDTTplG7ks6ZkL8b+F2UW37jfN3lap27oAZq1otU4FOP4EVvoMmNSdI4uzCPi7VgcI3AcIkdjZIwbpYf9XQwvFwMxYvdBPGhPtc/t8Lslgs97r
MkES4ZciNI/NwjKp0fW4kCiSBSUQUAUcp6vgqg2vVL9naqRHhXNRJuweaRtOO60z0mBkTgCnAvscdr2ymErrWDZArHosYXJZrXghjNmXvu+rS8GvT
vTc sign=MEYCIQCXhbJeFTv8oR8a51aU3Os9w2Vs9w67mzYk584Gy+MdbgIhAOD0bAwU0Vw1xOmR2oerqWKLFeawZe5rONmDMZFbJoB')
Fragment verification unsuccessful!
Adding all fragments in this CB to Inconsistent.txt
('No Valid CB for this instance:', 1)
In-Valid Hash list from SBs written to InvalidHashes_fromSB_forInst1.txt
        Log entry found in Tampered SB. Line no:6 SeqNo:1
        Log entry found in Tampered SB. Line no:8 SeqNo:2
        Log entry found in Tampered SB. Line no:9 SeqNo:3
        Log entry found in Tampered SB. Line no:10 SeqNo:4
        Log entry found in Tampered SB. Line no:11 SeqNo:5
        Log entry found in Tampered SB. Line no:12 SeqNo:6
        Log entry found in Tampered SB. Line no:13 SeqNo:7
        Log entry found in Tampered SB. Line no:14 SeqNo:8
        Log entry found in Tampered SB. Line no:15 SeqNo:9
        Log entry found in Tampered SB. Line no:16 SeqNo:10
        Log entry found in Tampered SB. Line no:17 SeqNo:11
        Log entry found in Tampered SB. Line no:19 SeqNo:12
        Log entry found in Tampered SB. Line no:20 SeqNo:13
        Log entry found in Tampered SB. Line no:21 SeqNo:14
        Log entry found in Tampered SB. Line no:22 SeqNo:15
        Log entry found in Tampered SB. Line no:23 SeqNo:16
        Log entry found in Tampered SB. Line no:24 SeqNo:17
        Log entry found in Tampered SB. Line no:25 SeqNo:18
        Log entry found in Tampered SB. Line no:26 SeqNo:19
        Log entry found in Tampered SB. Line no:27 SeqNo:20
        Log entry found in Tampered SB. Line no:28 SeqNo:21
No entry in SB for event @ Line:30 : Seq No:22 --
No entry in SB for event @ Line:31 : Seq No:23 --

Verified cert blocks written to./LogResult/CBs.txt
Sig blocks written to./LogResult/SBs.txt
Log messages written to./LogResult/AllEvents.txt
Anything that did not match (incl.
invalid cert blockfragments)  written to :./LogResult/Inconsistent.txt
```

25.7.2. Program Architecture

The program takes and reads the input syslog file containing the log messages. It optionally sets the ESN of module for which log events are to be validated, if this was passed in. If an ESN is not provided as input then the first ESN found in the syslog will be processed.

The verifier calls its parse function which segregates the messages based on ESN, and cre-

ates lists of Certifier fragments, Signature Blocks and Log events, based on matching with regular expressions.

Syslog may have gathered logs from multiple sources. As such, the verifier has a concept of a *logging world*, which represents a set of logs, sigblocks and certblocks that belong together, from a Security World. Based on Reboot Sequence ID, Sequence Number of the Log event, Global block counter of the Signature block and Fragment index of the Certifier block, a logging world is identified and a logging instance is created.

All records are thus given a log-instance number, such that records with the same instance number belong together.

Each event can thus be uniquely identified via a tuple. For the log messages, signature blocks and certifier blocks these are respectively (`rsid` and `sequence number`), (`rsid` and `gbc`) and (`rsid` and `findex`).

The `reconstruct_CBs` function is then called to validate the certifier fragments (using calls to an nShield HSM for crypto functionality). It then reconstructs the certifier blocks from the certifier fragments.

> ℹ️ This does not require the HSM to be in the same Security World as the HSM that first generated the logs.

A list of valid and verified Certifier Blocks is created.

For any log instance one valid Certifier Block is enough to validate the events, so further cer tifier blocks are ignored after the first.

Next the `process_sbs` function is called. Signature Blocks for a supplied ESN are validated per log instance (once again via calls to the module for crypto functionality), using the KAL value taken from the Certifier block previously.

The validated Signature block hashes are maintained as a dictionary of hashes with keys as unique ids. These unique ids per instance are generated based on rsid and sequence num- bers.

The `process_logs` function is finally called. This generates the hash of each of the log events and matches against hashes from corresponding signature blocks. Verified and Tam pered log events are then written to different files in the *LogResult* folder.

25.7.3. Extended Verification

While the example verifier uses an HSM for cryptographic operation, it would be possible to use 3rd party cryptographic libraries to provide this functionality. This is outside the scope

of this document.

Currently the log messages are verified against the hash in the signature blocks, and the sig
nature of the signature blocks is verified against the key extracted from the certifier block.
The certifier block its self is not verified. A potential extension to the verifier tool would be
to verify the certifier block. The certifier block is signed by KLF2. This can be checked
against the KLF2 value found within the module's warrant. This would complete the chain of
trust.

Additionally, the example verifier does not cope with fields that rotate back around to zero
when their max size is exceeded, for example with the `gbk`, `rsid` or `seqno` fields. Currently
logs, SBs and CBs are uniquely identified by (`rsid` and `sequence number`), (`rsid` and `gbc`) and
(`rsid` and `findex`). This means that, if any of those values rotate back around to zero, we are
no longer able to uniquely identify them. As a potential extension, RTC or line number val-
ues could be used to solve this.

The example verifier does not detect missing/deleted log messages in the case where a
complete group of log messages are deleted, along with their corresponding `Signature-
Block`. Given that the `SeqNo` field increases for each log message, spotting missing `SeqNos`
would reveal missing or deleted log messages. This is a potential extension.

The example verifier expects a static, unchanging log file to be supplied to it. This would be
compatible with verifying a batch of log files at the end of each day, for example. A possible
extension would be to extend the verifier to cope with a live stream of logs, continuously
verifying them as they are generated.

26. Key generation options and parameters

This appendix describes the various options and parameters that you can set when running the `generatekey` utility to control the application type and other properties of a key being generated.

> **ℹ** For information about generating keys with the `generatekey` utility, see [Generating keys with the command line](#).

26.1. Key application type (APPNAME)

The `APPNAME` parameter specifies the name of the application for which `generatekey` can generate keys. Specifying an application can restrict your choice of key type. A value for `APPNAME` must follow any `OPTIONS` and must precede any parameters specified for the key:

Parameter	Description
`simple`	Specifying the `simple` application type generates an nShield-native key. No special action is taken after the key is generated.
`custom`	Specifying the `custom` application type generates a key for custom applications that require the key blob to be saved in a separate file.
	Specifying `custom` also causes the generation of a certificate request and self-signed certificate. However, we recommend that you specify the `simple` (instead of `custom`) application type whenever possible.
`pkcs11`	Specifying the `pkcs11` application type generates keys that are formatted for use with PKCS #11 applications and are given a suitable identifier. The set of possible supported key types is currently limited to: • `DES3` • `DH` • `DSA` • `ECDH` • `ECDSA` • `Ed25519` • `HMACSHA1` • `RSA` • `Rijndael` (AES) • `X25519` Some key types are only available if the features that support them have been enabled for the module, if the Security World is not compliant with FIPS 140 Level 3, or if you do not set the `--no-verify` option.

Parameter	Description
embed	Specifying the `embed` application type generates a PEM-format RSA/DSA key file that points to a key in `NFAST_KMDATA` so a software application can then use the HSM-protected key. In applications that use Security World software older than v12.60 and would use the legacy OpenSSL CHIL engine with `hwcrhk`: • The `plainname` specified in the `generatekey` command is used as the prefix for all 3 generated files (`.key`, `_req`, `_selfcert`) • `.key` is appended to all 3 files • The `embedsavefile` specified in the `generatekey` command is the destination for all 3 files In applications that use v12.60 or later Security World software : • The `plainname` specified in the `generatekey` command is used as the prefix for only the `.key` file, the prefix for the `_req` and `_selfcert` file is `embed<hash>` • `.key` is not appended to the `_req` and `_selfcert` files • The `embedsavefile` is the destination only for the `.key` file, `_req` and `_selfcert` are created in the directory from which `generatekey` was run from
kpm	Specifying the `kpm` application type generates a key for delivery by an nForce Ultra key server. The `generatekey` utility automatically creates a special ACL entry that permits a `kpm` to be delivered to an nForce Ultra's enrolled internal hardware security module.
seeinteg	Specifying the `seeinteg` application type generates an SEE integrity key. The DSA, RSA, ECDSA and KCDSA algorithms are supported. SEE integrity keys are always protected by an OCS and cannot be imported. You cannot retarget an existing key as an SEE integrity key.
seeconf	Specifying the `seeconf` application type generates an SEE confidentiality key. Both the Triple DES and AES algorithms are supported for this key type. SEE confidentiality keys are module-protected by default and cannot be imported. You cannot retarget an existing key as an SEE confidentiality key.

26.2. Key properties (NAME=VALUE)

The `NAME=VALUE` syntax is used to specify the properties of the key being generated.

> **ℹ** If a parameter's argument contains spaces, you must enclose the argument within quotation marks (`"  "`).

You can supply an appropriate *VALUE* for the following *NAME* options:

Option	Description
`alias`	The *VALUE* for `alias` specifies an alias to assign to the key.
`blobsavefile`	When using the `custom` application type, the *VALUE* for `blobsavefile` specifies a file name of the form *FILENAME_req.ext* to which the key blob is saved. Additionally, a text file containing information about the key is saved to a file whose name has the form *ROOT_inf.txt*; for asymmetric key types, the public key blob is also saved to a file whose name has the form *ROOT_pub.EXT*.
`cardset`	The *VALUE* for `cardset` specifies an OCS that is to protect the key (if `protect` is set to `token`). In interactive mode, if you do not specify an OCS, you are prompted to select one at card-loading time. The default is the OCS to which the card currently inserted in the slot belongs (or the first one returned by `nfk minfo`).
`certreq`	Setting `certreq` enables you to generate a certificate request when generating a PKCS #11 key (RSA keys only). The default behavior is to not generate a certificate request. To generate a certificate request you must set the *VALUE* for `certreq` to `yes`, which makes `generatekey` prompt you to fill in the extra fields required to generate a key with a certificate request. The resultant certificate request is saved to the current working directory with a file name of the form *FILENAME req.ext* (where *FILENAME* is a name of your choice). An extra file with a name of the form *FILENAME.ext* is also generated for use as a pseudo-key-header. This file can be removed after the certificate request has been generated. You can use `certreq` with the `--retarget` option to generate a self-signed certificate for an existing key.
`checks`	For RSA key generation only, this specifies the number of checks to be performed. Normally, you should leave *VALUE* empty to let the module pick an appropriate default.
