



ENTRUST

nShield Security World

nShield 5s v13.5.1 User Guide

16 February 2024

Table of Contents

| | |
|---|----|
| 1. Introduction | 1 |
| 1.1. Read this guide if | 1 |
| 1.1.1. Terminology | 1 |
| 1.2. Conventions | 1 |
| 1.2.1. Model numbers | 1 |
| 1.2.2. Security World Software | 2 |
| 1.2.3. Setting the PATH for nShield utilities | 4 |
| 1.3. Further information | 5 |
| 1.4. Security advisories | 5 |
| 1.5. Recycling and disposal information | 5 |
| 2. Security Worlds | 6 |
| 2.1. Security | 7 |
| 2.1.1. Smart cards | 7 |
| 2.1.2. Remote Operator | 8 |
| 2.1.3. Remote Administration | 9 |
| 2.1.4. Client cooperation feature | 9 |
| 2.1.5. NIST SP800-131A | 9 |
| 2.1.6. FIPS 140 compliance | 10 |
| 2.2. Platform independence | 10 |
| 2.3. Application independence | 11 |
| 2.4. Flexibility | 11 |
| 2.4.1. Using the Security World key: module-protected keys | 12 |
| 2.4.2. Using Operator Card Sets: OCS-protected keys | 12 |
| 2.4.3. Using passphrases for extra security | 15 |
| 2.4.4. Using softcard-protected keys | 16 |
| 2.5. Scalability | 17 |
| 2.5.1. Load-sharing | 18 |
| 2.6. Robustness | 19 |
| 2.6.1. Backup and recovery | 19 |
| 2.6.2. Replacing a hardware security module | 19 |
| 2.6.3. Replacing the Administrator Card Set | 20 |
| 2.6.4. Replacing an Operator Card Set or recovering keys to softcards | 21 |
| 2.7. Audit Logging | 22 |
| 2.8. KeySafe and Security Worlds | 22 |
| 2.9. Applications and Security Worlds | 23 |
| 2.10. The nShield PKCS #11 library and Security Worlds | 23 |
| 2.11. Risks | 24 |

| | |
|--|----|
| 3. Platform services and ncoreapi | 26 |
| 3.1. ncoreapi service | 26 |
| 3.1.1. Multi-tenancy ready | 26 |
| 3.2. Platform services | 26 |
| 3.3. Separation of services | 27 |
| 4. nShield 5s modes of operation | 28 |
| 4.1. Modes of operation | 28 |
| 4.2. Check and change the mode of operation | 29 |
| 4.3. Return to factory state | 30 |
| 4.3.1. Purpose of factory state | 30 |
| 4.3.2. Enter and exit the factory state | 31 |
| 4.4. Recovery mode | 31 |
| 4.4.1. Restrictions in recovery mode | 32 |
| 4.4.2. Entry into recovery mode | 32 |
| 4.4.3. Exit from recovery mode | 33 |
| 5. Software installation | 34 |
| 5.1. After software installation | 34 |
| 6. Managing the system clock | 35 |
| 6.1. System interaction with the system clock | 35 |
| 6.2. Checking the system clock | 36 |
| 6.3. Adjusting the system clock | 36 |
| 6.4. Setting the system clock | 36 |
| 6.4.1. Restrictions on setting the clock | 37 |
| 7. Set up communication between host and module | 38 |
| 7.1. Overview of SSH keys | 38 |
| 7.2. Installation of SSH keys as part of software installation | 38 |
| 7.3. Installation of SSH keys independently of a software installation | 38 |
| 7.4. Viewing installed SSH keys | 39 |
| 7.5. Changing installed SSH keys | 39 |
| 7.6. Making a backup of installed SSH keys | 40 |
| 7.7. Restoring SSH keys from backup | 40 |
| 7.8. Preparing an HSM for use in another host | 40 |
| 8. Administration of platform services | 43 |
| 8.1. hsmadmin | 43 |
| 8.1.1. hsmadmin factorystate | 44 |
| 8.1.2. hsmadmin status | 44 |
| 8.1.3. hsmadmin npkginfo | 45 |
| 8.1.4. hsmadmin upgrade | 45 |
| 8.1.5. hsmadmin reset | 46 |

| | |
|---|----|
| 8.1.6. hsmadmin enroll | 47 |
| 8.1.7. hsmadmin keys | 48 |
| 8.1.8. hsmadmin logs | 51 |
| 8.1.9. hsmadmin info | 55 |
| 8.1.10. hsmadmin settime | 55 |
| 8.1.11. hsmadmin gettime | 56 |
| 8.1.12. hsmadmin setminvsn | 57 |
| 8.1.13. hsmadmin getenvstats | 58 |
| 8.1.14. hsmadmin cs5 | 60 |
| 9. System logging | 62 |
| 9.1. Maximum log size | 62 |
| 9.2. Interaction with the system clock | 63 |
| 9.3. Init log | 63 |
| 9.3.1. Retrieving the init log | 63 |
| 9.3.2. Clearing the init log | 64 |
| 9.4. System log | 64 |
| 9.4.1. Signed system logs | 64 |
| 9.4.2. Retrieving the system log | 65 |
| 9.4.3. Clearing the system log | 66 |
| 10. Client Software and module configuration | 69 |
| 10.1. About user privileges | 69 |
| 10.2. Setting up client cooperation | 69 |
| 10.2.1. Useful utilities | 72 |
| 10.2.2. Setting environmental variables | 74 |
| 10.2.3. Logging and debugging | 75 |
| 10.2.4. Configuring Java support for KeySafe | 75 |
| 10.3. Configuring the hardserver | 75 |
| 10.3.1. Overview of hardserver configuration file sections | 77 |
| 10.3.2. Using multiple modules | 80 |
| 10.4. Stopping and restarting the hardserver | 82 |
| 10.5. Configure the hardserver to export the module for guest VM usage | 83 |
| 11. Enabling optional features | 86 |
| 11.1. Available optional features | 86 |
| 11.1.1. Elliptic Curve | 87 |
| 11.1.2. Elliptic Curve activation | 87 |
| 11.1.3. Elliptic Curve support on the nShield product line | 87 |
| 11.1.4. nShield software / API support required to use elliptic curve functions | 88 |
| 11.1.5. Named Curves | 89 |
| 11.1.6. Custom curves | 89 |

| | |
|---|-----|
| 11.1.7. Further information on using elliptic curves | 89 |
| 11.1.8. Secure Execution Engine (SEE) | 90 |
| 11.1.9. Remote Operator support | 90 |
| 11.1.10. ISO smart card Support (ISS) | 91 |
| 11.1.11. Korean algorithms | 91 |
| 11.1.12. Fast RNG for ECDSA | 91 |
| 11.2. Ordering additional features | 91 |
| 11.3. Enabling features | 92 |
| 11.3.1. Viewing enabled features | 92 |
| 11.3.2. Enabling features with a smart card | 92 |
| 11.3.3. Enabling features without a smart card | 93 |
| 12. Creating and managing a Security World | 95 |
| 12.1. Creating a Security World | 95 |
| 12.1.1. The creation process | 95 |
| 12.1.2. Security World Files | 96 |
| 12.1.3. Security World options | 98 |
| 12.1.4. Creating a Security World using new-world | 102 |
| 12.1.5. After you have created a Security World | 110 |
| 12.2. Displaying information about your Security World | 111 |
| 12.2.1. Displaying information about a Security World with nfkminfo | 111 |
| 12.2.2. Displaying information about a Security World with kmfile-dump | 112 |
| 12.3. Adding or restoring an HSM to the Security World | 112 |
| 12.3.1. Adding an HSM to a Security World with the CSP or CNG wizard (Windows) | 114 |
| 12.3.2. Adding an HSM to a Security World with new-world | 115 |
| 12.4. Security World migration | 116 |
| 12.4.1. Pre-requisites for migrating keys | 116 |
| 12.4.2. Restrictions on migrating keys | 117 |
| 12.4.3. About the migration utility | 118 |
| 12.4.4. Migrating keys | 120 |
| 12.4.5. Migrating keys process | 120 |
| 12.4.6. Verifying the integrity of the migrated keys | 121 |
| 12.4.7. Migrating keys using custom protection pairs | 122 |
| 12.4.8. Troubleshooting | 123 |
| 12.5. Migrating KMDATA (Windows) | 125 |
| 12.6. Erasing a module from a Security World | 126 |
| 12.6.1. Erasing a module with new-world | 127 |
| 12.6.2. Erasing a module with KeySafe | 127 |
| 12.6.3. Erasing a module with initunit | 128 |

| | |
|---|------------|
| 12.7. Deleting a Security World | 128 |
| 13. Managing card sets and softcards | 130 |
| 13.1. Creating Operator Card Sets (OCSs) | 130 |
| 13.1.1. Persistent Operator Card Sets | 131 |
| 13.1.2. Time-outs | 132 |
| 13.1.3. FIPS 140 Level 3-compliant Security Worlds | 132 |
| 13.1.4. Creating an Operator Card Set using the command line | 132 |
| 13.1.5. Creating an Operator Card Set with KeySafe | 134 |
| 13.1.6. Creating an Operator Card Set with the CSP or CNG wizard (Windows) .. | 137 |
| 13.2. Creating softcards | 139 |
| 13.2.1. Creating a softcard with ppmk | 140 |
| 13.2.2. Creating softcards with KeySafe | 140 |
| 13.2.3. Creating a softcard with the CNG wizard (Windows) | 141 |
| 13.3. Erasing cards and softcards | 143 |
| 13.3.1. FIPS 140 Level 3-compliant Security Worlds | 144 |
| 13.3.2. Erasing cards with KeySafe | 144 |
| 13.3.3. Erasing cards using the command line | 145 |
| 13.3.4. Erasing softcards | 145 |
| 13.4. Viewing cards and softcards | 146 |
| 13.4.1. Viewing card sets with KeySafe | 147 |
| 13.4.2. Viewing card sets using the command line | 148 |
| 13.4.3. Viewing softcards | 148 |
| 13.4.4. Verifying the passphrase of a card or softcard | 150 |
| 13.5. Changing card and softcard passphrase | 151 |
| 13.5.1. Changing known passphrase | 151 |
| 13.5.2. Changing unknown or lost passphrase | 154 |
| 13.6. Replacing Operator Card Sets | 156 |
| 13.6.1. Replacing OCSs with KeySafe | 157 |
| 13.6.2. Replacing OCSs or softcards with rocs | 159 |
| 13.7. Replacing the Administrator Card Set | 167 |
| 13.7.1. Replacing an ACS with KeySafe | 168 |
| 13.7.2. Replacing an Administrator Card Set using racs | 169 |
| 14. Application interfaces | 171 |
| 14.1. nCipherKM JCA/JCE CSP | 171 |
| 14.1.1. Installing the nCipherKM JCA/JCE CSP | 172 |
| 14.1.2. Named Modules in Java 11 and Java 17 | 177 |
| 14.1.3. keytool | 177 |
| 14.1.4. Using keys | 178 |
| 14.1.5. System properties | 178 |