`curve`	For ECDH and ECDSA key generation only, the *VALUE* for `curve` specifies which curves from the supported range to use. Supported curves are: ANSI-B163v1, ANSIB191v1,BrainpoolP160r1, BrainpoolP160t1, BrainpoolP192r1, BrainpoolP192t1, BrainpoolP224r1, BrainpoolP224t1, BrainpoolP256r1, BrainpoolP256t1, BrainpoolP320r1, BrainpoolP320t1, BrainpoolP384r1, BrainpoolP384t1, BrainpoolP512r1, BrainpoolP512t1, NISTP192, NISTP224, NISTP256, NISTP384, NISTP521, NISTB163, NISTB233, NISTB283, NISTB409, NISTB571, NISTK163, NISTK233, NISTK283, NISTK409, NISTK571, SECP160r1 and SECP256k1
`embedconvfile`	The *VALUE* for `embedconvfile` specifies the name of the PEM file that contains the RSA key to be converted.

Option	Description
`embedsavefile`	When using the `embed` application type, the `VALUE` for `embedsavefile` specifies the name for the file where the fake RSA private key is to be saved. The file has the same syntax as an RSA private key file, but actually contains the key identifier rather than the key itself, which remains protected. A certificate request and a self-signed certificate are also written. If the file-name is *ROOT*`.EXT` then the request is saved to *ROOT*`_req.EXT` and the self-signed certificate is saved to *ROOT*`_selfcert.EXT`.
`from-application`	When retargeting a key, the *VALUE* for `from-application` specifies the application name of the key to be retargeted. Only applications for which at least one key exists are acceptable.
`from-ident`	When retargeting a key, the *VALUE* for `from-ident` specifies the identifier of the key to be retargeted (as displayed by the `nfkminfo` command-line utility).
`hexdata`	The *VALUE* for `hexdata` specifies the hex value of DES or Triple DES key to import. The hex digits are echoed to the screen and can appear in process listings if this parameter is specified in the command line.
`ident`	The *VALUE* for `ident` specifies a unique identifier for the key in the Security World. For applications of types `simple`, this is the key identifier to use. For other application types, keys are assigned an automatically generated identifier and accessed by means of some application-specific name. The following characters are allowed in key IDs: • digits 0-9 • lower-case letters a-z • hyphen (-)
`keystore`	The *VALUE* for `keystore` specifies the file name of the key store to use. This must be an nShield key store.
`keystorepass`	The *VALUE* for `keystorepass` specifies the password to the key store to use.
`logkeyusage`	The *VALUE* for `logkeyusage` specifies if usage of the generated key in cryptographic operations is subject to audit logging. If set to **yes** the ACL of the generated key will predicate audit-logging entries to be made for cryptographic usages of the key. The default is **no**.
`module`	The *VALUE* for `module` specifies a module to use when generating the key. If there is more than one usable module, you are prompted to supply a value for one of them. The default is the first usable module (one in the current Security World and in the operational state). ℹ You can also specify a module by setting the `--module` option.

Option	Description
`paramsreadfile`	The *VALUE* for `paramsreadfile` specifies the name of the group parameters file that contains the discrete log group parameters for Diffie-Hellman keys only. This should be a PEM-formatted PKCS#3 file. If a *VALUE* for `paramsread-file` is not specified, the module uses a default file.
`pemreadfile`	The *VALUE* for `pemreadfile` specifies the name of the PEM file that contains the key to be imported. When importing an RSA key, this is the name of the PEM-encoded PKCS #1 file to read it from. Password-protected PEM files are not supported.
`plainname`	The *VALUE* for `plainname` specifies the key name within the Security World. For some applications, the key identifier is derived from the name, but for others the name is just recorded in `kmdata` (**Linux**) or `%NFAST_KMDATA%` (**Windows**) and not used otherwise.
`protect`	The *VALUE* for `protect` specifies the protection method, which can be `module` for security-world protection, `softcard` for softcard protection or `token` for Operator Card Set protection. The default is `token`, except for `seeconf` keys, where the default is `module`. `seeinteg` keys are always token-protected. The `softcard` option is only available when your system has at least one softcard present.
`pubexp`	For RSA key generation only, the *VALUE* for `pubexp` specifies (in hexadecimal format) the public exponent to use when generating RSA keys. We recommend leaving this parameter blank unless advised to supply a particular value by Support.
`recovery`	The *VALUE* for `recovery` enables recovery for this key and is only available for card-set protected keys in a recovery-enabled world. If set to `yes`, the key is recoverable. If set to `no`, key is not recoverable. The default is `yes`. Non-recoverable module-protected keys are not supported.
`seeintegname`	If present, the *VALUE* for `seeintegname` identifies a `seeinteg` key. The ACL of the newly generated private key is modified to require a certificate from the `seeinteg` key for its main operational permissions, such `Decrypt` and `Sign` (`DuplicateHandle`, `ReduceACL`, and `GetACL` are still permitted without certification.) If you use `seeintegname` to specify a key that has been recovered with the `rocs` utility, you must also use the `-N` option with `generatekey`.
`selfcert`	The *VALUE* for `selfcert` enables you to generate a self-signed certificate when generating a PKCS #11 key (RSA keys only). To generate a self-signed certificate request you must set `selfcert` to `yes`, which makes `generatekey` prompt you to fill in the extra fields required to generate a key with a self-signed certificate. The resultant certificate is saved to the current working directory with a file name of the form *FILENAME*`.ext`. You can use this parameter with the `--retarget` option to generated a self-signed certificate for an existing key.

Option	Description
size	For key types with variable-sized keys, the *VALUE* for size specifies the key size in bits. The range of allowable sizes depends on the key type and whether the --no-verify option is used. The default depends on the key type; for information on available key types and sizes, see Cryptographic algorithms. This parameter does not exist for fixed-size keys, nor for ECDH and ECDSA keys which are specified using curve.
strict	For DSA key generation only, setting the *VALUE* for strict to yes enables strict verification, which also limits the size to 2048 or 3072 bits. The default is no.
type	The *VALUE* for type specifies the type of key. You must usually specify the key type for generation and import (though some applications only support one key type, in which case you are not asked to choose). Sometimes the type must also be specified for retargeting; for information on available key types and sizes, see Cryptographic algorithms. The --verify option limits the available key types.
x509country	The *VALUE* for x509country specifies a country code, which must be a valid 2-letter code, for the certificate request.
x509dnscommon	The *VALUE* for x509dnscommon specifies a site domain name, which can be any valid domain name, for the certificate request.
x509email	The *VALUE* for x509email specifies an email address for the certificate request.
x509locality	The *VALUE* for x509locality specifies a city or locality for the certificate request.
x509org	The *VALUE* for x509org specifies an organization for the certificate request.
x509orgunit	The *VALUE* for x509orgunit specifies an organizational unit for the certificate request.
x509province	The *VALUE* for x509province specifies a province for the certificate request.
xsize	The *VALUE* for xsize specifies the private key size in bits when generating Diffie-Hellman keys. The defaults are 256 bits for a key size of 1500 bits or more or 160 bits for other key sizes.

26.3. Available key properties by action/application

The following table shows which actions (generate, import, and retarget) are applicable to the different *NAME* options:

Property	generate	import	retarget
alias	X	X	X

Property	generate	import	retarget
blobsavefile	X	X	X
cardset	X	X	
certreq			
checks	X		
curve	X		
embedconvfile		X	
embedsavefile	X	X	X
from-application			X
from-ident			X
hexdata		X	
ident	X	X	
keystore	X	X	X
keystorepass	X	X	X
module	X	X	
nvram	X	X	
paramsreadfile	X		
pemreadfile		X	
plainname	X	X	X
protect	X	X	
pubexp	X		
qsize	X		
recovery	X	X	
seeintegname			
selfcert			
size	X		
strict	X		
type	X		
x509country	X	X	X
x509dnscommon	X	X	X

Property	generate	import	retarget
x509email	X	X	X
x509locality	X	X	X
x509org	X	X	X
x509orgunit	X	X	X
x509province	X	X	X
xsize	X		

The following table shows which applications are applicable to the different *NAME* options:

Property	custom	embed	hwcrhk	pkcs 11	seeconf	seeinteg	seessl	simple	kpm
alias									
blobsavefile	X								
cardset	X	X	X	X				X	X
certreq				X					
checks	X	X	X	X				X	X
curve	X	X	X	X	X	X		X	
embedconvfile		X							
embedsavefile		X		X					
from-application	X	X	X	X				X	X
from-ident	X	X	X	X				X	X
hexdata	X	X	X	X				X	
ident			X					X	X
keystore									
keystorepass									
module	X	X	X	X			X	X	X
nvram	X	X	X	X				X	
paramsreadfile	X	X	X	X	X	X		X	
pemreadfile	X		X					X	X
plainname	X	X		X	X	X	X	X	X
protect	X	X	X	X	X	X	X	X	X

Property	custom	embed	hwcrhk	pkcs 11	seeconf	seeinteg	seessl	simple	kpm
pubexp	X	X	X	X				X	X
qsize	X	X	X	X				X	X
recovery	X	X	X	X	X	X		X	X
seeintegname	X						X	X	
selfcert				X					
size	X	X	X	X	X	X	X	X	X
strict	X	X	X	X				X	
type	X	X	X	X	X	X	X	X	X
x509country		X							X
x509dnscommon		X							X
x509email		X							X
x509locality		X							X
x509org		X							X
x509orgunit		X							X
x509province		X							X
xsize	X	X	X	X				X	

27. Checking and changing the mode on the HSM

This appendix tells you how to check and change the mode on the nShield HSM. You must change the mode to perform certain configuration tasks.

27.1. Front panel controls

See The front panel interface for a description of the nShield HSM user interface, including the front panel controls.

> ℹ️ We recommend that you use a keyboard to manage the front panel menu options and enter text. See Using a keyboard to control the unit for more information.

27.2. Available modes

The following modes are available:

Operational   The default setting for day-to-day use.

Initialization   Sets the nShield HSM to start in pre-initialization mode. This allows you to use the nShield HSM to create a Security World or add the module to an existing one.

Maintenance   You cannot select this mode manually. It is managed by the nShield HSM and cannot be set by a user.

27.3. Identifying the current mode

You can check the current mode of the nShield HSM:

- At the nShield HSM itself
- By using the `enquiry` command-line utility from a client computer
- By using KeySafe from a client computer

27.3.1. Checking the mode at the nShield HSM

27.3.1.1. The status LED

The nShield HSM Status LED indicates the operational status of the module.

Status LED	Description
On, occasionally blinks off.	**Status: Operational mode** The module is in Operational mode and accepting commands. The more frequently the Status LED blinks off, the greater the load on the module.
Flashes two short pulses, followed by a short pause.	**Status: Initialization mode** Existing Security World data on the module has been erased. The module is automatically placed in Initialization mode after a Security World is created.
Flashes two long pulses followed by a pause.	**Status: Maintenance mode** Used for reprogramming the module with new firmware. The module only goes into Maintenance mode during a software upgrade.

27.3.1.2. The front panel display screen

The nShield HSM screen shows a color-coded footer at the bottom of the display when it is not in Operational mode.

Footer color	Text in footer	Meaning
Yellow	Initialization	The system is rebooting or waiting for an Administrator Card to be inserted.
Blue	Maintenance	An administrative task is being performed. This mode is only entered during firmware upgrades.
Red	HSM Failed	The internal module has failed.

27.3.2. Checking the mode using enquiry

You can use the `enquiry` command-line utility to display information about the hardserver and the status of the nShield HSM. The `enquiry` utility is in the `bin` subdirectory of the `nCipher` directory. This is usually `/opt/nfast` (**Linux**) or `C:\Program Files\nCipher\nfast` (**Windows**)

To check the mode using **enquiry**:

1. Sign in to the client computer as a user, and open a command window.
2. Run the command:

Linux