| | |
|---|-----|
| 14.1.6. Compatibility | 181 |
| 14.2. nShield PKCS #11 library | 182 |
| 14.2.1. Choosing functions | 182 |
| 14.2.2. PKCS #11 library with Security Assurance Mechanism | 184 |
| 14.2.3. Using the nShield PKCS #11 library | 186 |
| 14.2.4. nShield PKCS #11 library environment variables | 188 |
| 14.2.5. Checking the installation of the nShield PKCS #11 library | 202 |
| 14.2.6. How the nShield PKCS #11 library protects keys | 204 |
| 14.3. nShield native and custom applications | 205 |
| 14.4. Microsoft CAPI CSP | 205 |
| 14.4.1. Installing the CAPI CSP | 205 |
| 14.4.2. Importing a key | 206 |
| 14.4.3. Supported algorithms | 206 |
| 14.4.4. Container storage format | 207 |
| 14.4.5. Utilities for the CAPI CSP | 208 |
| 14.4.6. Uninstalling the CAPI CSP | 209 |
| 14.5. Microsoft Cryptography API: Next Generation (CNG) | 209 |
| 14.5.1. Configuring the nShield CNG CSP | 210 |
| 14.5.2. Supported algorithms for CNG | 214 |
| 14.5.3. Migrating keys for CNG | 216 |
| 14.5.4. Using CAPI keys in CNG | 220 |
| 14.5.5. Utilities for CNG | 220 |
| 15. Working with CodeSafe | 228 |
| 15.1. CodeSafe applications | 228 |
| 15.2. The csadmin utility tool | 228 |
| 15.2.1. Manage loadable CodeSafe images | 229 |
| 15.2.2. Load a CodeSafe application | 230 |
| 15.2.3. Start a CodeSafe application | 231 |
| 15.2.4. Stop a CodeSafe application | 232 |
| 15.2.5. Destroy a CodeSafe application | 233 |
| 15.2.6. List CodeSafe applications | 233 |
| 15.2.7. Configure CodeSafe 5 | 234 |
| 15.2.8. Manage CodeSafe 5 logging | 236 |
| 15.2.9. Obtain machine-specific statistics | 238 |
| 15.2.10. Manage container SSHD | 240 |
| 15.2.11. Manage identities | 246 |
| 16. Remote Operator | 249 |
| 16.1. About Remote Operator | 249 |
| 16.2. Configuring Remote Operator | 249 |

| | |
|--|------------|
| 16.2.1. Overview of configuring Remote Operator | 250 |
| 16.2.2. Configuring HSMs for Remote Operator | 250 |
| 16.2.3. Configuring slot import and export | 251 |
| 16.2.4. Using Remote Operator with applications requiring cards in slot 0 | 253 |
| 16.2.5. Using Remote Operator on Remapped Slots | 253 |
| 16.2.6. Configuration Example for Using Remote Administration and Remote Operator Concurrently | 254 |
| 16.2.7. Using Remote Operator with Remote Administration with Older Versions of the Software | 255 |
| 16.3. Creating OCSs and keys for Remote Operator | 256 |
| 16.3.1. Creating OCSs for use with Remote Operator | 256 |
| 16.3.2. Loading Remote Operator Card Sets | 257 |
| 16.3.3. Generating keys for use with Remote Operator | 257 |
| 16.3.4. Configuring the application | 257 |
| 17. Working with keys | 259 |
| 17.1. Generating keys | 259 |
| 17.1.1. Generating keys using the command line | 259 |
| 17.1.2. Generating keys with KeySafe | 261 |
| 17.1.3. Generating NVRAM-stored keys | 262 |
| 17.2. Importing keys | 263 |
| 17.2.1. Importing keys from the command line | 264 |
| 17.2.2. Importing keys with KeySafe | 265 |
| 17.3. Listing supported applications with generatekey | 266 |
| 17.4. Retargeting keys with generatekey | 266 |
| 17.5. Viewing keys | 267 |
| 17.5.1. Viewing keys with KeySafe | 267 |
| 17.5.2. Viewing keys using the command line | 267 |
| 17.6. Verifying Key Generation Certificates with nfkmverify | 269 |
| 17.6.1. Usage | 270 |
| 17.7. Discarding keys | 271 |
| 17.8. Restoring keys | 271 |
| 18. Using KeySafe | 272 |
| 18.1. Setting up KeySafe | 272 |
| 18.2. Starting KeySafe | 273 |
| 18.3. About the KeySafe window | 273 |
| 18.3.1. Sidebar | 274 |
| 18.3.2. Menu buttons | 274 |
| 18.3.3. Menus | 275 |
| 18.3.4. Module Status tree | 275 |

| | |
|---|-----|
| 18.3.5. Main panel area | 278 |
| 18.4. Errors | 279 |
| 18.4.1. Unable to establish KeySafe session. | 280 |
| 18.4.2. Unable to generate key | 280 |
| 19. Warrant Management | 282 |
| 19.1. Warrant management for nShield Solo and nShield Edge | 282 |
| 19.1.1. Warranting steps for nShield Solo and nShield Edge | 282 |
| 19.1.2. nfwarrant command-line utility | 283 |
| 19.2. Warrant management for nShield Connect + and nShield Connect XC | 285 |
| 19.3. Warrant management for nShield 5s and nShield 5c | 285 |
| 20. Supplied utilities | 286 |
| 20.1. Utilities for general operations | 286 |
| 20.1.1. hsmadmin | 286 |
| 20.1.2. enquiry | 286 |
| 20.1.3. checkmod | 286 |
| 20.1.4. cfg-mkdefault | 287 |
| 20.1.5. cfg-remoteslots | 287 |
| 20.1.6. cfg-reread | 287 |
| 20.1.7. fet | 287 |
| 20.1.8. ncdate | 287 |
| 20.1.9. ncversions | 287 |
| 20.1.10. nopclearfail | 288 |
| 20.1.11. nvram-backup | 288 |
| 20.1.12. nvram-sw | 288 |
| 20.1.13. pubkey-find | 288 |
| 20.2. randchk | 289 |
| 20.2.1. retrievewarrants | 289 |
| 20.2.2. rtc | 289 |
| 20.2.3. slotinfo | 289 |
| 20.2.4. snmpbulkwalk snmpget snmpgetnext snmpset snmpstat snmptranslate snmpwalk | 289 |
| 20.2.5. stattree | 289 |
| 20.3. Test analysis tools | 290 |
| 20.4. Security World utilities | 291 |
| 20.5. CodeSafe utilities | 294 |
| 20.6. PKCS #11 | 294 |
| 20.7. MSCAPI utilities (Windows) | 295 |
| 20.8. CNG (Windows) | 296 |
| 20.9. Developer-specific utilities | 298 |

| | |
|--|------------|
| 20.10. Utilities that require a privileged connection | 298 |
| 21. Using silent installations | 299 |
| 21.1. Installing using the silent install functionality | 299 |
| 21.2. Uninstalling using the silent install functionality | 300 |
| 22. Using nShield commands from PowerShell | 301 |
| 22.1. Install and configure PowerShell | 301 |
| 22.2. Calling nShield commands at the PowerShell prompt | 302 |
| 22.3. PowerShell modes: interactive and batch | 302 |
| 22.4. Input pipelines | 304 |
| 22.5. Secure strings | 304 |
| 23. Preload Utility | 306 |
| 23.1. Overview | 306 |
| 23.2. Using Preload | 306 |
| 23.2.1. Preload Commands | 307 |
| 23.2.2. Preload file location | 307 |
| 23.2.3. Preload Command Line Arguments | 307 |
| 23.2.4. Pattern Matching | 309 |
| 23.3. Preload File | 309 |
| 23.4. Softcard Support | 310 |
| 23.4.1. No Cardset Keys | 311 |
| 23.5. FIPS Auth | 311 |
| 23.6. Admin Keys | 311 |
| 23.6.1. Listing | 311 |
| 23.6.2. Loading | 312 |
| 23.7. High Availability | 312 |
| 23.7.1. Prerequisites for high availability mode | 313 |
| 23.7.2. Differences from legacy behaviour | 313 |
| 23.7.3. Conditions for Management/Reloading | 313 |
| 23.7.4. Merged Keys in the Preload File | 314 |
| 23.7.5. Polling Interval | 314 |
| 23.7.6. Key timeouts and use limits | 315 |
| 23.7.7. Multiple Preload instances in high availability mode | 315 |
| 23.7.8. FIPS Auth in High Availability mode | 317 |
| 23.7.9. PKCS #11 and JCE | 317 |
| 23.7.10. Unsupported options | 318 |
| 23.8. Logging | 318 |
| 23.9. Using preloaded objects - Worked example | 319 |
| 24. Environment variables | 321 |
| 25. Logging, debugging, and diagnostics | 325 |

| | |
|---|-----|
| 25.1. Logging and debugging | 325 |
| 25.1.1. Environment variables to control logging | 325 |
| 25.1.2. Logging from the nShield CSP and CNG (Windows) | 329 |
| 25.1.3. Logging and debugging information for PKCS #11 | 331 |
| 25.1.4. Hardserver debugging | 332 |
| 25.1.5. Debugging information for Java | 332 |
| 25.1.6. appliance-cli: system logs utility | 334 |
| 25.2. Diagnostics and system information | 335 |
| 25.2.1. nfdiag: diagnostics utility | 335 |
| 25.2.2. nfkminfo: information utility | 336 |
| 25.2.3. perfcheck: performance measurement checking tool | 345 |
| 25.2.4. stattree: information utility | 349 |
| 25.3. How data is affected when a module loses power and restarts | 356 |
| 26. Hardserver configuration files | 357 |
| 26.1. Hardserver configuration files | 357 |
| 26.2. General hardserver configuration settings | 358 |
| 26.2.1. server_settings | 358 |
| 26.2.2. hardserver loglevel | 360 |
| 26.2.3. server_performance | 362 |
| 26.2.4. module_settings | 362 |
| 26.2.5. server_remotecomms | 362 |
| 26.2.6. server_startup | 363 |
| 26.2.7. load_seemachine | 364 |
| 26.2.8. slot_imports | 365 |
| 26.2.9. slot_exports | 366 |
| 26.2.10. dynamic_slots | 366 |
| 26.2.11. slot_mapping | 367 |
| 26.2.12. dynamic_slot_timeouts | 367 |
| 26.2.13. audit_logging | 367 |
| 26.3. Sections only in client configuration files | 368 |
| 26.3.1. nethsm_imports | 368 |
| 26.3.2. rfs_sync_client | 368 |
| 26.3.3. remote_file_system | 369 |
| 26.3.4. remote_administration_service_slot_server_startup | 370 |
| 27. Cryptographic algorithms | 371 |
| 27.1. Symmetric algorithms | 371 |
| 27.2. Asymmetric algorithms | 371 |
| 27.3. FIPS information | 372 |
| 27.4. Compatibility of Security World versions with FIPS | 373 |

| | |
|---|-----|
| 28. Audit Logging | 374 |
| 28.1. Aims of audit logging | 374 |
| 28.1.1. Audit logging and CEF audit logging | 374 |
| 28.2. Configuring audit logging | 375 |
| 28.2.1. Enabling audit logging | 375 |
| 28.2.2. Disabling audit logging | 375 |
| 28.2.3. Key usage logging | 375 |
| 28.3. Commands audited | 376 |
| 28.4. Audit log contents | 380 |
| 28.4.1. Audit segment header contents | 380 |
| 28.4.2. Audit log contents | 381 |
| 28.5. Audit log verification | 383 |
| 28.6. Audit log administration | 383 |
| 28.6.1. Audit logging and firmware upgrade | 383 |
| 28.6.2. Audit logging and the system clock | 384 |
| 29. Key generation options and parameters | 385 |
| 29.1. Key application type (APPNAME) | 385 |
| 29.2. Key properties (NAME=VALUE) | 386 |
| 29.3. Available key properties by action/application | 390 |
| 30. Maintenance of nShield Hardware | 394 |
| 30.1. Voltage Monitoring for Battery Replacement | 394 |
| 30.2. Temperature Monitoring for Airflow Validation | 395 |
| 31. System upgrade | 396 |
| 31.1. Terminology | 396 |
| 31.2. Software and firmware compatibility | 396 |
| 31.3. System upgrade procedure | 396 |
| 31.3.1. Software upgrade procedure | 397 |
| 31.3.2. Firmware upgrade procedure | 398 |
| 32. Uninstalling Security World software | 399 |
| 32.1. Uninstalling the Security World software on Windows | 399 |
| 32.1.1. Removal of Security World data | 400 |
| 32.2. Uninstalling the Security World software on Linux | 400 |
| 33. Upgrading firmware | 402 |
| 33.1. Primary, recovery and bootloader firmware | 402 |
| 33.2. Firmware version control | 402 |
| 33.2.1. Version Security Number (VSN) | 402 |
| 33.2.2. Bootloader version | 403 |
| 33.3. Firmware on the installation media | 404 |
| 33.3.1. Recognising firmware files | 404 |

| | |
|--|------------|
| 33.4. Firmware installation overview | 404 |
| 34. SNMP monitoring agent | 407 |
| 34.1. Install and activate the nShield SNMP monitoring agent | 407 |
| 34.1.1. Default installation settings | 407 |
| 34.1.2. Do you already have an SNMP agent running? | 407 |
| 34.1.3. Starting the SNMP agent | 408 |
| 34.2. Basic configuration | 409 |
| 34.2.1. Protecting the SNMP installation | 409 |
| 34.2.2. Configuring the SNMP agent | 409 |
| 34.2.3. Create the configuration files (Windows) | 410 |
| 34.2.4. The SNMP agent persistent configuration file | 411 |
| 34.2.5. Agent Behaviour | 412 |
| 34.2.6. agentaddress directive | 412 |
| 34.2.7. agentgroup and agentuser directives (Linux) | 412 |
| 34.2.8. System information (Linux) | 413 |
| 34.3. USM users | 413 |
| 34.4. Traditional access control | 415 |
| 34.5. VACM configuration | 417 |
| 34.6. Trap Configuration | 421 |
| 34.6.1. SNMPv1 and SNMPv2 traps | 421 |
| 34.6.2. SNMPv3 traps | 422 |
| 34.7. Using the SNMP agent with a manager application | 422 |
| 34.7.1. Manager configuration | 423 |
| 34.7.2. MIB module overview | 423 |
| 34.7.3. MIB functionality | 424 |
| 34.7.4. Memory usage monitoring | 425 |
| 34.7.5. Administration sub-tree overview | 426 |
| 34.7.6. netHSMNetwork sub-tree | 436 |
| 34.7.7. Statistics sub-tree overview | 437 |
| 34.8. SNMP agent command-line | 442 |
| 34.8.1. SNMP agent (snmpd) switches | 442 |
| 34.8.2. Using the SNMP command-line utilities | 443 |
| 35. HSM status and error codes | 444 |
| 35.1. LED status | 444 |
| 35.2. LED error states | 444 |
| 35.2.1. Error codes shown on the LED | 445 |
| 35.3. Error codes accessed remotely | 447 |
| 35.3.1. Runtime library errors | 447 |
| 35.3.2. Hardware driver errors | 447 |

| | |
|---|------------|
| 35.3.3. Operational mode errors | 449 |
| 35.4. Other errors | 449 |
| 36. Application Performance Tuning | 451 |
| 36.1. Job Count | 451 |
| 36.2. Client Configuration | 451 |
| 36.3. Highly Multi-threaded Client Applications | 452 |
| 36.4. File Descriptor Limits (Linux) | 452 |
| 37. Merged Keys Concept | 453 |
| 38. Product returns | 455 |
| 39. SSH Client Key Protection | 456 |
| 39.1. SSH Services | 456 |
| 39.2. SSH Client Key Encryption | 456 |
| 39.2.1. Available SSH Key Protection Options | 456 |
| 39.3. Setting Protections on SSH keys | 457 |
| 39.4. Permissions on SSH keys | 458 |

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.

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.

1.1.1. Terminology

The nShield 5s 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 |

1.2. Conventions

1.2.1. Model numbers

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

| Model number | Used for |
|--------------|-------------------|
| NC5536E-B | nShield 5s Base |
| NC5536E-M | nShield 5s Medium |
| NC5536E-H | nShield 5s High |

1.2.2. Security World Software

The hardserver software controls communication between applications and Entrust nShield product line HSMs, which may be installed locally or remotely. It runs as a daemon (**Linux**) or a service (**Windows**) on the host computer.

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

The nShield HSM is supplied with the latest version of the HSM firmware installed. For more information about:

- Upgrading the firmware, see [Upgrading 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](#).

1.2.2.1. 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`.

1.2.2.2. 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 | <code>/opt/nfast/</code> | <code>NFAST_HOME</code> | <code>C:\Program Files\nCipher\nfast</code> |
| Key Management Data | <code>/opt/nfast/kmdata/</code> | <code>NFAST_KMDATA</code> | <code>C:\ProgramData\nCipher\Key Management Data</code> |
| Dynamic Feature Certificates | <code>/opt/nfast/femcerts/</code> | <code>NFAST_CERTDIR</code> | <code>C:\ProgramData\nCipher\Feature Certificates</code> |
| Static Feature Certificates | <code>/opt/nfast/kmdata/hsm-ESN/features</code> | | <code>%NFAST_KMDATA\features</code> |
| Log Files | <code>/opt/nfast/log</code> | <code>NFAST_LOGDIR</code> | <code>C:\ProgramData\nCipher\Log Files</code> |
| User Log Files | <code>/home/<user>/nshieldlogs</code> | <code>NFAST_USER_LOGDIR</code> | <code>C:\Users\<user>\nshieldlogs</code> |
| Remote Static Feature Certificates | <code>opt/nfast/kmdata/hsm-ESN/features</code> | | <code>%NFAST_KMDATA\features</code> |
| Remote Dynamic Feature Certificates | <code>opt/nfast/kmdata/hsm-ESN/features</code> | | <code>%NFAST_KMDATA\features</code> |



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.

1.2.2.3. 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.

1.2.3. 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.

1.3. 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.

1.4. 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.

1.5. 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
- An *Administrator Card Set (ACS)*
A set of Administrator smart cards used to control access to the Security World configuration, 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



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*.



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).



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 hardware 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.

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. Client cooperation feature

The client cooperation feature allows nShield HSM host computers to automatically update the Security World and key data stored on a remote file system (RFS). For more information, see [Setting up client cooperation](#).

2.1.5. NIST SP800-131A

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

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.



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.



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.



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](#)).



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`

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](#).



Some applications do not support the use of passphrases.

2.4.3.1. Maximum passphrase length



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 softcard is loaded, it remains valid for loading its keys until its **KeyID** is destroyed.

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 servers:

- Ensure each server has at least one hardware security module fitted
- Copy the host data to each server, or make it available on a shared disk
- Use the recovery and replacement data with the ACS to load the required cryptographic keys securely onto every hardware security module.

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 data are propagated to the remaining servers. If your servers are part of a cluster, then the tools provided by the cluster should synchronize the data. If the servers are connected by a network, then they could all access the same copy of 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 servers.