```
opt/nfast/bin/enquiry
```

Windows

```
enquiry
```

Example output:

```
Server:
enquiry reply flags      none
enquiry reply level      Six
serial number            ####-####-####-####
mode                     operational
version                  #.#.#
speed index              ###
rec. queue               ##..##
...
version serial           #
remote port (IPv4)       ####

Module #1:
enquiry reply flags      none
enquiry reply level      Six
serial number            ####-####-####-####
mode                     operational
version                  #.#.#
speed index              ###
rec. queue               ##..##
...
rec. LongJobs queue      ##
SEE machine type         PowerPCSXF
```

In this example, the `mode` line shows that the nShield HSM is in `operational` mode.

27.3.3. Checking the mode by using KeySafe

You can use the **Module Status tree** of the KeySafe GUI to identify the current mode of the nShield HSM.

To check the mode using KeySafe:

1. Start KeySafe on a client computer.
2. Locate the **Module Status tree** (part of the **Security World status** panel) positioned to the bottom left of the KeySafe window.
3. Expand the **Security World** and/or **Outside Security World** nodes as required.
4. Locate the appropriate nShield HSM (**Module**).
   The current mode of the module is displayed in the **State** field.

See Using KeySafe for more about using KeySafe. See Module information for more about

---

checking the mode.

27.4. Changing the mode

You can change the mode using:

- The front panel controls of the nShield HSM
- The `nopclearfail` command-line utility from a client computer

27.4.1. Changing the mode using the front panel controls

To change the mode, use the front panel menu screens and dialogs to do the following:

1. Navigate to **HSM** > **Set HSM mode**.
2. Select **Initialisation** or **Operational** as required.

27.4.2. Changing the mode using remote mode and nopclearfail

You can enable or disable changing the mode remotely, see `enable_remote_mode` in the `server_settings` section or the *Top-level menu* chapter of the HSM *Install Guide*. Once you have enabled remote mode changes, you can change the mode of the nShield HSM from a computer using the `nopclearfail` command, without accessing the unit itself.

27.4.2.1. Available commands

You can use the following commands to change the mode of a module:

Command	Resulting mode
`nopclearfail --operational | -O`	Operational
`nopclearfail --initialization | -I`	Pre-initialization

To change the mode, do the following:

1. Run either:

   a. The `nopclearfail --operational | -O` command.
      or:

   b. The `nopclearfail --initialization | -I` command.
      When finished, the system responds with `OK`.

> **ℹ** The system responds with `OK`, regardless of whether the mode of the nShield HSM has changed or not. To confirm that state of the module, do the following:

2. Run the `enquiry` command.

   The `mode` line of the `Module` section displays the current mode.

28. Maintenance of nShield Hardware

This chapter describes maintenance steps for your nShield hardware installation.

After installing your nShield HSM by following the *Installation Guide*, Entrust recommend that you use some of the provided software utilities to monitor your installation. Specifically, the `stattree` command allows reporting of voltages and temperatures from your module.

For more information regarding `stattree`, see stattree: information utility.

28.1. Temperature Monitoring for Airflow Validation

Temperatures within a module are monitored to protect against potential attacks, and to prevent overheating. The temperature of the internal ambient air of an nShield HSM is reported as `CurrentTempC` and `CurrentTemp2C` under the `HostEnvStats` node tag of `stattree`.

The table below documents the expected normal operating ranges for the temperatures of your module. Module temperatures would be expected to be within these values when installed with sufficient cooling in an approximately 20-30°C ambient air temperature environment. Calculated `stattree` statistics such as minima and maxima are reset on module reboot.

> ℹ️ The temperatures in this table do not cover operation of the product across the full temperature range specified in the *Warnings & Cautions* and *Installation Guide*. This is because these values are recommendations to ensure a long product lifetime, thus are specified for 20-30°C ambient air operation.

`stattree` **Statistic**	**Description**	**Minimum expected in optimum environment**	**Maximum expected in optimum environment**
`CurrentTempC`	Internal temperature 1	10°C	45°C
`CurrentTemp2C`	Internal temperature 2	10°C	45°C
`MaxTempC`	Maximum of internal temperature 1	-	45°C
`MaxTemp2C`	Maximum of internal temperature 2	-	45°C

> ⚠️ If any of the above temperatures are reporting higher than their specified maximum it is likely your nShield hardware does not have sufficient

cooling. Please refer to the *Installation Guide* to confirm your cooling setup.

29. Upgrading the image file and associated firmware

29.1. Version Security Number (VSN)

Each Connect image file has a Version Security Number (VSN). In addition, the internal module (HSM) firmware has its own individual VSN. This number is increased whenever we improve the security of the image file and/or firmware.

We supply several versions of the module firmware. You can always upgrade to firmware with an equal or higher VSN than that currently installed on your module.

The Version Security Number (VSN) stands as a safeguard to prevent earlier and potentially less secure images and accompanying firmware from being loaded onto the nShield Connect. It prevents the loading of an image file with a lower VSN than the existing VSN. The VSN is not incremented with every release, but only in the event of a significant security enhancement to the nShield Connect and / or its internal module.

> **ℹ** The Connect image VSN is only available from the nShield Connect Front Panel UI.

29.2. Key data

During an upgrade Security World and key data are preserved on the RFS host computer. Once you have upgraded the Connect image you must restore the unit to the Security World if you wish to continue using the key data.

> **ℹ** Adding or restoring a module will require authorisation from a quorum of Administrator cards.

> **ℹ** When upgrading the nShield HSM image file, client licenses and features activations on the HSM will persist. However, if you factory state the unit dynamic features are lost and must be re-enabled.

For more information, see Adding or restoring an HSM to the Security World.

> **!** Ensure you have a quorum of the ACS and that the ACS is available and operating correctly prior to commencing any firmware upgrade. If you do not, you will not be able to reload your Security World on your nShield Connect and you will not be able to use any of your keys.

29.3. Upgrading the Connect image

The Connect image (identified by the file extension .nff) is located on the .iso or DVD in the `nethsm-firmware` directory. The image file contains all necessary components required to upgrade the nShield Connect.

Before upgrading, copy the directory containing the `.nff` file from the .iso or DVD to the `nethsm-firmware` directory on the Remote File System of the nShield Connect.

The following sections describe how to load this firmware package onto your nShield Connect.

29.4. Upgrading the Connect image using the front panel

Before upgrading your Connect image ensure that you have a working quorum of Administrator Cards from the ACS. You need these together with the files in `/local` to restore your Security World on your nShield Connect after the upgrade.

To upgrade the Connect image:

1. Ensure that the Connect image file is named `nCx3N.nff` and located in the following directory:

    ◦ Windows: `%NFAST_HOME%\nethsm-firmware\<version>\nCx3N.nff`.

    ◦ Linux: `/opt/nfast/nethsm-firmware/<version>/nCx3N.nff`.

    Where `<version>` is a subfolder containing the firmware image to be used for the upgrade. There can be more than one `<version>` subfolder.

The directory `<version>` should match the image version string identified on the firmware ISO. If you are not sure on the details, contact Entrust nShield Support, https://nshieldsupport.entrust.com.

1. From the main menu on the unit, select **System** > **Upgrade system**.
2. Confirm that you want to upgrade the image file.
3. Select the directory that contains the image file or firmware that you require. If multiple Connect image directories are displayed, scroll to the relevant directory and select it.

    You are informed that the files are being transferred. The nShield Connect will disconnect from the network during the upgrade procedure and reconnect once the upgrade is complete.

4. Verify the image version, HSM (firmware) version, and image VSN that are displayed,

and confirm the upgrade when prompted.

29.5. Upgrading the nShield Connect from a privileged client

The following description assumes the RFS and Client are separate machines which an nShield HSM has already been configured to use. If you are using a combined RFS/Client, then apply the following instructions to the same machine. The Client must have privileged access to the nShield HSM.

The image upgrade file may be supplied as a separate item that must be copied into the sub folder for its respective version. The default file name is `nCx3N.nff`.

1. Ensure that the new image file is in the following folder on the RFS:

    ◦ Windows: `%NFAST_HOME%\nethsm-firmware\<version>`

    ◦ Linux: `/opt/nfast/nethsm-firmware/<version>`

    Where `<version>` is a subfolder containing the image for the respective version. There can be more than one `<version>` subfolder. The string `<version>` should match the name of the version folder in which the image is located on the version's firmware ISO.

    If the `<version>` subfolder does not already exist on the RFS, it must be created by a user with the necessary privileges.

2. List the image file(s) available on the RFS, run the following command from the Client:

    ```
    >nethsmadmin —m<n> -s <RFS_IP> -l
    ```

    Where:

    ◦ `<n>` is the module number for the target nShield HSM

    ◦ `<RFS_IP>` is the IP address of the RFS.

    ◦ Additionally the `--rfs-hkneti=<RFS_HKNETI>` and `--rfs-esn=<RFS_ESN>` options can be set to enable secure authentication of the RFS. There are three possible cases:

        ▪ **Without secure authentication**: The authentication of the RFS will be based on the IP address only if the `--rfs-hkneti` and `--rfs-esn` options are not specified.

        ▪ **Software-based authentication**: The `--rfs-hkneti` option specifies the software KNETI hash of the RFS. The `--rfs-esn` option shall not be specified.

`<RFS_HKNETI>` can be obtained by running `anonkneti -m0 localhost` on the RFS.

- **nToken authentication**: Only if an nToken (or local HSM) is installed in the RFS. The `--rfs-hkneti` and `--rfs-esn` options specify the KNETI hash and ESN of the nToken.

`<RFS_HKNETI>` and `<RFS_ESN>` can be obtained by running `ntokenenroll -H` on the RFS.

For example, when the image file is located in the appropriately named `<version>` folder:

```
>nethsmadmin -m1 -s 194.28.158.146 -l
Initiating RFS nethsm image check on 194.28.158.146...

Checking the nethsm-firmware directory on the RFS.
nethsm-firmware/VersionName/nCx3N.nff
nethsm-firmware/AnotherVersionName/nCx3N.nff

Images were successfully found on the RFS (194.28.158.146).
```

For example, if the version folder does not exist or its name is not correct, the `nethsmadmin` command cannot find the image:

```
>nethsmadmin -m1 -s 194.28.158.146 -l
Initiating RFS nethsm image check on 194.28.158.146...

Checking the nethsm-firmware directory on the RFS.
No images found on the RFS (194.28.158.146).
```

3. In order to load (or upgrade) the Connect image run the following command from the Client:

```
>nethsmadmin —m<n> -s <RFS_IP> --upgrade-image=nethsm-firmware/<selected-image-version>/nCx3N.nff
```

Where:

- `<n>` is the module number for the target nShield HSM
- `<selected-image-version>` specifies the version subfolder on the RFS containing the firmware image you wish to load (upgrade) onto the nShield HSM.
- `<RFS_IP>` specifies the IP address of the RFS using the `-s` argument. For example

```
>nethsmadmin -m1 -s 194.28.158.14 --upgrade-image=nethsm-firmware/VersionName/nCx3N.nff
```

ℹ️ Copy the path to the required image file as provided by the

> available image list above. (Linux style path separators are used irrespective of whether the Client or RFS are Windows or Linux based).

For example:

```
>nethsmadmin -m1 -s 194.28.158.14 --upgrade-image=nethsm-firmware/VersionName/nCx3N.nff
Initiating appliance image upgrade using file nethsm-firmware/VersionName/nCx3N.nff...
Upgrade operation state changed to: Image Transfer Initiated
Upgrade operation state changed to: Image Transferred
Upgrade operation state changed to: Image Verified
Not able to contact appliance because of reason(23): CrossModule,#1-ExplicitRequest,#2-Mode
Upgrade operation final state: Image Verified
Image upgrade completed.
Please wait for appliance to reboot.
Please wait for approximately half an hour for the appliance to internally upgrade.
```

The following line is expected and requires no action:

```
Not able to contact appliance because of reason(23): CrossModule,#1-ExplicitRequest,#2-Mode
```

The notification appears because the RFS/client cannot contact the nShield Connect. Once the image is copied across, the nShield Connect will disconnect from the network for the duration of the upgrade and reconnect once the upgrade is completed.

> ℹ️ If the nShield HSM suffers a loss of power while you are upgrading the image file or internal module firmware, exit the `nethsmadmin` utility, wait until power is restored to the HSM, then try to restart the process as shown above.

4. After the image upgrade has completed, run the `enquiry` utility to check the image version of the target nShield HSM is as expected.

29.5.1. Enabling and disabling remote upgrade

You can enable or disable upgrading an nShield HSM remotely, see `enable_remote_mode` in the server_settings section or the *Top-level menu* chapter of the HSM *Install Guide*. Once you have enabled remote upgrade, you can upgrade an nShield HSM from a computer using the `nethsmadmin` command, without accessing the unit itself.

29.6. After firmware installation

After you have installed new firmware and initialized the HSM, you can create a new Secu-

rity World with the HSM or reinitialize the HSM into an existing Security World.

If you are initializing the HSM into a new Security World, see Creating a Security World.

If you are re-initializing the HSM into an existing Security World, see Adding or restoring an HSM to the Security World.

30. SNMP monitoring agent

This appendix describes the Simple Network Management Protocol (SNMP) monitoring agent. The SNMP monitoring agent provides you with components that you can add to your (third-party) SNMP manager application.

SNMP was developed in 1988 and revised in 1996. It is currently regarded as the standard method of network management. It is widely supported and offers greater interoperability than traditional network management tools (for example, `rsh` or `netstat`). This makes it ideal for use for the large array of platforms that we support and also avoids the overhead of remote login and execution, helping to reduce network congestion and improve performance.

SNMP defines a collection of network management functions allowing management stations to gather information from, and transmit commands to, remote machines on the network. Agents running on the remote machines can take information gathered from the system and relay this information to the manager application. Such information is either requested from the underlying operating system or gained by interrogating the hardware.

> Every SNMP manager adds monitor components differently. Consult the documentation supplied with your SNMP Manager application for details on how to add the MIB files.

SNMP defines the following SNMP messages:

Message	Description
get	This message is sent by a manager to retrieve the value of an object at the agent.
set	This message is sent by a manager to set the value of an object at the agent.
trap	This message is sent by an agent to notify a management station of significant events.

The SNMP monitoring agent is based on the open-source Net-SNMP project, version 5.7.3. More information on SNMP in general, and the data structures used to support SNMP installations, is available from the NET-SNMP project Web site: https://net-snmp.sourceforge.io/.

This site includes some support information and offers access to discussion e-mail lists. You can use the discussion lists to monitor subjects that might affect the operation or security of the SNMP agent or command-line utilities.

> Discuss any enquiries arising from information on the NET-SNMP Web site with Support before posting potentially sensitive information to the

NET-SNMP Web site.

30.1. Installing and activating the SNMP agent

On Linux, the SNMP agent is installed with the installation of the Security World Software and starts automatically.

On Windows, the SNMP agent can be installed and activated separately. After installing the SNMP components, an activation command can be issued.

30.1.1. Default installation settings

When installing Security World Software, you may be prompted to select Security World Software components from a list. If you select `all` components, then the SNMP agent is installed as part of a full Security World Software installation. The default installation directory for the nShield Management Information Base (MIB) and the SNMP configuration files (`snmp.conf` and `snmpd.conf`) is `/opt/nfast/etc/snmp/` (**Linux**) or `%NFAST_HOME%\etc\snmp\` (**Windows**).

30.1.2. Do you already have an SNMP agent running?

If you already have another SNMP agent running, you must configure the ports used by the agents in order to avoid conflicts before enabling the SNMP agent. A port is assigned by editing the `agentaddress` entry in the `snmpd.conf` file or by editing the `defaultPort` entry in `snmp.conf` file. If both files have been edited, the `agentaddress` entry is `snmpd.conf` file takes priority for snmpd, and the `defaultPort` entry in `snmp.conf` is ignored.

If no existing SNMP agent is found, the SNMP agent runs on the default port 161. If an existing SNMP agent is detected, and no SNMP agent configuration files are found (implying a fresh installation), the installer automatically configures the SNMP agent to use the first unused port above 161 by creating a new `snmpd.conf` configuration file with the appropriate directive. It then displays a message indicating the number of the port that is has selected.

If an existing SNMP agent is found and an existing SNMP agent installation exists, the installer checks the existing configuration files for an appropriate directive and warns you if one does not exist. If you need to edit these configuration files yourself, a port is assigned by editing the `agentaddress` entry in `snmpd.conf` file or editing the `defaultPort` entry in `snmp.conf` file. If both files have been edited, the `agentaddress` entry in `snmpd.conf` file takes priority for snmpd, and the `defaultPort` entry in `snmp.conf` is ignored.

30.1.3. Starting the SNMP agent

Linux

The SNMP agent is started automatically however it can be stopped and started manually. To stop, start, or restart (stop and immediately start again) the SNMP daemon:

```
/opt/nfast/scripts/init.d/ncsnmpd stop|start|restart
```

See The SNMP configuration file: snmp.conf for more information on additional parameters accepted by snmpd.

Windows

To register the SNMP agent as a Windows service, enter the following command with administrative privileges:

```
%NFAST_HOME\bin\snmpd -register [params]
```

See The SNMP configuration file: snmp.conf for more information on additional parameters accepted by snmpd.

This installs the agent as a Windows Service but does not start it automatically.

> **ⓘ** By default, the SNMP agent logs start-up and shut-down to the Event Viewer. More detailed logging can be configured by providing additional parameters when running the SNMP agent either from the command line or when registering as a service.

To unregister the SNMP agent as a Windows service, enter the following command:

```
snmpd -unregister
```

The SNMP agent can be started and stopped from the services control panel or from the command prompt using:

```
net start "nCipher SNMP Agent"
net stop "nCipher SNMP Agent
```

30.2. Basic configuration

30.2.1. Protecting the SNMP installation

The SNMP agent allows other computers on the network to connect to it and make requests for information. The SNMP agent is based on the NET-SNMP code base, which has been tested but not fully reviewed by Entrust. We strongly recommend that you deploy the SNMP agent only on a private network or a network protected from the global Internet by appropriate network protection systems (such as a firewall, a network Intrusion Detection/Prevention System, etc.).

The default nShield SNMP installation allows read-only access to the Management Information Base (MIB). There is no default write access to any part of the MIB.

Every effort has been taken to ensure the confidentiality of cryptographic keys even when the SNMP agent is enabled. In particular, the nShield module is designed to prevent the theft of keys even if the security of the host system is compromised, provided that the Administrator Cards are used only with trusted hosts. Care must be used when changing the configuration of the SNMP agent.

> **!** We strongly advise that you use the SNMP User-based Security Model (USM) with Authentication and Privacy protocols selected, to ensure only authorised users can obtain information from the SNMP agent and the confidentiality and data integrity of the transferred information is protected.

Care has also been taken to ensure that malicious attackers are unable to inundate your module with requests by flooding your SNMP agent. Command results from administration or statistics commands are cached, and thus the maximum rate at which the SNMP agent sends commands to the module is throttled. For more information on setting the cache time-outs, see The SNMP configuration file: snmp.conf.

30.2.2. Configuring the SNMP agent

The Security World Software package uses various configuration files to configure its applications. This section describes the overall nature of the configuration files for the SNMP agent.

If you are installing the SNMP agent to a host that has an existing SNMP agent installation, you may need to edit the SNMP configuration files (`snmpd.conf` and `snmp.conf`) associated with the SNMP agent to change the port on which the agent listens for SNMP requests. For more information, see Do you already have an SNMP agent running?.

> **i** Make sure you protect access to the configuration files, since these contain information that defines the security parameters of the SNMP system. The default location for the nShield SNMP configuration files is

> /opt/nfast/etc/snmp/ (**Linux**) or %NFAST_HOME%\etc\snmp\ (**Windows**).

30.2.3. Create the configuration files (Windows)

On Windows, the snmp.conf and snmpd.conf files are not created automatically by the instal
lation. Instead, example files (example.snmp.conf and example.snmpd.conf) are created in
that location, which you can copy, rename (to snmp.conf and snmpd.conf), and edit with
your desired configuration settings.

> **ℹ** The sample snmpd.conf file includes agentuser and agentgroup direc-
> tives, however these are inoperative in Windows.

> **ℹ** You can override the default search path by setting the environment
> variable SNMPCONFPATH to a colon-separated (":") list of directories for
> which to search.

30.2.3.1. Re-reading SNMP configuration files

The SNMP agent reads its configuration files on startup, and any changes made after this
point will have no effect. If new directives are added and need to be applied, the SNMP
agent can be forced to re-read its configuration files with:

- An snmp set of integer(1) to enterprises.nCipher.reloadCon-
  fig.0(.1.3.6.1.4.1.7682.999.0)
- kill -HUP signal sent to the snmpd agent process
- stop then restart the SNMP agent.

30.2.3.2. The SNMP configuration file: snmp.conf

The snmp.conf configuration file contains directives that apply to all SNMP applications.
These directives can be configured to apply to specific applications. The snmp.conf configu
ration file is not required for the agent to operate and report MIB entries.

30.2.3.3. The SNMP agent configuration file: snmpd.conf

The snmpd.conf configuration file defines how the SNMP agent operates. It is required only
if an agent is running.

The snmpd.conf file can contain any of the directives available for use in the snmp.conf file
and may also contain the following Security World Software-specific directives:

Directive	Description
statstimeout	This directive specifies the length of time for which statistics commands are cached. The default is 5 seconds.
admintimeout	This directive specifies the length of time for which administrative commands are cached. The default is 60 seconds.
keytable	This directive sets the initial state of the key table to none, all, or query. See listKeys in Administration sub-tree overview.
enable_trap_zero_suffix	This directive appends the '.0' suffix to object identifiers (OIDs) for backward compatibility. The default is 0 (disabled): the directive can be set to 1 to restore the suffix. Valid values are 0 and 1.
memoryUsageOkThreshold	This directive specifies the threshold (as a percentage) below which HSM memory usage is considered to be ok. The default is 0. See Memory usage monitoring for more details.
memoryUsageHighThreshold	This directive specifies the threshold (as a percentage) at which HSM memory usage is considered to be too high. The default is 0. See Memory usage monitoring for more details.

ℹ️ There may be a tolerance gap between the memoryUsageOkThreshold and the memoryUsageHighThreshold values.

ℹ️ The timeouts should be set to values that achieve a balance between receiving up to date information whilst preventing excessive load.

30.2.4. The SNMP agent persistent configuration file

On running the SNMP agent for the first time, the persist directory will be created. This contains configuration files that are maintained by the SNMP agent. This directory will be created in /opt/nfast/etc/snmp/persist (**Linux**) or %NFAST_HOME%\etc\snmp\persist (**Windows**).

Modifications should only be made to the persist folder's snmp.conf file in order to create users. The files within this directory should otherwise only be managed by the SNMP agent itself.

User creation can be performed with the createUser directive. See USM users. On initialization of the agent the information is read from the file and the lines are removed (eliminating the storage of the master password for that user) and replaced with the key that is derived from it. This key is a localised key, so that unlike the password, if it is stolen it can not be used to access other agents.

ℹ️ Do not modify the persistent snmpd.conf file while the agent is running.

> The file is only read on initialization of the agent and it is overwritten when the SNMP agent terminates. Any changes made to this file while the SNMP agent directives is running will be lost. The SNMP agent should be stopped prior to adding `createUser` directories to the configuration file.

30.2.5. Agent Behaviour

There are a small number of directives that control the behaviour of the SNMP Agent when considering it as a daemon providing a network service.

30.2.6. agentaddress directive

The listening address(es) that the SNMP Agent will use are defined by the `agentaddress` directive. It takes a comma separated list of address specifiers where an address specifier consists of one or more of:

- a transport specifier udp: or tcp
- a hostname or IPv4 address
- a port number (for example, :161 or :1161).

The default behaviour is to listen on UDP port 161 on all IPv4 interfaces (i.e. equivalent to udp:161).

```
agentaddress localhost : 161,tcp:1161
```

`agentaddress` will listen on UDP port 161, but only on the loopback interface (the port specification ":161" is not strictly necessary as this is the default port). It will also listen on TCP port 1161 on all IPv4 interfaces.

30.2.7. agentgroup and agentuser directives (Linux)

The user and group that the SNMP Agent changes to after opening the listening port(s) are defined using the `agentgroup` and `agentuser` directives. The following must be used:

```
agentgroup ncsnmpd
agentuser ncsnmpd
```

30.2.8. System information (Linux)

Most of the scalar objects in the .iso.org.dod.internet.mgmt.mib-2.system sub-tree can be configured.

```
sysLocation STRING
sysContact STRING
sysName STRING
```

The three directives above set the system location, contact or name for the SNMP Agent respectively. Ordinarily these objects are writable via a suitably authorised SNMP SET request, however, specifying one of these directives in the configuration file makes the cor responding object read-only.

```
sysServices INTEGER
```

Sets the value of the sysService.0 object. RFC1213 defines how the integer value is calculated.

```
sysDescr STRING
sysObjectID OID
```

The two directives above set the system description and object ID for the agent. These objects are not SNMP-writable, but these directives can be used by a network administrator to configure suitable values for them.

30.3. USM users

The SNMPv3 protocol supports a User based Security Model as defined in RFC-3414. USM provides authentication and privacy (encryption) functions and operates at the message level allowing for the following security level to be used with SNMPv3:

- Communication without authentication and privacy (`noauth`)
- Communication with authentication and without privacy (`auth`)
- Communication with authentication and privacy (`priv`).

Within this document the three possible security levels are referred to as `noauth`, `auth` and `priv`. However, other forms are sometimes used within the NET-SNMP and the equivalents are:

Security level	Equivalents
noauth	noauthnopriv
auth	authnopriv

Security level	Equivalents
priv	authpriv

Users can be added to the SNMP configuration with the `createUser` directive, defining the security mechanisms to be used.

```
createUser [-e ENGINEID] username [SHA authpassphrase] [AES privpassphrase]
```

It would not normally be necessary to specify the engine ID, but if it is specified, `ENGINEID` is defined as a hexadecimal string of octets starting with the 0x prefix. The encoding of the engine ID is defined in the description of `SnmpEngineID` from RFC3411.

The following recommendations should be followed when defining the security parameters for SNMPv3:

- Select a 'Security Level' of Priv, (`authpriv`) or auth (`authNoPriv`).
  - `Priv` is the preferred 'Security Level', since this will provide both data source authentication and confidentially protection for the SNMP messages.
  - `auth` is the minimum 'Security Level' that should be selected, since this will ensure that SNMP data sent/received has not been tampered with and has been sent from an authorised entity.
- Define separate `authpassphrase` and `privpassphrase`.
  - It is good security practice to have key separation.
- Use randomly generated passphrases which contain upper and lower case characters, numbers and symbols (for example, ASCII characters 0x20 - 0x7E).
  - This should give an entropy per character of 6.57bits,
- Use either 15 char for 96 bits of security strength keys and 20 char for 128 bits security strength keys.
  - The minimum length of both `Auth` and `Priv` passphrases is eight characters.
  - If a random passphrase is not used, consult NIST SP800-63-2 - Appendix A to determine the security strength of the password and the resultant keys. See https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-2.pdf.

> ℹ MD5 and DES are not supported or enabled in the nShield distribution of SNMP. Only SHA may be used for authentication, and only AES may be used for privacy (encryption).

It is strongly recommended that `createUser` directives be added to the `persist/snmpd.conf` file, so that the passphrases are not available after the SNMP agent is installed. See [n5c-ug:snmp-monitor:::usm-user]. The user can then be referenced in access control direc-

tives(s) after which it can be used.

30.4. Traditional access control

Most simple access control requirements can be specified using the directives `rouser` /`rwuser` (for SNMPv3) or `rocommunity`/`rwcommunity` (for SNMPv1 or SNMPv2c).

```
rouser [-s usm] USERNAME [noauth | auth | priv [OID | -V VIEW [CONTEXT]]
rwuser [-s usm] USERNAME [noath | auth | priv [OID | -V VIEW [CONTEXT]]
```

These directives specify that an SNMPv3 user (USERNAME) will be allowed read-only or read-write access respectively. The default (unspecified) security level is `auth`, which is the recommended minimum security level (see above). It is not recommended to use the usm security level `noauth`, where all SNMP messages are unauthenticated and any tampering of the message cannot be detected. Using `noauth` will reduce the security of the SNMP messages to the level of SNMPv1 or SNMPv2c.

OID restricts access for that user to the subtree rooted at the given OID.

`VIEW` restricts access for that user to the specified View-based Access Control Model (VACM) view name. An optional context can also be specified, or `context` to denote a context prefix. If no context field is specified (or the token * is used), the directive will match all possible contexts. (Contexts are a mechanism within SNMPv3 whereby an agent can support parallel versions of the same MIB objects, referring to different underlying data sets.)

A security model can be specified with `-s SECMODEL` however the default security model `usm` is the only security model which is supported in the nShield distribution of SNMP.

Example:

- Read-only user with access to the full OID tree requiring authentication as a minimum:

```
rouser userl
```

Or

```
rouser -s usm user1 auth .1
```

- Read-only user with access to the nShield MIB allowing unauthenticated requests:

```
rouser user2 noauth .1.3.6.1.4.1.7682
```

- Read-write user with access to the full OID tree requiring authentication as a minimum:

```
rwuser user3
```

Or

```
rwuser user3 auth .iso
```

- Read-write user with access to the snmpVacmMIB subtree requiring authentication and encryption:

```
rwuser user4 priv snmpVacmMIB
```

Or

```
rwuser user4 priv .1.3.6.1.6.3.16
```

```
rocommunity COMMUNITY [SOURCE [ OID | -V VIEW [CONTEXT]]
rwcommunity COMMUNITY [SOURCE [ OID | -V VIEW [CONTEXT]]
```

Specifies an SNMPv1 or SNMPv2c community that will be allowed read-only (`GET` and `GET-NEXT`) or read-write (`GET`, `GETNEXT` and `SET`) access respectively. By default, this will provide access to the full OID tree for such requests, regardless of where they were sent from. `SOURCE` allows access either from a particular range of source addresses, or globally (`default`). A restricted source can either be a specific hostname or address (for example, `localhost` or 127.0.0.1), or a subnet - represented as IP/MASK (for example, 10.10.10.0/255.255.255.0), or IP/BITS (for example, 10.10.10.0/24).

OID `VIEW` and `CONTEXT` are as defined for `rouser` and `rwuser`.

Example:

- Setting up a read-only community named `public` that can be accessed by any user with the community name:

```
rocommmunity public
```

- Setting up a read/write community named `private` that can only be accessed from the machine on which the agent is running:

```
rocommmunity private localhost
```

In each case, only one directive should be specified for a given SNMPv3 user, or community string. It is not appropriate to specify both `rouser` and `rwuser` directives referring to the same SNMPv3 user (or equivalent community settings). The `rwuser` directive provides all the access of `rouser` (as well as allowing `SET` support). The same applies to `rwcommunity` and `rocommunity`.

More complex access requirements (such as access to two or more distinct OID subtrees, or different views for `GET` and `SET` requests) should use VACM configuration directives.

30.5. VACM configuration

The full flexibility of the VACM, for example allowing access to two or more distinct OID sub trees, or different access requirements for reading and writing, is available using four config uration directives - `com2sec`, `group`, `view` and `access`. The directives essentially define who has access and what they have access to using four directives. The first two directives (`com sec2sec` and `group`) define the who, while the last two (`view` and `access`) define the what.