You can configure the host computer of an nShield HSM to:

- Access a Remote File System (RFS) as used by an nShield HSM.
- Share Security World and key data stored in `/opt/nfast/kmdata/local` (**Linux**) or `%NFAST_KMDATA%\local` (**Windows**).

Client hardware security modules that access data in this way are described as *cooperating clients*. For more information, see [Setting up client cooperation](#).



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).

2.5.1. Load-sharing

If you have more than one hardware security module on your system, 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.



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. fallback 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 host 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 a new hardware security module by using the ACS and the recovery data to load the Security World key securely. Use the same mechanism to reload the Security World key if you need to upgrade the firmware in the hardware security module or if you need to add extra

hardware security modules to the Security World.

If you have more than one hardware security module on your system and you use one of the load-sharing modes identified above, then your system is resilient to the failure of individual hardware security modules.

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 KeySafe **Replace Administrator Card Set** option
- `racs` utility (see [Replacing the Administrator Card Set](#)).



You should also use `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 a given card set, and to recover keys protected by one softcard to another softcard, 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

The Audit Logging facility can be enabled on nShield HSMs where there is a requirement to provably log events. 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
- Public key log verification

The format of the audit logs depends on the version of firmware loaded on the HSMs.

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](#).



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.

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 host computers, 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.

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.



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.



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. Platform services and ncoreapi

The nShield HSM firmware provides multiple services. These are divided into platform services and the `ncoreapi` service.

3.1. ncoreapi service

The `ncoreapi` service provides cryptographic services to the end user. This can either be via custom applications created by the end user accessing services using the `ncoreapi` service, as described in *nCore API Documentation* and *Cryptographic API*, or by using the utilities provided on the installation media.

3.1.1. Multi-tenancy ready

The system has been prepared for use in multi-tenant systems. In the current firmware version, only one instance of the `ncoreapi` service is allowed to run at any one time. Future versions of firmware will allow multiple instances of the `ncoreapi` service to run concurrently.

3.2. Platform services

Several platform services are provided which perform the tasks associated with the installation, commissioning, and maintenance of the HSM firmware and hardware. These run independently of the `ncoreapi` service.

The platform services are

| Service name | Function |
|-----------------------|--|
| <code>updater</code> | This services provides functions to upgrade the HSM firmware |
| <code>setup</code> | This service provides functions to view the HSM 'lifetime' data installed in the factory and to return the HSM to factory settings |
| <code>monitor</code> | This service provides functions to retrieve and clear logs stored within the HSM |
| <code>sshadmin</code> | This service provides functions to manage the SSH keys used by the platform services and the <code>ncoreapi</code> service |
| <code>launcher</code> | Launcher service. On versions with CodeSafe 5 support, this is used for starting CodeSafe 5 applications on the HSM. |

The administration of platform services is described in [Administration of platform services](#)

An interlock mechanism prevents most platform services from being accessed when the `ncoreapi` service is in operational mode:

- Non-invasive services that only access information, such as log retrieval or a firmware version check, can be used while `ncoreapi` is running.
- Invasive services that would change the platform's state, such as log clearing or firmware updates, cannot be used while `ncoreapi` is running.

To access invasive platform services the `ncoreapi` service must be put into maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

For example:

```
>nopclearfail -M -m 1
Module 1, command ClearUnitEx: OK
```

3.3. Separation of services

Each of the platform services and the `ncoreapi` service has its own communication channel with the host PC that is protected by use of SSH encryption. The procedure for installing the necessary SSH keys is described in [Set up communication between host and module](#).

4. nShield 5s modes of operation

This chapter describes the use of nShield 5s modes of operation:

- [Modes of operation](#)
- [Check and change the mode of operation](#)
- [Return to factory state](#)
- [Recovery mode](#)

4.1. Modes of operation

The status of the nShield 5s HSM can only be one of the following:

| Status | Description |
|---------------------------|---|
| Starting up | The nShield 5s HSM is booting up and performing self tests. After all tests complete successfully, the HSM enters Operational mode. |
| Operational mode | The nShield 5s HSM is working and ready to perform cryptographic operations. An initialized HSM enters Operation mode automatically after it is powered up and all pre-tests are successfully completed. To enter Operational mode manually, see Check and change the mode of operation . |
| Emulated maintenance mode | The nShield 5s HSM is ready to receive maintenance commands, or is processing a maintenance command. The HSM remains in Emulated maintenance mode until you change mode manually, see Check and change the mode of operation . |
| Pre-initialization mode | The nShield 5s HSM is ready to receive initialization commands. For example, initialization commands to set the root-of-trust key (KNSO), to create a Security World, or load an existing Security World. To enter Pre-initialization mode, see Check and change the mode of operation . |
| Initialization mode | The nShield 5s HSM is processing an initialization command. After the command completes, the HSM will return to Pre-initialization mode. |
| Uninitialized mode | The nShield 5s HSM was booted with no root-of-trust key (KNSO) set. This typically happens after leaving a factory state, see Return to factory state . To resolve this, switch to Pre-initialization mode, set the KNSO and reboot the HSM. |
| Error | The nShield 5s HSM is in an error state, see HSM status and error codes . No cryptographic operations can be performed until this error has been cleared. |

| Status | Description |
|---------------|---|
| Recovery mode | The nShield 5s HSM is running on the recovery image instead on the primary image. See Recovery mode . |
| Factory state | The nShield 5s HSM is in a factory state. See Return to factory state . |

4.2. Check and change the mode of operation

You must change the mode on the nShield 5s HSM to perform certain maintenance and configuration tasks. The nShield 5s HSM does not have a physical mode switch. Switch between modes using the `nopclearfail` utility.

Use the following commands to change the mode of an nShield 5s HSM:

| Command | Resulting mode |
|---|---------------------------|
| <code>nopclearfail --maintenance -M</code> | Emulated maintenance mode |
| <code>nopclearfail --operational -O</code> | Operational |
| <code>nopclearfail --initialization -I</code> | Pre-initialization |

1. Run the `nopclearfail` command specifying the module number and the new mode.

When finished, the system responds with `OK`. This message is not confirmation that the module has changed mode.

```
nopclearfail --maintenance --module 1
Module 1, command ClearUnitEx: OK
```

2. Confirm the new mode of the module by running the `enquiry` command.

The `mode` line of the `Module` section displays the current mode.

```
enquiry -m1
Module #1:
enquiry reply flags none
enquiry reply level Five
serial number XXXX-XXXX-XXXX
mode Emulated maintenance mode. hsmadmin may be used to perform module management whilst in
this mode.
module type code 14
product name NC5536E/NC5536N
device name #1 Secure Shell nshield-XXXX-XXXX-XXXX.local
hardware status OK
```

4.3. Return to factory state

nShield 5s HSMs that are delivered from the factory contain no data relating to the `ncoreapi` service. A small amount of 'lifetime' data, which is used by the platform services, is pre-installed. This data is for personalisation and identification of the individual HSM, such as its ESN.

You can perform a reset operation that returns the data stored in an HSM to the state it was in when it left the factory. This erases user credentials and information, leaving only the 'lifetime' data.

When an HSM is in this state it will not support any user commands other than `hsmadmin enroll` and it will be necessary to follow the process described in [Installation of SSH keys](#) before any further actions can be taken.



Returning to factory state will erase any optional features that were not installed at the factory. See, [Enabling optional features](#).



Returning to factory state will change the key used to sign system logs. You should make a record of the new log verification key as soon as possible after returning an HSM to factory state. See [Verifying Signed Logs](#) for more information. Signed system logs are only available from firmware version 13.5 onwards so this is not necessary for HSMs running older firmware.

4.3.1. Purpose of factory state

The main reason for returning an nShield 5s HSM to factory state is to securely erase all user secrets. This is important when, for example:

- The HSM is being taken out of service.
- The HSM is being moved from one domain to another, where it is important to ensure that there is no possibility of secrets being leaked between domains.
- The HSM is being returned to Entrust for servicing or warranty.
- You have lost the SSH keys used to communicate with the HSM, see [Recovery from loss of SSH keys](#)

4.3.1.1. Recovery from loss of SSH keys

Returning a unit to factory state will be necessary if you have lost possession of the SSH keys used to communicate with the HSM and you have not previously made a backup of

those keys with `hsmadmin keys backup` (or `hsmadmin keys backup --passphrase` if the HSM is being re-installed in a different machine). If this happens, returning the HSM to factory state will allow `hsmadmin enroll` to successfully create new keys and re-establish communication with the HSM.

4.3.2. Enter and exit the factory state

The nShield 5s HSM can be returned to factory state in one of two ways. Either by use of `hsmadmin factorystate` or by placing the HSM in [Recovery mode](#).

If the SSH keys used to communicate with the HSM have been lost, only the [Recovery mode](#) option is possible. Both of the above methods include a reboot of the HSM.



The command `hsmadmin factorystate` is prohibited if the system logs have exceeded a maximum size, see [maximum log size](#) or if the system clock is invalid, see [System interaction with the system clock](#). In these situations you can only return to factory state by placing the HSM in [Recovery mode](#).

The HSM is taken out of factory state by use of `hsmadmin enroll`.

4.4. Recovery mode

nShield 5s HSMs are loaded with two different firmware images:

- The Primary image.
- The Recovery image.

During normal operation, the HSM is running firmware that is loaded from the Primary image.

If required, the HSM can be forced into recovery mode to run firmware loaded from the Recovery image. Entry into recovery mode performs the same actions as `hsmadmin factorystate`

Recovery mode is useful in the following cases:

- To return the HSM to a known good state for disaster recovery.
- To retrieve the `init` log if the HSM fails to boot into primary mode, see [Retrieving the init log](#)
- To clear the system log if the HSM is prohibiting actions because it has exceeded the maximum log size, see [Maximum log size](#)

- To restore communication with the HSM if the SSH keys have been lost and no backup is available, see [Set up communication between host and module](#).

4.4.1. Restrictions in recovery mode

The main purpose of recovery mode is to allow essential maintenance activities that are not possible in when the nShield 5s is running the primary image firmware.