```
Com2sec [-Cn CONTEXT] SECNAME SOURCE COMMUNITY
```

Maps an SNMPv1 or SNMPv2c community string to a security name. As it defines the community and maps it to a security name, `rocommunity`/`rwcommunity` directives are not required when using the directive.

`SECNAME` is the security name to be defined.

`SOURCE` is as defined for the `rocommunity`/`rwcommunity` directives above.

`COMMUNITY` defines the community name to be mapped to the security name. The same com munity string can be specified in several separate directives with different source tokens, and the first source/community combination that matches the incoming request will be selected. Various source/community combinations can also map to the same security name.

`CONTEXT` if defined (using `-Cn`), means that the community string will be mapped to a secu rity name in the named SNMPv3 context. Otherwise the default context ("") will be used.

Example:

Creating three SNMPv1/v2c community names (`private`, `public` and `ltd`), where `private` and `ltd` only allow requests from the machine on which the SNMP Agent is running (note lines beginning with a # in `snmpd.conf` are treated as comments):

```
#        [-Cn CONTEXT] SECNAME     SOURCE    COMMUNITY
```

```
com2sec    ""          sec_private localhost private
com2sec                sec_public  default   public
com2sec                sec_limited localhost ltd
```

```
group GROUP v1 | v2c | usm SECNAME
```

Maps a security name (in the specified security model) into a named group. Several group
directives can specify the same group name, allowing a single access setting to apply to
several users and /or community strings. Note that groups must be set up for the two com-
munity-based models separately - a single `com2sec` directive will typically be accompanied
by two `group` directives.

- `GROUP` is the group name being defined/added to.
- `v1`, `v2c`, or `usm` defines the security model to which the definition relates.
- `SECNAME` is the security (USM) user name or security name defined by `com2sec` to be
  added to the group.

Example:

Creating three groups (`grp_private`, `grp_public`, `grp_limited`) for three USM users (`user1`,
`user2` and `user3`) and the three communities shown in the `com2sec` example above:

```
# GROUP v1|v2c|usm SECNAME
group grp_private v1 sec_private
group grp_private v2c sec_private
group grp_private usm user1

group grp_public v1 sec_public
group grp_public v2c sec_public
group grp_public usm user2

group grp_limited v1 sec_limited
group grp_limited v2c sec_limited
group grp_limited usm user3
```

```
view VNAME included | excluded OID [MASK]
```

Defines a named `view` - a subset of the overall OID tree. This is most commonly a single sub
tree, but several `view` directives can be given with the same view name (`VNAME`), to build up
a more complex collection of OIDs. An optional mask can also be specified, providing a
means of indicating which parts of the OID must be matched.

`VNAME` is the view being modified.

`included | excluded` allows you to define whether the view includes or excludes the subtree,
allowing the definition of a more complex view (for example, by excluding certain sensitive
objects from an otherwise accessible subtree).

MASK is an optional list of hex octets (separated by '.' or ':') whose bits indicate which OID sub-identifiers to match against. So for example if we assume we have on OID with 11 sub-identifiers (`.1.3.6.1.x.y.z.table.entry.column.1`) where the last four relate to a table, an entry, a column and index 1, specifying a MASK value of " `FF.A0` " (i.e. 1111111110100000) maps to this OID as follows:

```
1.3.6.1.x.y.z.table.entry.column.1
1 1 1 1 1 1 1   1    1    0    1
```

i.e. this mask means all parts of the OID except the column must match, therefore defining a view to any column of the first row of the table.

By including and excluding various aspects of the full OID tree, it is possible to define fine grained visibility within a view's definition.

Example:

Creating five views where vw_sysContact only has access to the system.sysContact.0 OID, vw_nCipher only has access to the MIB, vw_global has access to the full OID tree, vw_nCipher_stats only has access to nCipher.nC-series.statistics and vw_nCipher_admin only has access to nCipher.nC-series.administration:

```
#      VNAME           included|excluded   OID                                    [MASK]
view   vw_sysContact   excluded            .1
view   vw_sysContact   included            system.sysContact.0                    FF.80

view   vw_nCipher      excluded            .iso
view   vw_nCipher      included            .1.3.6.1.4.1.7682

view   vw_global       included            .1

view   vw_nCipher_stats excluded           .1
view   vw_nCipher_stats included           enterprises.nCipher.nC-series.statistics

view   vw_nCipher_admin excluded           .1
view   vw_nCipher_admin included           enterprises.nCipher.nC-series.administration
```

```
access GROUP CONTEXT any | v1 | v2c | usm noauth | auth | priv exact | prefix READ WRITE NOTIFY
```

Maps from a group of users/communities (with a particular security model and minimum security level, and specific context) to one of three views, depending on the request being processed.

GROUP is a group name defined by the group directive and specifies the group that access is being defined for.

CONTEXT specifies the context for the access (the default context is the empty string ""). The context of incoming requests must match against the context either exactly or by prefix, as

specified by the choice of `exact` | `prefix` made in this directive.

`any`, `v1`, `v2c`, or `usm` define the security model to which this definition relates.

`noauth` | `auth` | `priv` define the security level to which this definition relates. For `v1` or `v2c` access, this will need to be `noauth` as these protocols do not support authentication.

`exact` | `prefix` specify how `CONTEXT` should be matched against the context of the incoming request, either an exact match to `CONTEXT`, or prefixed by `CONTEXT`.

`READ`, `WRITE` and `NOTIFY` specifies the view to be used for `GET*`, `SET` and `TRAP`/`INFORM` requests (although the `NOTIFY` view is not currently used). The keyword `none` is used if there is to be no access for that type of request.

Example:

Specifying that:

- SNMPv1 requests using the public community only have read access to the enterprises.nCipher.nC-series.statistics subtree,
- SNMPv2c requests using the public community only have read access to the enterprises.nCipher.nC-series.administration.subtree,
- SNMPv3 requests using the user2 USM user, must as a minimum be authenticated, and have read, notify access to the nShield MIB (i.e. enterprises nCipher)
- SNMPv3 requests using the user1 USM user, must as a minimum be authenticated and encrypted, and have read, write and notify access to the full OID tree. Note that since requests must be authenticated and encrypted as a minimum, SNMPv1 and v2c requests using the private community cannot be made even though the community is included in grp_private.
- SNMPv1 and SNMPv2 requests using the ltd community and SNMPv3 requests using the user3 USM user, do not require to be authenticated or encrypted, and have read, write access to the system.sysContact.0 OID.

```
#       GROUP       CONTEXT SECMODEL LEVEL   PREFIX READ           WRITE      NOTIFY
access grp_public  ""      v1       noauth  exact  vw_nCipher_stats none      none
access grp_public  ""      v2c      noauth  exact  vw_nCipher_admin none      none
access grp_public  ""      usm      auth    exact  vw_nCipher       none      vw_nCipher
access grp_private ""      any      priv    exact  vw_global        vw_global
vw_global
access grp_limited ""      any      noauth  exact  vw_sysContact    vw_sysContact
none
```

30.6. Trap Configuration

The distribution of SNMP supports SNMPv1, SNMPv2 and SNMPv3 traps. Control over

these traps is defined with a number of directives:

30.6.1. SNMPv1 and SNMPv2 traps

```
trapcommunity COMMUNITY
```

Defines the default community to be used when sending SNMPv1 or SNMPv2 traps. Note that this directive must be used prior to a `trapsink` or `trap2sink` directive that wishes to use this community.

`COMMUNITY` the community name to be used.

Example:

```
trapcommunity traps
```

```
trapsink HOST [COMMUNITY [PORT]]
trap2sink HOST [COMMUNITY [PORT]]
```

Defines a destination for SNMPv1 or SNMPv2 traps generated by the agent.

`HOST` is an address specifier defining the network target that traps will be sent to. It consists of an optional transport specifier (`udp` (default if not specified) or `tcp`), followed by a hostname or IPv4 address, followed by an optional port number. The address components are separated by colons ":". For example, `localhost` or `tcp:192.168.137.2:163`.

`COMMUNITY` if specified will be the community name used for the traps. If it is not specified, the most recently specified `trapcommunity` will be used.

`PORT` allows for port-number to be defined if it is not present as part of the `HOST` specification. If no port is defined, the default port number of 162 will be used.

When a TCP transport specifier is used the SNMP agent establishes the TCP connection with the trap manager at start-up. Therefore the trap manger must be started before the SNMP agent otherwise an error is reported for the line in the `snmpd.conf` file which defines the trap manager.

Likewise when the TCP connection between the SNMP agent and the trap manager is dropped, traps are lost. Therefore it is inadvisable to use TCP instead of UDP for the transport specifier of trap managers.

If TCP is used for the connection between the SNMP agent and the trap manager and the connection is lost, to re-establish the connection the SNMP agent must be restarted, with

the trap manager running and able to accept a TCP connection from the SNMP agent.

For issues with the trap manager accepting TCP connections from a SNMP agent refer to trap manager documentation.

Example:

```
trap2sink udp:192.168.137.220:162 traps
```

30.6.2. SNMPv3 traps

```
trapsess [SNMPCMD_ARGS]  HOST
```

Defines the configuration for a trap. This is the only way to define SNMPv3 traps and it is an alternative method for defining SNMPv1 and SNMPv2 traps.

`SNMPCMD_ARGS` are arguments that would be used for an equivalent `snmptrap` command. So for example to send an SNMPv3 trap as USM user `user1` with authentication and encryption, the value `-v3 -u user1 -1 priv` would be used.

`HOST` see host definition for `trap2sink` above. Example:

```
trapsess -v3 -u user1 -1 priv udp:192.168.137.220:162
trapsess -v2c -c public 192.168.137.221:162
```

30.7. Using the SNMP agent with a manager application

The nShield SNMP monitoring agent provides MIB files that can be added to your (third-party) SNMP manager application.

30.7.1. Manager configuration

The manager application is the interface through which the user is able to perform network management functions. A manager communicates with agents using SNMP primitives (`get`, `set`, `trap`) and is unaware of how data is retrieved from, and sent to, managed devices. This form of encapsulation creates the following:

- The manager is hidden from all platform specific details
- The manager can communicate with agents running on any IP-addressable machine.

As a consequence, manager applications are generic and can be bought off the shelf. You

may already be running SNMP managers, and if so, you can use them to query the SNMP agent.

> ℹ️ The manager is initially unaware of the MIB tree structure at a particular node. Managed objects can be retrieved or modified, but only if their location in the tree is known.

It is more useful if the manager can *see* the MIB tree present at each managed node. The MIB module descriptions for a particular node must be parsed by a manager-specific MIB compiler and converted to configuration files. These files are read by the manager application at run time.

The SNMP agent is designed to monitor all current nShield modules, working with all supported versions of nShield firmware (contact Support for details of supported firmware).

30.7.2. MIB module overview

A large proportion of the SNMP system is fully specified by the structure of the MIB; the behavior of the agent depends on relaying information according to the layout of the MIB.

The MIB module resides at a registered location in the MIB tree determined by the Internet Assigned Numbers Authority (IANA). The private enterprise number of 7682 designated by the IANA corresponds to the root of the branch, and by convention this (internal) node is the company name.

The MIB module groups logically related data together, organizing itself into a classification tree, with managed objects present at leaf nodes. The nC-series node (`enterprises.nCipher.nC-series`) is placed as a sub-tree of the root (`enterprises.nCipher`); this allows future product lines to be added as additional sub-trees. The structure of the tree underneath the registered location is vendor-defined, and this specification defines the structure chosen to represent Security World Software-specific data.

The MIB file is /opt/nfast/etc/snmp/mibs/ncipher-mib.txt (**Linux**) or `%NFAST_HOME%\etc\snmp\mibs\ncipher-mib.txt` (**Windows**).

30.7.3. MIB functionality

The MIB module separates module information into the following categories:

- Retrieval of status and information about installed nC-series modules
- Retrieval of live statistics of performance of installed nC-series modules

These categories form the top-level nodes of the sub-tree; the functionality of the first category is in the administration sub-tree, and the second category is in the statistics sub-tree. The top-level tree also contains three items that it would be useful to check at-a-glance:

Node name	R/W	Type	Remarks
hardserverFailed	R	TruthValue	True if the remote hardserver is not running. If the hardserver is not running, then most of the rest of the information is unreliable or missing.
modulesFailed	R	TruthValue	True if any modules have failed.
load	R	Unsigned32	Percentage of total available capacity currently utilized.

30.7.3.1. Traps

The traps sub-tree (`enterprises.nCipher.nC-series.nC-traps`) contains traps that the SNMP agent sends when certain events occur. For details on configuring traps, see USM users.

The following table gives details of the individual traps:

Node name	Description
hardserverAlert	This trap is sent when the hardserver fails or is shut down.
hardserverUnAlert	This trap is sent when the hardserver restarts.
moduleAlert	This trap is sent when a module fails.
moduleUnAlert	This trap is sent when a module is restarted after a failure.
psuAlert	This trap is sent when a PSU fails.
psuUnAlert	This trap is sent when a previously-failed PSU is working again.
fanfailureAlert	This trap is sent when a fan fails.
fanfailureUnAlert	This trap is sent when a previously-failed fan is working again.
memoryUsageHighAlert	This trap is sent when the HSM memory usage high threshold has been reached or exceeded by an HSM. See section on Memory usage monitoring below for more details.
memoryUsageOkAlert	This trap is sent when the memory usage for an HSM falls below the HSM memory usage ok threshold. See section on Memory usage monitoring below for more details.

ⓘ | Some traps can take up to five minutes to be received.

> ℹ️ Other generic Net-SNMP traps may also be received. These include the two below, see Net-SNMP project website for more details.

Net-SNMP trap name	Description
`SNMPv2-MIB::coldStart`	This trap is sent when the SNMP agent is started
`NET-SNMP-AGENT-MIB::nsNotifyShutdown`	This trap is sent when the SNMP agent is stopped

30.7.4. Memory usage monitoring

The HSM memory usage thresholds and memory usage traps provide a mechanism to monitor HSM memory usage for HSMs in which the SNMP agent's client computer are enrolled.

With memory usage monitoring enabled, there will be a `memoryUsageHighAlert` trap sent each time the currently in-use `memoryUsageHighThreshold` is reached or exceeded by an HSM.

With memory usage monitoring enabled, a `memoryUsageHighAlert` trap is also sent:

- If the SNMP agent starts up and recognises that there are HSMs in a high memory usage state or,
- If HSMs in a high memory usage state are enrolled or,
- If the SNMP agent loses and then re-gains contact with the local hardserver which is connected to HSMs in a high memory usage state or,
- If failed HSMs in a high memory usage state then recover.

For each of the four scenarios above, one `memoryUsageHighAlert` trap will be sent for each HSM in a high memory usage state.

With memory usage monitoring enabled, there will be a `memoryUsageOkAlert` trap sent each time the memory usage for an HSM falls below the currently in-use `memoryUsageOkThreshold`.

The value for `memoryUsageOkThreshold` is read from the snmpd.conf file on starting the SNMP agent and is used provided it contains an integer value in the range 0 to 100 (inclusive); otherwise, the default value of 0 is used. The value for `memoryUsageHighThreshold` is processed in the same way.

Memory usage monitoring is enabled unless the in-use values for `memoryUsageOkThreshold` and `memoryUsageHighThreshold` are both 0 or the in-use values are such that `memoryUsageOkThreshold` > `memoryUsageHighThreshold`.

For example, in `snmpd.conf`, if `memoryUsageOkThreshold` is assigned an invalid value and `memo`

`ryUsageHighThreshold` is assigned a valid value of say 75%, then memory usage monitoring will be enabled and the values 0% and 75% will be used respectively.

An example of memory usage monitoring by an SNMP agent on a client computer enrolled with 2 HSMs is given below:

30.7.5. Administration sub-tree overview

The administration sub-tree (`enterprises.nCipher.nC-series.administration`) contains information about the permanent state of the hardserver and the connected modules. It is likely that most of the information in this branch rarely changes over time, unlike the `statistics` branch. The information given in the administration sub-tree is mostly acquired by the `NewEnquiry` command and is supplied both per-module and (where appropriate) aggregated over all modules.

The following table gives details of the individual nodes in the administration sub-tree:

Node name	R/W	Type	Remarks
`hardserverRunning`	R	Enum 1: Running 2: NotRunning	This variable reflects the current state of the hardserver (`Running` or `NotRunning`).
`noOfModules`	R	Gauge32	Number of nC-series modules.
`hsVersion`	R	DisplayString	Hardserver version string.
`globalSpeedIndex`	R	Gauge32	Number of 1024-bit signatures each second.

Node name	R/W	Type	Remarks
globalminQ	R	Gauge32	Minimum recommended queue.
globalmaxQ	R	Gauge32	Maximum recommended queue.
SecurityWorld	R	TruthValue	True if a Security World is installed and operational.
swState	R	DisplayString	Security World display flags, as reported by nfkminfo.
listKeys	R/W	Integer 1: none 2: all 3: query 4: resetquery	Controls the behavior of the key table (switch off, display all keys, enable individual attribute queries, clear the query fields). Displaying all keys can result in a very long list.
serverFlags	R	DisplayString	Supported hardserver facilities (the NewEnquiry level 4 flags).
remoteServerPort	R	Gauge32	TCP port on which the hardserver is listening.
swGenTime	R	DisplayString	Security World's generation time.
swGeneratingESN	R	DisplayString	ESN of the module that generated the Security World.

listKeys can be preset using the keytable config directive in snmpd.conf file (see The SNMP configuration file: snmp.conf).

30.7.5.1. Security World hash sub-tree

The following table gives details of the nodes in the Security World hash sub-tree (enterprises.nCipher.nC-series.administration.swHashes):

Node name	R/W	Type	Remarks
hashKNSO	R	MHash	Hash of the Security Officer's key.
hashKM	R	MHash	Hash of the Security World key.
hashKRA	R	MHash	Hash of the recovery authorization key.
hashKRE	R	MHash	Hash of the recovery key pair.
hashKFIPS	R	MHash	Hash of the FIPS authorization key.

Node name	R/W	Type	Remarks
hashKMC	R	MHash	Hash of the module certification key.
hashKP	R	MHash	Hash of the passphrase replacement key.
hashKNV	R	MHash	Hash of the nonvolatile memory (NVRAM) authorization key.
hashKRTC	R	MHash	Hash of the Real Time Clock authorization key.
hashKDSEE	R	MHash	Hash of the SEE Debugging authorization key.
hashKFTO	R	MHash	Hash of the Foreign Token Open authorization key.

30.7.5.2. Security World quorums sub-tree

The following table gives details of the nodes in the Security World quorums sub-tree (`enterprises.nCipher.nC-series.administration.swQuorums`):

Node name	R/W	Type	Remarks
adminQuorumK	R	Gauge32	The default quorum of Administrator cards.
adminQuorumN	R	Gauge32	The total number of cards in the ACS.
adminQuorumM	R	Gauge32	The quorum required for module reprogramming.
adminQuorumR	R	Gauge32	The quorum required to transfer keys for OCS replacement.
adminQuorumP	R	Gauge32	The quorum required to recover the passphrase for an Operator card.
adminQuorumNV	R	Gauge32	The quorum required to access non volatile memory (NVRAM).
adminQuorumRTC	R	Gauge32	The quorum required to update the Real Time Clock.
adminQuorumDSEE	R	Gauge32	The quorum required to view full SEE debug information.

Node name	R/W	Type	Remarks
adminQuorumFTO	R	Gauge32	The quorum required to use a Foreign Token Open Delegate Key.

30.7.5.3. Module administration table

The following table gives details of the nodes in the module administration table (`enter-prises.nCipher.nC-series.administration.moduleAdminTable`):

Node name	R/W	Type	Remarks
moduleAdminIndex	R	Gauge32	Module number of this row in the table.
mode	R	Integer 1: Operational 2: Pre-init 3: Init 4: Pre-maint 5: Maint 6: AccelOnly 7: Failed 8: Unknown	Current module state.
fwVersion	R	DisplayString	Firmware version string.
speedIndex	R	Gauge32	Speed index (approximate number of 1024-bit modulo exponentiation operations possible per second) of module
minQ	R	Gauge32	Module minimum recommended queue length
maxQ	R	Gauge32	Module maximum recommended queue length
serialNumber	R	DisplayString	Module Electronic Serial Number (ESN).
productName	R	DisplayString	
hwPosInfo	R	DisplayString	Hardware bus/slot info (such as PCI slot number).

Node name	R/W	Type	Remarks
moduleSecurityWorld	R	TruthValue	Indicates whether or not the module is in the current SW.
smartcardState	R	DisplayString	Description of smart card in slot (empty, unknown card, admin/operator card from current SW, failed). N/A for acceleration only modules.
moduleSWState	R	Integer 1: Unknown 2: Usable 3: MaintMode 4: Uninitialized 5: Factory 6: Foreign 7: AccelOnly 8: Failed 9: Unchecked 10: InitMode 11: PreInitMode 12: Unverified 13: UnusedTableEntry	Current module and Security World state.
moduleSWFlags	R	DisplayString	Security World flags for this module.
hashKML	R	MHash	Hash of the module's secret key.
moduleFeatures	R	DisplayString	Features enabled on this module.
moduleFlags	R	DisplayString	Like serverFlags, but for each module.
versionSerial	R	Gauge32	Firmware Version Security Number (VSN); see Version Security Number (VSN).
hashKNETI	R	MHash	K_{NETI} hash, if present.
longQ	R	Gauge32	Max. rec. long queue.

Node name	R/W	Type	Remarks
connectionStatus	R	DisplayString	Connection status (for imported modules).
connectionInfo	R	DisplayString	Connection information (for imported modules).
machineTypeSEE	R	DisplayString	SEE machine type.

30.7.5.4. Slot administration table