The `ncoreapi` and `launcher` services don't run when the nShield 5s is in recovery mode. Only the platform services are available, meaning that only the commands described in [Administration of platform services](#) are available.

If you run `hsmadmin enroll` in recovery mode, a warning will appear. This is because the certificates for the SSH keys described in [Set up communication between host and module](#) are not created in recovery mode. You can ignore this warning.

Commands that use `ncoreapi` or `launcher` service do not run and may show error messages.

4.4.2. Entry into recovery mode

Boot the nShield 5s HSM into recovery mode by holding down the recovery mode button on the back panel of the HSM and then rebooting the HSM. You must continue holding down the button for 60 seconds after initiating the reboot. The button is non-latching.



You must hold down the recovery mode button while the HSM is rebooting. If you reboot the HSM and then press and hold down the button, you will miss the part of the reboot process in which you can change the mode of the HSM.

See the appropriate *Installation Guide* for your nShield HSM for the location of the recovery mode button. You can trigger a reboot with `hsmadmin reset` or by power cycling the host machine containing the HSM.

If you cannot reach the recovery mode button and enter the reboot command simultaneously, you might need to connect a keyboard, mouse, and monitor to the back of the server hosting the HSM. If this is not possible, you need a second person to pass the command to the HSM while you hold down the button, or to hold down the button while you pass the command.

Entering and exiting recovery mode return the HSM to factory state. You must run `hsmadmin enroll` after the boot has completed before any further actions can be performed.

Run `hsmadmin status` to verify that the HSM is in recovery mode. If you are still in primary

mode, try the process again, making sure that the recovery mode button is pressed down before or as soon as the reboot command is passed, and that it is held for the allotted time.

4.4.3. Exit from recovery mode

Exit recovery mode by booting the nShield 5s HSM without the recovery mode button held down. If the firmware is changed whilst in recovery mode using `hsmadmin upgrade`, the unit automatically reboots.

When the unit next boots into primary mode it will be in factory state. You must run `hsmadmin enroll` again before any further actions can be performed.

If you exited recovery mode using `hsmadmin reset`, or as part of a firmware upgrade, you must restart the hardserver/nFast server after running `hsmadmin enroll`.

Run `hsmadmin status` to verify that the HSM is in the correct mode.

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

After you have successfully installed the Security World Software, as described in the *Installation 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. If the SSH keys have not been set up, create the communication path between the host machine and the HSM, as described in [Set up communication between host and module](#).



If you followed the steps in the *Installation Guide* when installing the software, this should already be set up.

3. If necessary, perform additional software and HSM configuration tasks, as described in [Client software and module configuration](#):
 - Set up client configuration, as described in [Setting up client cooperation](#)
 - Set nShield specific environment variables, as described in [Setting environment variables](#)
 - Configure logging and debugging parameters, as described in [Logging and debugging](#)
 - Configure Audit Logging, as described in [Configuring audit logging](#)
 - Configure Java support for KeySafe, as described in [Configuring Java support for KeySafe](#)
 - Configure the hardserver, as described in [Configuring the hardserver](#)
4. Create and configure a Security World, as described in [Creating a Security World](#).
5. Create an OCS, as described in [Creating Operator Card Sets \(OCSs\)](#).

6. Managing the system clock

The HSM's date and time are used to validate security certificate expiry dates and to provide accurate timestamps for system and audit logs. Thus ensuring that an HSM's system clock is closely synchronized with an external time source is critical for maintaining robust security and audit capabilities.

The external time source used is the clock of the host PC in which the HSM is installed.



Make sure that the date and time on the host machine are set correctly according to the documentation for the operating system on the host machine.

The system clock should have been set as part of the installation process, see [Setting the system clock](#) but the clocks on the HSM and the host PC can drift apart over time due to hardware and environmental differences and must be periodically adjusted to keep them in synchronization.

The system clock may also be incorrect if the HSM has been running on battery power for some time.

6.1. System interaction with the system clock

It is important that the system clock is accurate so that timestamps in the system logs can be correlated across the whole network in which the HSM is operating.

For HSMs running firmware version 13.5 or later, if the system clock is lost, for instance due to the HSM running on battery power for an extended period of time, the following administrator actions will be prohibited until the system clock is restored:

- Factory state, see [Return to Factory State](#)
- Firmware upgrade, see [Firmware upgrade](#)
- Setting the minimum VSN, see [Version Security Number](#)
- Setting SSH keys, see [Set up communication between host and module](#)
- Log expiry, see [Expiring signed logs](#)
- Log export, see [Retrieving signed logs](#)
- Adjusting system time, see [Adjusting the system clock](#)
- Administration of CodeSafe, see [The csadmin tool](#)

See [Setting the system clock](#) for more information on setting the system clock.

6.2. Checking the system clock

The system clock can be checked using the command `gettime`.

If the system clock is incorrect it can be set by using the command `settime`.

For small errors in the system clock it is recommended to adjust the clock as described in [Adjusting the system clock](#). For larger errors in the system clock it may be necessary to use the procedures described in [Setting the system clock](#) since use of the `--adjust` option would result in the clock being incorrect for the long period of time required to converge the clocks.



Following the procedures in [Setting the system clock](#) will require the HSM to be taken out of operational mode.

6.3. Adjusting the system clock

The system clock is adjusted by using the `settime` command with the `--adjust` command line option. This will gradually reduce any difference in time between the host and the HSM's clocks, preventing large jumps or discontinuities in time.

Use of the `--adjust` parameter allows the system time to be adjusted whilst the HSM is in operational mode.

The procedure gradually converges the clocks at the rate of a few seconds per day and thus it may take a long time to correct a large error.

When you execute `hsmadmin settime adjust`, the command immediately returns `HSM system time calibration in progress` to acknowledge that the calibration process has started. There is no notification when the calibration is complete.

The frequency at which the `hsmadmin settime --adjust` command should be used depends on multiple factors, for example the precision of the internal clock of the HSM host PC and the extent of drift between the host's clock and the HSM's clock.

The recommended starting point for most systems would be to issue the command once per day. Experimentation is required to find the optimal frequency.

6.4. Setting the system clock

The system clock is set by using the `settime` command without the `--adjust` command line option. This will immediately set the HSM system clock to that of the host PC, but can only be performed when the HSM is in maintenance mode.



The `settime` command uses UTC date and time format.

6.4.1. Restrictions on setting the clock

To prevent malicious actors from tampering with the HSM's system clock by moving it backward, the HSM is designed to prevent the setting of a time and date that is earlier than a previously set time and date. This ensures that the HSM's system clock remains secure and accurate and helps prevent unauthorized access that could occur if the system clock were tampered with.

When the `settime` command is issued without the `--adjust` parameter, the new time is saved within the HSM. Subsequent `settime` commands are prohibited from setting a time earlier than this saved time.

It is only possible to set the HSM system clock to an earlier time than the saved time by returning the HSM to factory state first, see [Return to Factory State](#). This will erase the saved time within the HSM and allow you to issue the `settime` command without restriction.



As a security measure it is never possible to set the HSM clock to a time earlier than its manufacturing time. The manufacturing time can be obtained with the command `hsmadmin info` and is shown as the `mfgtime` entry.



Setting the system date and time automatically resets the HSM.

7. Set up communication between host and module

7.1. Overview of SSH keys

Communications between the host and the HSM are protected by use of SSH secure channels. To allow mutual authentication of the endpoints, the SSH protocol uses separate key pairs in the host and the HSM. The functionality within the HSM is divided into different services that use separate SSH channels. See [Platform services and ncoreapi](#). You need to install the SSH keys for each service before you can use those services.



From firmware versions 13.5 onwards the SSH keys are further protected by an internally generated certificate. This certificate binds an SSH key to the ESN of the module which generated the key. Existing warrants validate that the HSM is a genuine Entrust module with that same ESN. The combination of the certificate and warrant provides a way to validate that the SSH keys have not been tampered with after being generated.

The internal certificates are generated each time the HSM is factory-stated. If your HSM has been upgraded from a firmware version earlier than 13.5 you must factory state your HSM to generate the certificates. See [factory state](#) for more information.



Entrust recommends that you back up the `sshadmin` key as described in [Making a backup](#) whenever SSH keys are installed or changed, if your security policy allows.

7.2. Installation of SSH keys as part of software installation

The `hsmadmin enroll` command automates the installation of SSH keys.

On Linux, this command is run automatically as part of the software installation script. On Windows, this command must be run immediately after installation of the software.

See the *Installation Guide* for your HSM for further details.

7.3. Installation of SSH keys independently of a software installation

If the HSM has been returned to factory state, either with the `hsmadmin factorystate` command or by booting the HSM in recovery mode, as described in [factory state](#), you must install the SSH keys with the `hsmadmin enroll` command before any other actions can be performed.

The `hsmadmin enroll` command can be run on a module in which SSH keys have already been installed. In such a system, the command detects that valid keys already exist and takes no action.

If you are installing SSH keys due to their accidental loss or erasure, and you have previously made a backup of the `sshadmin` key using `hsmadmin keys backup`, then you can install them without returning your HSM to factory state by passing the path to the backed-up `sshadmin` key to `hsmadmin keys restore`.

The `hsmadmin enroll` command automatically validates certificates as part of the enrollment process and produces a warning if it fails to find a certificate for any service. This warning is expected if the HSM:

- is in recovery mode
- is running a firmware version prior to 13.5
- has been upgraded to a firmware version of 13.5 or later but has not performed a factory state operation since the upgrade.

If you receive this warning in any other circumstance you should contact Entrust support.

7.4. Viewing installed SSH keys

The SSH keys installed on the host and each connected HSM can be viewed using the command `hsmadmin keys show`. All the keys shown are public keys. Private keys are not viewable with this command.

The command also shows the date and time at which the client (host) keys were installed.

7.5. Changing installed SSH keys

If your security policy requires you to change the client (host) SSH keys, you can achieve this with the following method.

1. Print the currently installed keys with the command `hsmadmin keys show`
2. Generate and install new client keys with the command `hsmadmin keys roll`. See [SSH Client Key Protection](#) for information about protection options that can be set on keys

during generation.