The following table gives details of the nodes in the slot administration table (`enter-prises.nCipher.nC-series.administration.slotAdminTable`):

Node name	R/W	Type	Remarks
slotAdminModuleIndex	R	Integer32	Module number of the module containing the slot.
slotAdminSlotIndex	R	Integer32	Slot number (1-based, unlike nCore which is 0-based).
slotType	R	Integer 1: Datakey 2: Smart card 3: Emulated 4: Soft token 5: Unconnected 6: Out of range 7: Unknown	Slot type.
slotFlags	R	DisplayString	Flags referring to the contents of the slot (from `slotinfo`).

Node name	R/W	Type	Remarks
slotState	R	Integer 1: Unused 2: Empty 3: Blank 4: Administrator 5: Operator 6: Unidentified 7: Read error 8: Partial 9: Out of range	Partial refers to cards in a partially-created card set.
slotListFlags	R	DisplayString	Flags referring to attributes of the slot (from getslotlist).
slotShareNumber	R	Gauge32	Share number of card currently in slot, if present.
slotSharesPresent	R	DisplayString	Names of shares present in card currently in slot.

30.7.5.5. Card set administration table

The following table gives details of the nodes in the card set administration table (enterprises.nCipher.nC-series.administration.cardsetAdminTable):

Node name	R/W	Type	Remarks
hashKLTU	R	MHash	Hash of the token protected by the card set.
cardsetName	R	DisplayString	
cardsetK	R	Gauge32	Required number of cards in the card set.
cardsetN	R	Gauge32	Total number of cards in the card set.
cardsetFlags	R	DisplayString	Other attributes of the card set.
cardsetNames	R	DisplayString	Names of individual cards, if set.

Node name	R/W	Type	Remarks
cardsetTimeout	R	Gauge32	Token time-out period, in seconds, or 0 if none.
cardsetGenTime	R	DisplayString	Generation time of card set.

30.7.5.6. Key administration table

The key administration table is visible as long as the `listKeys` node in the administration sub-tree is set to a value other than `none`.

The following table gives details of the nodes in the key administration table (`enter-prises.nCipher.nC-series.administration.keyAdminTable`):

Node name	R/W	Type	Remarks
keyAppname	R	DisplayString	Application name.
keyIdent	R	DisplayString	Name of key, as generated by the application.
keyHash	R	MHash	
keyRecovery	R	Integer 1: Enabled 2: Disabled 3: No key 4: Unknown 5: Invalid 6: Unset	The value `unset` is never returned by the key table. If you set the value `unset`, the keys are not filtered based on any of the attributes.
keyProtection	R	Integer 1: Module 2: Cardset 3: No key 4: Unknown 5: Invalid 6: Unset	The value `unset` is never returned by the key table. If you set the value `unset`, the keys are not filtered based on any of the attributes.

Node name	R/W	Type	Remarks
keyCardsetHash	R	MHash	Hash of the card set protecting the key, if applicable.
keyFlags	R	DisplayString	Certificate and public key flags.
keyExtraEntries	R	Gauge32	Number of extra key attributes.
keySEEInteg	R	DisplayString	SEE integrity key, if present.
keyGeneratingESN	R	DisplayString	ESN of the module that generated the key, if present.
keyTimeLimit	R	Gauge32	Time limit for the key, if set.
keyPerAuthUseLimit	R	Gauge32	Per-authentication use limit for the key.

30.7.5.7. Key query sub-tree

The key query sub-tree is used if the `listKeys` node in the administration sub-tree is set to `query`.

If these values are set, they are taken as required attributes for filtering the list of available keys; if multiple attributes are set, the filters are combined (AND rather than OR).

The following table gives details of the nodes in the key query sub-tree (`enterprises.nCipher.nC-series.administration.keyQuery`):

Node name	R/W	Type	Remarks
keyQueryAppname	R/W	DisplayString	Application name.
keyQueryIdent	R/W	DisplayString	Name of key, as generated by the application.
keyQueryHash	R/W	DisplayString	
keyQueryRecovery	R/W	Integer 1: Enabled 2: Disabled 3: No key 4: Unknown 5: Invalid 6: Unset	The value `unset` is never returned by the key table. If you set the value `unset`, the keys are not filtered based on any of the attributes.

Node name	R/W	Type	Remarks
keyQueryProtection	R/W	Integer 1: Module 2: Cardset 3: No key 4: Unknown 5: Invalid 6: Unset	The value unset is never returned by the key table. If you set the value unset, the keys are not filtered based on any of the attributes.
keyQueryCardsetHash	R/W	DisplayString	Hash of the card set protecting the key, if applicable.
keyQueryFlags	R/W	DisplayString	Certificate and public key flags.
keyQueryExtraEntries	R/W	Gauge32	Number of extra key attributes.
keyQuerySEEInteg	R/W	DisplayString	SEE integrity key, if present.
keyQueryGeneratingESN	R/W	DisplayString	ESN of the module that generated the key, if present.
keyQueryTimeLimit	R/W	Gauge32	Time limit for the key, if set (0 for no limit).
keyQueryPerAuthUseLimit	R/W	Gauge32	Per-authentication use limit for the key (0 for no limit).

30.7.6. Statistics sub-tree overview

The statistics sub-tree (enterprises.nCipher.nC-series.statistics) contains rapidly changing information about such topics as the state of the nShield modules, the work they are doing, and the commands being submitted.

> Do not rely on information returned from the agent to change instantaneously on re-reading the value. To avoid loading the nShield module with multiple time-consuming statistics commands, the agent can choose to cache the values over a specified period. You can configure this period in the agent configuration file see Basic configuration.

30.7.6.1. Statistics sub-tree

The following table gives details of the nodes in the statistics sub-tree, and the module sta-

tistics table (`enterprises.nCipher.nC-series.statistics.moduleStatsTable`):

Node name	R/W	Type	Remarks
moduleStatsIndex	R	Integer	Module number of this row (for `moduleStatsTable`).
hsuptime	R	Counter32	Uptime of the hardserver.
cmdCountAll	R	Counter32	Returned aggregated for all modules and all commands.
cmdBytesAll	R	Counter32	
cmdErrorsAll	R	Counter32	Returned as for `cmdCount`, returned value is combined module errors added to hardserver marshalling/unmarshalling errors.
replyCountAll	R	Counter32	
replyBytesAll	R	Counter32	
replyErrorsAll	R	Counter32	See notes above for `cmdErrors`.
clientCount	R	Gauge32	
maxClients	R	Counter32	
deviceFails	R	Counter32	
deviceRestarts	R	Counter32	
outstandingCmds	R	Counter32	Total number of outstanding commands over all modules.
load[All]	R	Counter32	

30.7.6.2. Module statistics table

The following table gives details of the nodes in the module statistics table (`enterprises.nCipher.nC-series.statistics.moduleStatsTable`):

Node name	R/W	Type	Remarks
moduleStatsIndex	R	Integer	Module number of this row (for `moduleStatsTable`).
uptime	R	Counter32	Uptime of the module.
cmdCount	R	Counter32	Returned aggregated for all commands.

Node name	R/W	Type	Remarks
cmdBytes	R	Counter32	
cmdErrors	R	Counter32	Returned as for cmdCount all the different error states aggregated into one counter.
replyCount	R	Counter32	
replyBytes	R	Counter32	
replyErrors	R	Counter32	See notes above for cmdErrors.
loadModule	R	Counter32	
loadModule	R	Counter32	
objectCount	R	Gauge32	
clock	R	DisplayString	Depending on the module settings, this can require K_{NSO} permissions to read (and therefore depend on the installation parameters of the agent).
nvRAMInUse	R	Gauge32	
volatileRAMInUse	R	Gauge32	
tempSP	R	DisplayString	
currentCPUTemp1	R	DisplayString	
currentCPUTemp2	R	DisplayString	
currentFanSpeed	R	DisplayString	
currentFanDuty	R	DisplayString	
CPUVoltage1	R	DisplayString	
CPUVoltage2	R	DisplayString	
CPUVoltage3	R	DisplayString	
CPUVoltage4	R	DisplayString	
CPUVoltage5	R	DisplayString	
CPUVoltage6	R	DisplayString	
CPUVoltage7	R	DisplayString	
CPUVoltage8	R	DisplayString	
CPUVoltage8	R	DisplayString	

Node name	R/W	Type	Remarks
CPUVoltage9	R	DisplayString	
CPUVoltage10	R	DisplayString	
CPUVoltage11	R	DisplayString	
nvmFreeSpace	R	Counter32	Free space available on the HSM's NVRAM, in bytes
nvmWearLevel	R	DisplayString	Wear level of the HSM's NVRAM
nvmWornBlocks	R	DisplayString	Worn blocks in the HSM's NVRAM

30.7.6.3. nShield HSM statistics table

The following table gives details of the nodes in the nShield HSM statistics table (`enterprises.nCipher.nC-series.statistics.netHSMStatsTable`):

Node name	R/W	Type	Remarks
netHSMStatsIndex	R	Integer	Table index (not module number).
netHSMUptime	R	Counter32	Host system uptime.
netHSMCPUUsage	R	Gauge32	CPU usage of unit host processor.
netHSMUserMem	R	Gauge32	Total user memory of unit.
netHSMKernelMem	R	Gauge32	Total kernel memory of unit.
netHSMCurrentTemp	R	DisplayString	Internal unit temperature (sensor 1).
netHSMMaxTemp	R	DisplayString	Maximum recorded temperature (sensor 1).
netHSMCurrentTemp2	R	DisplayString	Internal unit temperature (sensor 2).
netHSMMaxTemp2	R	DisplayString	Maximum recorded temperature (sensor 2).
netHSMVoltage5V	R	DisplayString	unit 5V power reading.
netHSMVoltage12V	R	DisplayString	unit 12V power reading.
netHSMFan1Speed	R	Gauge32	Fan 1 speed (RPM).
netHSMFan2Speed	R	Gauge32	Fan 2 speed (RPM).
netHSMFan3Speed	R	Gauge32	Fan 3 speed (RPM).
netHSMIPAddress	R	IpAddress	IP address of unit.

Node name	R/W	Type	Remarks
netHSMDescription	R	DisplayString	Textual description of module (for example, `Local module 3`).
netHSMFan4Speed	R	Gauge32	Fan 4 speed (RPM).
netHSMVoltage3p3V	R	DisplayString	3.3V Supply Rail Voltage
netHSMCurrent3p3V	R	DisplayString	3.3V Supply Rail Current
netHSMCurrent5V	R	DisplayString	5V Supply Rail Current
netHSMCurrent12V	R	DisplayString	12V Supply Rail Current
netHSMVoltage5VSB	R	DisplayString	5V Supply Rail Voltage (Standby)
netHSMCurrent5VSB	R	DisplayString	5V Supply Rail Current (Standby)
netHSMTamperBattery1	R	DisplayString	Voltage of Tamper Battery 1
netHSMTamperBattery2	R	DisplayString	Voltage of Tamper Battery 2
netHSMPSUfailure	R	TruthValue	Power Supply failure status

30.7.6.4. Per connection statistics table

The following table gives details of the nodes in the per connection statistics table (`enterprises.nCipher.nC-series.statistics.connStatsTable`):

Node name	R/W	Type	Remarks
connStatsIndex	R	Integer32	Index of this entry.
connNumber	R	Integer32	Hardserver connection number.
connUptime	R	Counter32	Uptime of the connection.
connCmdCount	R	Counter32	Number of commands submitted through this connection.
connCmdBytes	R	Counter32	Number of bytes submitted through this connection.
connCmdErrors	R	Counter32	Number of marshalling errors on commands through this connection.
connReplyCount	R	Counter32	Number of replies received by this connection.
connReplyBytes	R	Counter32	Number of bytes received by this connection.

Node name	R/W	Type	Remarks
connReplyErrors	R	Counter32	Number of marshalling errors on replies through this connection.
connDevOutstanding	R	Gauge32	Number of commands outstanding on this connection.
connQOutstanding	R	Gauge32	Number of commands outstanding in the hardserver queue.
connLongOutstanding	R	Gauge32	Number of long jobs outstanding for this connection.
connRemoteIPAddress	R	IpAddress	IP Address of connection client.
connProcessID	R	Integer32	Process identifier reported by connection client.
connProcessName	R	DisplayString	Process name reported by connection client.
connObjectTotal	R	Gauge32	The total object count for a connection

30.7.6.5. Module/connection statistics table

The following table gives details of the nodes in the per module, per connection statistics table (`enterprises.nCipher.nC-series.statistics.connModuleStatsTable`).

Node name	R/W	Type	Remarks
connModuleStatsConnId	R	Integer	Identity of this connection
connModuleStatsModuleIndex	R	Integer	Index of the module entry
connModuleStatsObjectCount	R	Gauge32	The object count on this module for this connection

30.7.6.6. Fan table

The fan table provides the speeds of each fan on the remote module (HSM). The following table gives details of the nodes in the fan table (`enterprises.nCipher.softwareVersions.netHSMFanTable`):

Node name	R/W	Type	Remarks
netHSMModuleIndex	R	Integer32	Module number
netHSMFanIndex	R	Integer32	Fan number

Node name	R/W	Type	Remarks
netHSMFanSpeed	R	Gauge32	Fan speed (RPM)

30.7.6.7. Software versions table

The following table gives details of the nodes in the software versions table (`enter-prises.nCipher.softwareVersions.softwareVersionsTable`):

Node name	R/W	Type	Remarks
compIndex	R	Integer	Table index.
compName	R	DisplayString	Component name.
compOutput	R	Component output name	Component name.
compMajorVersion	R	Gauge	
compMinorVersion	R	Gauge	
compPatchVersion	R	Gauge	
compRepository	R	DisplayString	Repository name.
compBuildNumber	R	Gauge	

30.8. SNMP agent command-line

30.8.1. SNMP agent (snmpd) switches

The SNMP agent that binds to a port and awaits requests from SNMP management software is snmpd. Upon receiving a request, snmpd processes the request, collects the requested information and/or performs the requested operation(s) and returns the information to the sender.

The SNMP agent supports a limited subset of command line switches that can be specified when starting the agent.

Usage