Linux-only

The hardserver must be restarted in order to be able to use the new `ncoreapi` SSH client key after performing this operation, for example, with `/opt/nfast/sbin/init.d-ncipher restart`.

3. Verify that the new keys have been installed with the command `hsmadmin keys show`

It is not possible to change the server (HSM) keys with this method. Should you be required to change the server keys, this can only be achieved by returning the unit to factory state with the `hsmadmin factorystate` command or by booting the HSM in recovery mode, see [nShield 5s modes of operation](#).

7.6. Making a backup of installed SSH keys

If your security policy allows it, make a backup of your private client key for the `sshadmin` service so that communication with the HSM can be re-established if the installed keys are erased or otherwise lost.

Do this with command `hsmadmin keys backup` for verbatim copy of the `sshadmin` key with its existing protections (by default, it is tied to the host machine). Use `hsmadmin keys backup --passphrase` to backup the `sshadmin` key with a user-supplied passphrase so that it can be restored on another machine or after a re-installation of the OS if necessary.



The backup key should be protected from unauthorized access. Refer to your security procedures for information on how to store the backup file.

7.7. Restoring SSH keys from backup

If you erase or lose your SSH keys, communication with the HSM will be lost. If you have previously made a backup of those keys using the command `hsmadmin keys backup` you can restore that backup with the command `hsmadmin keys restore`. This command will restore the private client key for the `sshadmin` service and then create keys for all other services.

7.8. Preparing an HSM for use in another host

The client (host) SSH keys must be the same for every HSM connected to the same host.

This will happen automatically if the HSMs are all installed together and are all in factory state. The `hsmadmin enroll` command installs the same client keys in each HSM.

Additional HSMs can be installed in a host at any time and, provided that the new modules are in factory state, the `hsmadmin enroll` command installs the same client keys in the new modules as are currently installed in any existing modules.

If it is necessary to be able to transfer a module from one host to another without returning it to factory state this can be achieved with the following method.

In the method below, the term 'source' refers to the host from which the module will be transferred and the term 'destination' refers to the host to which the module will be transferred.

1. Backup the private `sshadmin` client key from the destination host to a location that can be accessed by the source host, such as a shared drive or a USB stick, with the following command:

```
hsmadmin keys backup --passphrase <FILE>
```

Where `<FILE>` specifies the location of the shared drive or USB stick. You will be prompted to enter and confirm a passphrase to use to protect the key.

2. Make sure that the HSM is in Maintenance mode, then install the destination host private `sshadmin` key on the source host with the following command:

```
hsmadmin keys migrate --privkeyfile <FILE>
```

Where `<FILE>` specifies the location of the file written in the previous step. You will be prompted to enter the passphrase of the key.

3. Remove the module from the source host and install in the destination host.



If the keys are not changed on the destination host, this step may be left indefinitely or until needed. For example, the module could be kept in storage as a cold standby unit.

4. Make sure that the source host is in Maintenance mode, then run `hsmadmin enroll` on it to refresh the list of installed nShield 5s HSMs.

5. Linux only

Restart the hardserver.

6. Make sure that all HSMs are in Maintenance mode and run `hsmadmin enroll` on the

destination host.

8. Administration of platform services

nShield 5s platform services are administered through the unified utility `hsmadmin`, which directs the command to the service that implements the command.

Some commands require elevated privileges by default because both the permissions and the protection settings have an impact on the usability of the keys by non-administrative users. Commands that create keys or modify configuration always require elevated privileges. Elevated privileges mean `root` on Linux, and the built-in local Administrators group (running in an elevated shell) on Windows. If a command requires elevated privileges, this is indicated in the command description.

You can modify the permissions and protection options on service keys to allow particular groups of users to execute commands that require the private key for a given service. See [Permissions on SSH keys](#) and [Setting protection on SSH keys](#).

All of the platform services are administered by a unified utility called `hsmadmin`

8.1. hsmadmin

The `hsmadmin` utility manages the administration of nShield HSMs using different subcommands.

```
hsmadmin <subcommand>
```

You can use one of the following subcommands each time you run `hsmadmin`:

- [factorystate](#)
- [status](#)
- [npkginfo](#)
- [upgrade](#)
- [reset](#)
- [enroll](#)
- [keys](#)
- [logs](#)
- [info](#)
- [settime](#)
- [gettime](#)
- [setminvsn](#)

- [getenvstats](#)
- [cs5](#)

8.1.1. hsmadmin factorystate

This command requires **root** privileges on Linux and the privileges of the built-in local Administrators group on Windows.

Before running this command, place the unit in maintenance mode using **nopclearfail -M -m <MODULEID> -w**.

This command returns an HSM to the state it was in when it left the factory. This securely erases all user credentials and information. It resets the **sshadmin** SSH credential to the default.

```
hsmadmin factorystate [-h] [--timeout <TIMEOUT>] [--esn <ESN>] [--verbose]
```

This command takes the following parameters:

| Parameter | Description |
|------------------|---|
| --timeout | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| --esn | Resets specific modules to factory state. You need to add --esn before each ESN you include in the command, for example: <pre>hsmadmin factorystate --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> If no ESNs are specified, the command resets all connected modules. |
| --verbose | Prints verbose logs. |

8.1.2. hsmadmin status

This command displays the ESN and currently loaded firmware version for discovered HSMs. It also displays whether the current image is a primary or a recovery image. When used with the **--json** option it displays primary firmware version, recovery firmware version, and uboot version.

```
hsmadmin status [-h] [--esn <ESN>] [--timeout <TIMEOUT>] [--verbose] [--json]
```

This command takes the following parameters:

| Parameter | Description |
|------------------------|---|
| <code>--verbose</code> | Prints verbose logs. |
| <code>--json</code> | Prints the HSM firmware version and image version in JSON format. |
| <code>--esn</code> | <p>Displays information for specified HSMs.</p> <p>You need to add <code>--esn</code> before each ESN you include in the command, for example:</p> <pre>hsmadmin status --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> <p>If you do not specify any ESNs, the command displays information for all connected HSMs.</p> |
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |

8.1.3. hsmadmin npkginfo

This command inspects the `npkg` file and displays the metadata.

```
hsmadmin npkginfo [--json] <NPKGFILE>
```

This command takes the following parameters:

| Parameter | Description |
|-------------------------------|--|
| <code>--json</code> | Prints metadata in JSON format. |
| <code><NPKGFILE></code> | Specifies the NPKG-format file to inspect. |

8.1.4. hsmadmin upgrade

This command requires `root` privileges on Linux and the privileges of the built-in local Administrators group on Windows.

Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

This command installs firmware packages in `npkg` format. The command can install both primary and recovery firmware.

```
hsmadmin upgrade [-h] --esn <ESN> [--timeout <TIMEOUT>] [--dry-run] [--force] [--verbose] [--json] <NPKGFILE>
```

This command takes the following parameters:

| Parameter | Description |
|-------------------------------|---|
| <code>--verbose</code> | Prints verbose logs. |
| <code>--json</code> | Prints metadata in JSON format. |
| <code>--dry-run</code> | Don't load the package, just validate it. |
| <code>--force</code> | Ignore warnings and force upgrade to proceed. |
| <code>--esn</code> | Specifies the HSMs in which to load the NPKG file. You need to add <code>--esn</code> before each ESN you include in the command, for example: <pre>hsmadmin upgrade --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321 <NPKGFILE></pre> |
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code><NPKGFILE></code> | Specifies the npkg file to load in to the HSMs. |

8.1.5. hsmadmin reset



Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`. If you run the command while in operational mode, it creates a failed state and you will need to run `nopclearfail -r -m <MODULEID>` to correct it.

This command resets the nShield HSM.

```
hsmadmin reset [-h] [--esn <ESN>]
```

This command takes the following parameters:

| Parameter | Description |
|--------------------|---|
| <code>--esn</code> | <p>Specifies the HSMs to reset.</p> <p>You need to add <code>--esn</code> before each ESN you include in the command, for example:</p> <pre>hsmadmin reset --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> <p>If you do not specify any ESNs, all connected HSMs will be reset.</p> |

8.1.6. hsmadmin enroll

This command requires `root` privileges on Linux and the privileges of the built-in local Administrators group on Windows.

Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

This command configures the SSH keys for the nShield HSM.



Linux-only

The install script calls this command automatically as `hsmadmin enroll --sshadmin-key /root/.ssh/id_nshield5_sshadmin`. This will generate SSH client keys and register them with the units if they have not been previously set up. If the `sshadmin` key is not found in its usual location under `/opt/nfast` then `/root/.ssh/id_nshield5_sshadmin` will be tried instead, so it is convenient to use `hsmadmin keys backup` to backup the key to this location.

If `hsmadmin enroll` is called after install to change the installed units, the hardserver will need to be restarted in order to pick up the configuration changes, for example, by running `/opt/nfast/sbin/init.d-ncipher restart`.

```
hsmadmin enroll [--timeout <TIMEOUT>] [--verbose] [--sshadmin-key <SSHADMIN_KEY>]
```

This command takes the following parameters:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |

| Parameter | Description |
|-----------------------------|---|
| <code>--verbose</code> | Prints verbose logs. |
| <code>--sshadmin-key</code> | Path to backup of <code>sshadmin</code> key to use if not present in the standard location. |

8.1.7. hsmadmin keys

This command requires `root` privileges on Linux and the privileges of the built-in local Administrators group on Windows.

Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

This command is used to manage the SSH keys currently loaded on a module.

```
hsmadmin keys [--timeout <TIMEOUT>] <subcommand>
```

This command takes the following parameter:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |

You can use one of the following subcommands with this command:

- `show`
- `migrate`
- `roll`
- `backup`
- `restore`
- `remote-set`
- `remote-remove`

8.1.7.1. hsmadmin keys show

This subcommand displays the public client and server keys used to communicate with the HSMs. For client keys, it also displays the time stamp held on the associated key file in the host file system.

```
hsmadmin keys show [--json] [--verbose]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|-------------------------------|
| <code>--json</code> | Prints output in JSON format. |
| <code>--verbose</code> | Prints verbose logs. |

8.1.7.2. hsmadmin keys migrate

This subcommand changes the SSHAdmin client key on all connected modules to match a public key. The public key is derived from the private key specified in the subcommand.