```
snmpd [-h] [-v] [-f] [-a] [-d] [-V] [-P PIDFILE):] [-q] [-D] [-p NUM] [-L] [-l LOGFILE] [-r]
```

This command can take the following options:

Option	Description
`-h`	This option displays a usage message.
`-H`	This option displays the configuration file directives that the agent understands.
`-v`	This option displays version information.
`-f`	This option specifies not forking from the calling shell.
`-a`	This option specifies logging addresses.
`-A`	This option specifies that warnings and messages should be appended to the log file rather than truncating it.
`-d`	This option specifies the dumping of sent and received UDP SNMP packets.
`-V`	This option specifies verbose display.
`-P`	PIDFILE This option specifies the use of a file (PIDFILE) to store the process ID.
`-q`	This option specifies that information be printed in a more easily parsed format (quick print).
`-D`	This option turns on debugging output.
`-p`	NUM This option specifies running on port NUM instead of the default: 161.
`-c`	CONFFILE This option specifies reading CONFFILE as a configuration file.
`-C`	This option specifies that the default configuration files not be read.
`-L`	This option prints warnings and messages to stdout and err.
`-s`	This option logs warnings/messages to syslog.
`-r`	This option specifies not exiting if root-only accessible files cannot be opened.
`-I`	[-]INITLIST This option specifies a list of MIB modules to initialize (or not). Run `snmpd` with the `-Dmib_init` option for a list.
`-l`	LOGFILE This option prints warnings/messages to a file LOGFILE (by default, `LOG-FILE=log/snmpd.log`).

30.8.2. Using the SNMP command-line utilities

As an alternative to using an SNMP manager application, we supply several command-line utilities to test your SNMP installation and enable you to obtain information about your nShield module from the SNMP agent. These utilities support the `-h` (display a usage message) as described in the table above.

Utility	Description
`snmptest`	This utility monitors and manages SNMP information.

Utility	Description
snmpget	This utility runs a single GET request to query for SNMP information on a network entity.
snmpset	This utility runs a single SET request to set SNMP information on a network entity.
snmpgetnext	This utility runs a single GET NEXT request to query for SNMP information on a network entity.
snmptable	This utility obtains and prints an SNMP table.
snmptranslate	This utility translates SNMP object specifications into human-readable descriptions.
snmpwalk	This utility communicates with a network entity using repeated GET NEXT requests.
snmpbulkwalk	This utility communicates with a network entity using BULK requests.

> 🛈 These tools are general purpose SNMP utilities and are configurable for use with other SNMP agents. For more information on configuring and using these tools, refer to the NET-SNMP project Web site: `http://net-snmp.sourceforge.net/`.

31. Morse code error messages

If a Hardware Security Module (HSM) encounters an unrecoverable error, it enters the error state. In the error state, the module does not respond to commands and does not write data to the bus.

The blue Status LED flashes the Morse distress code (SOS: three short pulses, followed by three long pulses, followed by three short pulses). The Morse distress code is followed by one of the error codes listed in the tables shown in this guide.

For internal security modules running firmware 2.61.2 and above, the error code listed in this chapter is also reported by the `enquiry` utility in the `hardware status field` of the `Module` and under `hardware errors` in the hardserver log.

Errors are a rare occurrence. If any module goes into the error state, except as a result of you issuing the `Fail` command, contact Support, and give full details of your set up and the error code.

Contact Support even if you successfully recover from the error by taking the recommended action. For troubleshooting information, see the relevant *Installation Guide* for your module type.

31.1. Reading Morse code

The following guidelines are useful when reading Morse code messages from the module:

- The duration of a dash (`-`) is 3 times the duration of a dot (`.`).
- The gap between components of a letter has the same duration as a dot.
- The gap between letters has the same duration as a dash.
- The duration of the gap between repeated series of letters (a Morse code word gap) is 7 times the duration of a dot.

The following table shows the error codes corresponding to numerals.

Numeral	Morse
1	. - - - -
2	. . - - -
3	. . . - -
4	. . . . -
5	. . . . .

Numeral	Morse
6	- . . . .
7	- - . . .
8	- - - . .
9	- - - - .
0	- - - - -

31.2. Runtime library errors

Memory failures can occur if the module is exposed to excessive heat. If you experience these errors, check the ventilation around the module. The module generates considerable heat and, if not well ventilated, may be operating at too high a temperature, even if the rest of your server room is at an appropriate temperature.

The runtime library error codes could be caused by firmware bugs or by faulty hardware. If any of these errors is indicated, reset the module.

Code				Meaning
O L C	- - -	. - . .	- . - .	SIGABRT: assertion failure and/or `abort()` called.
O L D	- - -	. - . .	- . .	Interrupt occurred when disabled.
O L E	- - -	. - . .	.	SIGSEGV: access violation.
O L I	- - -	. - . .	. .	SIGSTAK: out of stack space.
O L J	- - -	. - . .	. - - -	SIGFPE: unsupported arithmetic exception (such as division by 0).
O L K	- - -	. - . .	- . -	SIGOSERROR: runtime library internal error.
O L N	- - -	. - . .	- .	SIGFATALPANIC: error in error handling code.

Codes `OLD`, and `OLE` are more likely to indicate a hardware problem than a firmware problem.

To reset a unit that is in an error state, turn off the unit and then turn it on again.

31.3. Hardware driver errors

In general, the hardware driver error codes described in the following table indicate that some form of automatic hardware detection has failed. As well as indicating simple hardware failure, one of these error codes could indicate that there is a bug in the firmware or

that the wrong firmware has been loaded.

> In the following table, the symbol "#" stands for a given numeral's Morse code representation.

If any of these errors is indicated, contact support.

Code					Meaning
HL	. . . .	. - . .			M48T37 NVRAM (or battery) failed
H B	. . . .	- . . .			Debug serial port initialization failed.
H C	. . . .	- . - .			Processing thread initialization failed.
HCP	. . . .	- . - .	. - - .		Card poll thread initialization failed.
H D	. . . .	- . .			Failure reading unique serial number.
H E	. . . .	.			EEPROM failed on initialization.
HF	. . . .	. . - .			Starting up crypto offload.
H I	. . . .	. .			Interrupt controller initialization failed.
H M	. . . .	- -			System hardware initialization failed.
H O	. . . .	- - -			Token interface initialization failed.
H R	. . . .	. - .			Random number generator failed.
HRS	. . . .	. - .	. . .		RNG startup failed.
HRTP	. . . .	. - .	-	. - - .	Periodic (scheduled daily) RNG selftest failed.
HRM	. . . .	. - .	- -		RNG data matched.
HS	. . . .	. . .			Unexpected error from SCSI controller or host interface initialization failed.
HV	. . . .	. . . -			Environment sensors failed (for example, temperature sensor)
HCV	. . . .	- . - .	. . . -		CPLD wrong version for PCI policing firmware.
H P P	. . . .	. - - .	. - - .		PCI Interface Policing failure.
HST	. . . .	. . .	-		Speed test failed.
H H R	. . . .	. . . .	. - .		RTC hardware detection failed or random number generator detection failed.
H R H	. . . .	. - .	. . . .		RNG hardware failed during operation

Chapter 31. Morse code error messages

Code						Meaning
KR	- . -	. - .				RSA selftest failed.
H M n	. . . .	- -	#			DSP n failed self-test at start up.
H C n C A	. . . .	- . - .	#	- . - .	. -	CPU n failed self-test; no memory for cached RAM test.
H C n C C	. . . .	- . - .	#	- . - .	- . - .	CPU n failed self-test; CPU ID check failed.
H C n C F	. . . .	- . - .	#	- . - .	. . - .	CPU n failed self-test; freeing memory for cached RAM test.
H C n C G	. . . .	- . - .	#	- . - .	- - .	CPU n failed self-test; setting up cached RAM test.
H C n C R	. . . .	- . - .	#	- . - .	. - .	CPU n failed self-test; read error during cached RAM test.
H C n C W	. . . .	- . - .	#	- . - .	. - -	CPU n failed self-test; write error during cached RAM test.
H C n H D	. . . .	- . - .	#	. . . .	- . .	DRBG n failed self-test.
H C n K A	. . . .	- . - .	#	- . -	. -	CPU n failed selftest - AES known-answer test.
H C n K B	. . . .	- . - .	#	- . -	- . . .	CPU n failed selftest - AES CMAC known-answer test.
H C n K C	. . . .	- . - .	#	- . -	- . - .	CPU n failed selftest - ECDSA known-answer test
H C n K E	. . . .	- . - .	#	- . -	.	CPU n failed self-test; DES known-answer test.
H C n K F	. . . .	- . - .	#	- . -	. . - .	CPU n failed self-test; Triple-DES known-answer test.
H C n K H	. . . .	- . - .	#	- . -	. . . .	CPU n failed self-test; SHA-1 known-answer test.
H C n K I	. . . .	- . - .	#	- . -	. .	CPU n failed selftest - HMAC-SHA512 known-answer test.
H C n K J	. . . .	- . - .	#	- . -	. - - -	CPU n failed selftest - HMAC-SHA256 known-answer test.
H C n K M	. . . .	- . - .	#	- . -	- -	CPU n failed self-test; HMAC-SHA1 known-answer test.
H C n K N	. . . .	- . - .	#	- . -	- .	CPU n failed selftest - HMAC-SHA224 known-answer test.

Code						Meaning
H C *n* K P	. . . .	- . - .	#	- . -	. - - .	CPU n failed selftest - HMAC-SHA384 known-answer test.
H C *n* K R	. . . .	- . - .	#	- . -	. - .	CPU n failed selftest - RSA known-answer test
H C *n* K S	. . . .	- . - .	#	- . -	. . .	CPU *n* failed self-test; DSA known-answer test.
H C *n* L C	. . . .	- . - .	#	. - . .	- . - .	CPU *n* failed self-test; locking check.
H C *n* P S	. . . .	- . - .	#	. - - .	. . .	CPU *n* failed self-test; test terminated at start.
H C *n* RT	. . . .	- . - .	#	. - .	-	CPU n failed selftest - RTC check.
H C *n* S A	. . . .	- . - .	#	. . .	. - - .	CPU *n* failed self-test; no memory for uncached RAM test.
H C *n* S F	. . . .	- . - .	#	. . .	. . - .	CPU *n* failed self-test; freeing memory for uncached RAM test.
H C *n* S R	. . . .	- . - .	#	. . .	. - .	CPU *n* failed self-test; read error during uncached RAM test.
H C *n* S W	. . . .	- . - .	#	. . .	. - -	CPU *n* failed self-test; write error during uncached RAM test.
H C *n* T S	. . . .	- . - .	#	-	. . .	CPU *n* failed self-test; could not start test.

31.4. Maintenance mode errors

The following error codes indicate faults encountered when a module is in the maintenance mode.

Code				Meaning	Action
I D	. .	- . .		Copies of metadata do not match when trying to run image.	Contact Support.
I H	. .	. . . .		Bad metadata: hash mismatch.	Repeat firmware upgrade.
I I	. .	. .		Execution image does not match metadata.	Contact Support.
I L	. .	. - . .		Bad metadata: either bad length or bad metadata when running loadboot application.	Repeat firmware upgrade.

Code				Meaning	Action
I M	. .	- -		Bad metadata: malformed ImageMetaData.	Repeat firmware upgrade.
I P	. .	. - - .		Bad metadata: bad padding.	Repeat firmware upgrade.
I R	. .	. - .		Bad metadata: extra bytes at end.	Repeat firmware upgrade.
I S	. .	. . .		Image entry point not found.	Contact Support.
I U	. .	. . -		Bad metadata: ROM blank.	Repeat firmware upgrade.
I X	. .	- . . -		Bad metadata: malformed header.	Repeat firmware upgrade.
J H	. - - -	. . . .		Both copies of metadata invalid.	Contact Support.
H Z E	. . . .	- - . .	.	Monitor checksum failed.	Contact Support.
K F E	- . -	. . - .	.	Flash sector erase failed.	Repeat firmware upgrade.
K F P	- . -	. . - .	. - - .	Flash sector program failed.	Repeat firmware upgrade.
M M D	- -	- -	- . .	No memory for download buffer.	Contact Support.

> ℹ️ For instructions on upgrading module firmware, see the appendix in the User Guide for your module type.

31.5. Operational mode errors

The following runtime library error codes could be caused by either bugs in the firmware or faulty hardware.

Code				Meaning	Action
D	- . .			Fail command received.	Reset module by turning it off and then on again.
T	-			Temperature of the module has exceeded the maximum allow-able.	Restart your host computer, and improve module cooling.
G G G	- - .	- - .	- - .	Failure when performing `ClearUnit` or `Fail` command.	Contact Support.
I J A	. .	. - - -	. -	Audit logging: failed to send audit log message.	Contact Support.

Code				Meaning	Action
I J B	..	. - - -	- ...	Audit logging: no module memory (therefore failed to send audit log message).	Contact Support.
I J C	..	. - - -	-.-.	Audit logging: key problem or FIPS incompatibility (therefore failed to sign audit log message).	Contact Support.
I J D	..	. - - -	-..	Audit logging: NVRAM problem (therefore failed to configure or send audit log message).	Contact Support.

> ℹ️ SOS IJA can occur for any type of log message (i.e. a log message, signature block or certifier block).

> ℹ️ To improve the cooling of your PCIe module, increase the distance between PCIe cards, and increase the airflow through your host computer.

32. Uninstalling Security World Software

This appendix describes how to uninstall Security World Software.

> ⛔ Do not uninstall the Security World Software unless either you are certain it is no longer required or you are going to upgrade it.

The uninstaller removes only those files that were created during the installation. To remove key data or Security World data, navigate to the installation directory and delete the files in the *%NFAST_KMDATA%* folder.

If you intend to remove your Security World before uninstalling the Security World Software, Entrust recommends that you erase the OCS before you erase the Security World or uninstall the Security World Software. Except where Remote Administration cards are used, after you have erased a Security World, you can no longer erase any cards that belonged to it.

1. Log in to the host computer as Administrator or as a user with local administrator rights.
2. Run the following command to erase the OCS:

   ```
   createocs -m# -s0 --erase
   ```

   where # is the module number.

3. Navigate to the Windows Control Panel, and double-click **Programs and Features**.
4. Select the Security World Software entry, then click **Uninstall** to remove the software.

33. Application Performance Tuning

33.1. Job Count

To achieve the best throughput of cryptographic jobs (such as Sign or Decrypt) in your application, arrange for multiple jobs to be on the go at the same time, rather than doing them one at a time. This is true even when using only a single HSM in your system.

When using an nShield HSM, Entrust recommend that you set the number of outstanding jobs within the `rec. queue` (recommended queue) range specified by the `enquiry` output for the module.

If you are sending single jobs synchronously from each thread of your client application, try to keep the number of threads within this `rec. queue` range for best throughput.

When using higher-level APIs, such as PKCS#11, your application could benefit from increasing the thread count above the `rec. queue` range or the number that gives the best throughput when using nCore directly.

If you are load-balancing across multiple HSMs and want to maximize throughput across all of them, then use the sum of all `rec. queue` ranges for each of the modules to set the target for the outstanding jobs.

The `ncperftest` utility supports performance measurements of a range of cryptographic operations with different job counts and client thread counts. You may find this useful to inform tuning of your application. Run `ncperftest --help` to see the available options.

33.2. Client Configuration

If your application is coded directly against nCore, you have a choice of sending multiple jobs asynchronously from a single client connection to the hardserver, or having multiple threads each with their own client connection to the hardserver with a single job sent synchronously in each. You can use the `--threads` parameter to the `ncperftest` utility to experiment with the performance impact of having more threads/connections with fewer jobs outstanding in each, or having fewer or just one thread/connection with more jobs outstanding in that connection.

When using higher-level APIs such as PKCS#11, all cryptographic operations are synchronous, so larger numbers of threads must be used to increase the job count and make full use of HSM resources. These APIs automatically create a hardserver connection for each thread. If many HSMs are being used, a great many threads may be required to achieve best

throughput. You can adjust the thread counts in the performance test tools for these APIs (for example, `cksigtest` for PKCS#11) to gauge how much concurrency is required for best throughput in your application.

33.3. Highly Multi-threaded Client Applications

If your application is highly multi-threaded, operating system defaults may not be optimal for best performance:

You may benefit from using a scalable memory allocator that is designed to be efficient in multi-threaded applications, examples include `tcmalloc`.

On some systems the default operating system scheduling algorithm is also not optimized for highly multi-threaded applications. A real-time scheduling algorithm such as the POSIX round-robin scheduler may yield noticeable performance improvements for your application.

33.4. File Descriptor Limits (Linux)

On Linux systems, large numbers of threads each with their own hardserver connection will require your application to make use of large numbers of file descriptors. It may be necessary to increase the file descriptor limit for your application. This can be done using `ulimit -n NewLimit` on most systems, but you may need to increase system-wide hard limits first.

34. Merged Keys Concept

A merged key is a level of abstraction higher than normal keys. It holds an internal list of normal key IDs, each associated with its corresponding module. When a command to the hardserver specifies a MergedKey ID instead of a normal (single) key ID, the hardserver chooses an HSM and corresponding single key ID from the list in the Merged Key. See diagram below. Which module is chosen may depend on multiple factors, including load sharing settings in the hardserver config.

Benefits of MergedKeys:

- A client need hold only a single M_KeyID reference instead of one for each HSM.

- That ID remains usable even while the key's actual IDs on HSMs can fluctuate.

- The hardserver can use heuristics to choose the most appropriate HSM (for example, the least heavily loaded one).

- If some HSMs become unavailable, the hardserver uses the remaining ones automatically.
    - A MergedKey can be updated, removing its entry for a defunct HSM and corre-

sponding single-key ID.

- New HSMs can be added: if a new HSM is made operational and added to the relevant security world, then

  ° the key can be loaded onto that HSM, thus creating a new single-key ID for that key on that HSM, and then

  ° the new (Key ID, HSM) pair can be added to the existing Merged Key.

35. Product returns

If you need to return your nShield HSM, contact Entrust nShield Support for instructions:
https://nshieldsupport.entrust.com

36. Remote File System Volumes

The hardserver (**Linux**) or nFast Server (**Windows**) service restricts the paths that can be shared as RFS (Remote File System) volumes using the `remote_file_system` section of the `config` file or using the `rserverperm --accessfiles` command-line configuration.

By default, the following paths are permitted:

- `/opt/nfast/kmdata` (**Linux**).
- `%NFAST_KMDATA%`, typically `C:\ProgramData\nCipher\Key Management Data` (**Windows**).
- Any path that was created by the `rfs-setup` utility and associated with RFS volumes to prepare an RFS for an nShield HSM or for use with the `rfs-sync` utility.
- Subdirectories of permitted paths.

If you want to add custom paths not included in this list as RFS volumes, you must add them to the list of permitted paths before starting the hardserver (**Linux**) or nFast Server (**Windows**) service. If you make these changes after starting the service, you need to restart it for the changes to take effect.

You can update the list of permitted paths by either setting the NFSERV_RFS_ALLOWED_-PATHS environment variable (see Allow custom RFS paths with an environment variable) or by creating an additional `config.secure` configuration file (see Allow custom RFS paths with a configuration file.)

36.1. Allow custom RFS paths with an environment variable

Linux

1. If the `/etc/nfast.conf` file does not already exist, create it.

   This file must only be writable by root. This is enforced by nShield start-up scripts.

2. Add the `NFSERV_RFS_ALLOWED_PATHS` environment variable to the `nfast.conf` file with a colon-separated list of paths (`/<path>/share`).

   For example, to share `path1` and `path 2` (spaces are permitted):

   ```
   export NFSERV_RFS_ALLOWED_PATHS=/path1/share:/path 2/share
   ```

Windows

   Create the `NFSERV_RFS_ALLOWED_PATHS` environment variable in the global system environment variables with a semicolon-separated list of paths (`\<path>\share`).

For example, to share `path1` and `path 2` (spaces are permitted):

```
C:\path1\share;D:\path 2\share
```

36.2. Allow custom RFS paths with a configuration file

1. Create the `config.secure` in `/opt/nfast/hardserver.d` (**Linux**) or the `%PROGRAM-DATA%\nCipher\hardserver.d` directory, which is typically `C:\ProgramData\nCipher\hardserver.d\config.secure`.

2. Add the paths as values in an `rfs_allowed_paths` JSON array. The JSON must be valid.

   For example, to share `path1` and `path 2` (spaces are permitted):

   Linux

   ```
   {
     "rfs_allowed_paths" : ["/path1/share", "/path 2/share"]
   }
   ```

   Windows

   ```
   {
     "rfs_allowed_paths" : ["C:\\path1\\share", "D:\\path 2\\share"]
   }
   ```

   > You must use a backslash (`\`) to escape the backslashes in the path.