```
hsmadmin keys migrate --privkeyfile <PRIVKEYFILE> [--json] [--verbose]
```

This subcommand takes the following parameters:

| Parameter | Description |
|----------------------------|--|
| <code>--json</code> | Prints output in JSON format. |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--privkeyfile</code> | Specifies the file containing the private key to be migrated to. |

8.1.7.3. hsmadmin keys roll

This subcommand changes the client keys for all services.

See [SSH Client Key Protection](#) for information about protection options that can be set on keys during generation.

```
hsmadmin keys roll [--json] [--verbose]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|-------------------------------|
| <code>--json</code> | Prints output in JSON format. |
| <code>--verbose</code> | Prints verbose logs. |

On Linux, the hardserver must be restarted in order to be able to use the new `ncoreapi` SSH client key after performing this operation, for example, with `/opt/nfast/sbin/init.d-ncipher restart`.

8.1.7.4. hsmadmin keys backup

This subcommand makes a backup of the private client key for the `sshadmin` service.



The backup key should be protected against unauthorized access. Refer to your security procedures for information on how to store the backup file.

```
hsmadmin keys backup [--passphrase] <FILE>
```

This subcommand takes the following parameters:

| Parameter | Description |
|-------------------------------|---|
| <code>--passphrase, -p</code> | Replace host key protection with passphrase protection. |
| <code><FILE></code> | Path to file in which to store backup. |

If the `--passphrase` option is not supplied, then the existing `sshadmin` key file will be copied verbatim with whatever existing protections it has. By default, the `sshadmin` key is tied to the host machine and OS install, and will not be usable on another machine.

If the `--passphrase` option is used, then the `sshadmin` key will be loaded and re-encrypted using a user passphrase that must be supplied at the prompt. If the existing `sshadmin` key was also protected with a user passphrase (this is not the case by default), then there will be a prompt for that key's passphrase too. The backup key will not be tied to the host machine in this case, and can be used to re-install the HSM on another machine.

On Linux, the backup file will be generated with owner and group matching the directory in which it is created, and readable by owner only.

8.1.7.5. hsmadmin keys restore

This subcommand restores the private client key for the `sshadmin` service from a backup file that has previously been created with the `hsmadmin keys backup` command.

Once the private client key for the `sshadmin` service has been successfully restored, this command will automatically configure all other SSH keys for the HSM.

```
hsmadmin keys restore <FILE>
```

This subcommand takes the following parameter:

| Parameter | Description |
|-----------|--|
| <FILE> | Path to file previously created by <code>hsmadmin keys backup</code> |

8.1.7.6. hsmadmin keys remote-set

This subcommand installs a specific SSH public key for remote access to one HSM service.

```
hsmadmin keys remote-set <SERVICE> <KEYTYPE> <KEYDATA>
```

This subcommand takes the following parameters:

| Parameter | Description |
|-----------|-------------------------------------|
| <SERVICE> | HSM service to be accessed remotely |
| <KEYTYPE> | SSH public key type |
| <KEYDATA> | SSH public key |

8.1.7.7. hsmadmin keys remote-remove

This subcommand removes a specific SSH public key that had previously been set for remote access and restores the local client key.

```
hsmadmin keys remote-remove <SERVICE>
```

This subcommand takes the following parameter:

| Parameter | Description |
|-----------|--|
| <SERVICE> | HSM service from which to remove remote access |

8.1.8. hsmadmin logs

This command manages the system logs of connected HSMs. These logs are separate from the `ncoreapi` logs. See [Platform services and ncoreapi](#) for more information about platform services and `ncoreapi`.

For more information about system logs, see [System logging](#).

For more information about managing `ncoreapi` logs, see [Audit Logging](#).

```
hsmadmin logs <subcommand>
```

You can use one of the following subcommands with this command:

- [get](#)
- [clear](#)
- [export](#)
- [expire](#)
- [getkey](#)

8.1.8.1. hsmadmin logs get

This subcommand retrieves logs from a connected HSM.



HSMs running firmware version 13.5 or later can produce logs in either a signed or unsigned format. This subcommand will retrieve unsigned logs. To retrieve logs in a signed format, use the [export](#) subcommand.

```
hsmadmin logs get [-h] [--verbose] [--timeout <TIMEOUT>] --esn <ESN> --log <LOG> [--json | --out <OUTFILE>]
```

This subcommand takes the following parameters:

| Parameter | Description |
|---------------------------|--|
| --timeout | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| --esn | Specifies the HSM from which to retrieve logs. Only one ESN can be used in the command to retrieve the logs of one specific HSM. |
| --json | Prints output in JSON format. |
| --out | Write logs to file specified by OUTFILE |
| --verbose | Prints verbose logs. |
| --log | Selects log to be retrieved. Options are system , init |

8.1.8.2. hsmadmin logs clear

This subcommand clears logs from connected HSMs.



The [system](#) log can only be cleared using this command on firmware versions earlier than 13.5. The [system](#) log on HSMs running firmware

version 13.5 or later is cleared using the `expire` command. See [Logging, debugging, and diagnostics](#) for more information. The `init` log can be cleared on all firmware versions using this command.

Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

```
hsmadmin logs clear [-h] [--verbose] [--timeout <TIMEOUT>] [--esn <ESN>] --log <LOG> [--json]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | Specifies the HSMs from which to clear logs. You need to add <code>--esn</code> before each ESN you include in the command, for example: <pre>hsmadmin logs clear --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321 --log <LOG></pre> If you do not specify any ESNs, logs will be cleared from all connected HSMs. |
| <code>--json</code> | Prints output in JSON format. |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--log</code> | Selects log to be cleared. Options are <code>system</code> , <code>init</code> |

8.1.8.3. hsmadmin logs export

This subcommand retrieves and validates signed logs from a connected HSM.



The directory used for storing the log files must exist before running this command.

```
hsmadmin logs export [-h] [--verbose] [--timeout <TIMEOUT>] [--esn <ESN>] [--saved] [--expire] [--json | --out <OUTDIR>]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | Specifies the HSM from which to export logs. |
| <code>--json</code> | Prints metadata in JSON format. |
| <code>--out</code> | Write logs to directory specified by OUTDIR |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--expire</code> | Expire the log after exporting it |
| <code>--saved</code> | If not expired, re-export a previously saved log |

8.1.8.4. hsmadmin logs expire

This subcommand expires saved system logs from a connected HSM.

```
hsmadmin logs expire [-h] [--verbose] [--timeout <TIMEOUT>] [--esn <ESN>] --seq <SEQ_NO> [--json]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | Specifies the HSM from which to expire logs. |
| <code>--json</code> | Prints output in JSON format. |
| <code>--seq</code> | expire the log identified by <code><SEQ_NO></code> |
| <code>--verbose</code> | Prints verbose logs. |

8.1.8.5. hsmadmin logs getkey

This subcommand retrieves the system log signing key from a connected HSM.

```
hsmadmin logs getkey [-h] [--verbose] [--timeout <TIMEOUT>] [--esn <ESN>] [--json | --out <OUTFILE>]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | Specifies the HSM from which to retrieve the log signing key |
| <code>--json</code> | Prints output in JSON format. |
| <code>--out</code> | Write key to file specified by <code><OUTFILE></code> . |
| <code>--verbose</code> | Prints verbose logs. |

8.1.9. hsmadmin info

This command requires `root` privileges on Linux and the privileges of the built-in local Administrators group on Windows.

This command returns information that was loaded in the HSM during manufacturing. This information is persistent even after returning the HSM to factory state.

```
hsmadmin info [-h] [--timeout <TIMEOUT>] [--esn <ESN>] [--verbose] [--json]
```

This command takes the following parameters:

| Parameter | Description |
|------------------------|---|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | Returns information for the HSM identified by <code><ESN></code> . You need to add <code>--esn</code> before each ESN you include in the command, for example: <pre>hsmadmin info --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> If no ESNs are specified, the command returns information for all connected modules. |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--json</code> | Prints output in JSON format. |

8.1.10. hsmadmin settime

This command is used to synchronize the HSM system clock with the clock in the host PC.

See [Managing the system clock](#) for more information on managing the system clock.

This command requires **root** privileges on Linux and the privileges of the built-in local Administrators group on Windows.

To use this command without the **--adjust** parameter, the HSM must be in maintenance mode.



Setting the system date and time without the **--adjust** parameter automatically resets the HSM.

```
hsmadmin settime [-h] [--adjust <adjust>] [--timeout <TIMEOUT>] [--esn <ESN>] [--verbose]
```

This command takes the following parameters:

| Parameter | Description |
|------------------|---|
| --adjust | Optional parameter. If specified an HSM System clock drift calibration is executed. |
| --timeout | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| --esn | <p>Sets the system date and time of specific modules.</p> <p>You need to add --esn before each ESN you include in the command, for example:</p> <pre>hsmadmin settime --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> <p>If no ESNs are specified, the command resets all connected modules. If the adjust parameter is specified, a module reset is not required.</p> |
| --verbose | Prints verbose logs. |

8.1.11. hsmadmin gettime

This command requires **root** privileges on Linux and the privileges of the built-in local Administrators group on Windows.

It returns the system date and time of the HSM.

```
hsmadmin gettime [-h] [--timeout <TIMEOUT>] [--esn <ESN>] [--verbose] [--json]
```

This command takes the following parameters:

| Parameter | Description |
|------------------------|---|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | <p>Returns information for the HSM identified by <ESN>.</p> <p>You need to add <code>--esn</code> before each ESN you include in the command, for example:</p> <pre>hsmadmin gettime --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> <p>If no ESNs are specified, the command returns the HSM system date and time for all connected modules.</p> |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--json</code> | Prints output in JSON format. |

8.1.12. hsmadmin setminvsn

This command requires `root` privileges on Linux and the privileges of the built-in local Administrators group on Windows.

Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

This command sets the minimum VSN number of the firmware which the HSM will in the future accept as an upgrade.

```
hsmadmin setminvsn [-h] [--timeout <TIMEOUT>] [--esn <ESN>] [--verbose] [--json] <VSN>
```

This command takes the following parameters:

| Parameter | Description |
|------------------------|--|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |

| Parameter | Description |
|------------------------|---|
| <code>--esn</code> | <p>Sets the minimum VSN on the HSM identified by <ESN>.</p> <p>You need to add <code>--esn</code> before each ESN you include in the command, for example:</p> <pre>>hsmadmin setminvsn --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321 2</pre> <p>If no ESNs are specified, the command sets the minimum VSN on all connected HSMs.</p> |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--json</code> | Prints output in JSON format. |
| <VSN> | <p>The minimum VSN to set.</p> <p>Once this command is executed, the HSM will no longer accept a command to upgrade to a firmware with a VSN lower than <VSN>.</p> <p>The new minimum VSN cannot be lower than the HSM's current VSN, and cannot be higher than the VSN of the firmware currently installed on the HSM.</p> |

8.1.13. hsmadmin getenvstats

This command returns the environmental monitoring statistics of the HSM.

Environmental monitoring statistics available depend on the model of the HSM, the hardware revision and the version of the firmware installed on the HSM.

For the nShield5 with firmware version 13.3 the available statistics are:

| | |
|----------------------|--|
| uptime | The time since the HSM was last rebooted, in seconds. |
| current_time | The current system time of the HSM. |
| mem_total | Total amount of physical RAM, in kilobytes. |
| misp_temp | Temperature recorded by the MSP sensor, in degrees C. |
| cpu_temp | Temperature recorded by the CPU sensor, in degrees C. |
| crypto_co_proc_temp | Temperature recorded by the cryptographic co-processor sensor, in degrees C. |
| voltage_t1022_core | Voltage drawn by the T1022 core chip. |
| voltage_t1022_ifc_io | Voltage drawn by the T1022 IFC I/O chip. |

| | |
|-------------------------------|---|
| voltage_t1022_serdes | Voltage drawn by the T1022 SERDES chip. |
| voltage_t1022_serdes_io | Voltage drawn by the T1022 SERDES I/O chip. |
| voltage_c292_serdes | Voltage drawn by the C292 SERDES chip. |
| voltage_fpga_serdes | Voltage drawn by the FPGA SERDES chip. |
| voltage_c292_serdes_io | Voltage drawn by the C292 SERDES I/O chip. |
| voltage_fpga_serdes_io | Voltage drawn by the FPGA SERDES I/O chip. |
| voltage_msp_avcc | MSP Analogue Vcc. |
| voltage_ddr4_io_access | Voltage drawn by the DDR4 I/O access chip. |
| voltage_ddr4_io | Voltage drawn by the DDR4 I/O chip. |
| voltage_battery | Voltage supplied by the on-board battery. |
| voltage_pci_bus | Voltage drawn by the PCI bus. |
| max_temp | Highest temperature recorded by any temperature sensor since statistics were reset. |
| min_temp | Lowest temperature recorded by any temperature sensor since statistics were reset. |
| ais31_preliminary_alarm_count | AIS31 (RNG) preliminary alarm count. |
| spi_retries | SPI protocol failure count. |
| sp_i2c_total_failures | MSP430 I2C total failures. |
| sp_i2c_slave_failures | MSP430 I2C slave failures. |
| sp_temp_failures | MSP430 temperature failures. |
| sp_voltage_failures | MSP430 voltage failures. |
| host_bus_exceptions | PCI0 (Host) NPE and PE error count. |
| crypto_bus_exceptions | PCI1 (Crypto) NPE error count. |
| sp_sensor_cmd_failures | Read security processor handshake line failure count. |
| nvm_free_space | Free space on user NVRAM. |
| nvm_wear_level | Wear level on user NVRAM. |
| nvm_worn_blocks | Worn block count on user NVRAM. |
| bios_code | Not used; always reports 'None' |
| dfs_throttling | Whether CPU performance is currently degraded due to excessive heat. |

```
hsmadmin getenvstats [-h] [--timeout <TIMEOUT>] [--esn <ESN>] [--verbose] [--json]
```

This command takes the following parameters:

| Parameter | Description |
|------------------------|---|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--esn</code> | <p>Returns information for the HSM identified by <ESN>.</p> <p>You need to add <code>--esn</code> before each ESN you include in the command, for example:</p> <pre>>hsmadmin getenvstats --esn 1A23-BC45-6789 --esn 9Z87-YX65-4321</pre> <p>If no ESNs are specified, the command returns the environmental monitoring statistics of all connected modules.</p> |
| <code>--verbose</code> | Prints verbose logs. |
| <code>--json</code> | Prints output in JSON format. |

8.1.14. hsmadmin cs5

This command is used to manage some aspects of CodeSafe SEE machines running on the HSM.

See also [csadmin](#) for additional commands related to managing CodeSafe SEE machines.

```
hsmadmin cs5 <subcommand>
```

You can use the following subcommand with this command:

- [stats](#)

The following subcommands are only relevant to the nShield 5c. For information on their usage refer to the *nShield 5c User Guide*

- clientinfo
- genclientinfo
- enroll
- unenroll
- list

8.1.14.1. hsmadmin cs5 stats

This subcommand gets statistics from active SEE machines.

```
hsmadmin cs5 stats [--timeout TIMEOUT] [-u UUID] [--esn ESN] [--json]
```

This subcommand takes the following parameters:

| Parameter | Description |
|------------------------|---|
| <code>--timeout</code> | Time to wait for service response, in seconds. Default 30 seconds, minimum 3s, maximum 120s. |
| <code>--u</code> | UUID of SEE machine from which to obtain statistics. If no UUID is specified, statistics will be retrieved for all running SEE machines. |
| <code>--esn</code> | Returns statistics for the HSM identified by <ESN>. If no ESNs are specified, the command returns statistics for all connected modules. |
| <code>--json</code> | Prints output in JSON format. |

9. System logging

This section describes how you can access the logging information generated by the platform. This information is separate from that produced by `ncoreapi`. See [Platform services and ncoreapi](#) for more information about platform services and `ncoreapi`.

For information about the logging and debugging information generated by `ncoreapi`, see [Logging, debugging, and diagnostics](#) and [Audit Logging](#).

Two system logs are automatically created by the system. These are the `init` log and the `system` log.

The `init` log records information created by the system when it boots.

The `system` log continuously records system level information whenever the system is running.

Both logs record information automatically and there is no user configuration required. The information recorded is determined by the system and there is no user configuration of the level of information recorded.

The commands used to retrieve and clear logs are described in [hsmadmin logs](#).

9.1. Maximum log size

The `init` and `system` log share a common non-volatile storage area within the HSM. This is capable of storing the logs for a long period of normal usage but if the logs are not cleared it could eventually become full.

For HSMs running firmware version 13.5 or later, once the logs approach maximum capacity the system will prevent the following actions and issue warnings until some space is freed by clearing the logs.

- Factory state, see [Return to Factory State](#)
- Firmware upgrade, see [Firmware upgrade](#)
- Setting the minimum VSN, see [Version Security Number](#)
- Setting SSH keys, see [Set up communication between host and module](#)
- Setting or adjusting system time, see [Setting the system clock](#) and [Adjusting the system clock](#)
- Administration of CodeSafe, see [The csadmin tool](#)



If the logs are not cleared promptly when the warnings are received the

logs may reach their capacity and the system will cease to operate.

The **system** log is normally much larger than the **init** log but Entrust recommend that you clear both logs at the same time. This is because there are no specific warnings produced by actions that write to the **init** log. From firmware versions 13.5 onwards the **init** log sized is fixed and it is not necessary to clear the **init** log.



It is not normally possible to clear the **system** log from HSMs running v13.5 or later firmware without first exporting the entries. However it is possible to do this in recovery mode. See [Recovery Mode](#).

9.2. Interaction with the system clock

It is important that the timestamps in the system logs are accurate so that events can be correlated across the whole network in which the HSM is operating.

For HSMs running firmware version 13.5 or later, if the system clock is lost, for instance due to the HSM running on battery power for an extended period of time, a number of administrator actions will be prohibited until the system clock is restored, including:

- Log expiry, see [Expiring signed logs](#)
- Log export, see [Retrieving signed logs](#)

See [System interaction with the system clock](#) for more information.

9.3. Init log

The **init** log records information that is produced each time the system boots.

The amount of information recorded depends on the firmware version loaded on the HSM.

The information is intended for use by Entrust Support to diagnose any issues that cause an HSM to fail to boot and thus the log is normally retrieved from recovery mode, see [Recovery Mode](#).

9.3.1. Retrieving the init log

The **init** log can be retrieved with the following command:

```
hsmadmin logs get --esn <ESN> --log init [--json | --out <OUTFILE>]
```

Entrust recommend that you always direct the output to a file using the `--out` parameter and save the file for forwarding to Entrust Support.

9.3.2. Clearing the init log

The `init` log can be cleared with the following command:

```
hsmadmin logs clear [--esn <ESN>] --log init
```

Before running this command, place the unit in maintenance mode using `nopclearfail -M -m <MODULEID> -w`.

From firmware versions 13.5 onwards it is not necessary to clear the `init` log because it is cleared automatically each time the system successfully boots.

9.4. System log

The `system` log records system level events and warnings.

For HSMs running firmware version 13.5 or later these logs are produced in a signed format. HSMs running firmware earlier than 13.5 produce logs in an unsigned format.

The procedures for managing logs differ depending on whether the logs are signed or unsigned.

9.4.1. Signed system logs

Signed system logs are produced by firmware version 13.5 or later and have a number of benefits.

Signed logs can be verified to ensure that they have not been tampered with. See [Verifying Signed Logs](#).

Signed logs contain a sequence number that prevents unintentional deletion of logs and can be used to help order stored logs and identify any gaps. See [Log Sequence Number](#).

9.4.1.1. Log sequence number

Signed system logs use a sequence number to prevent the unintentional loss of logs and to aid in sorting logs.

The HSM holds an internal sequence number for the system log currently being written. This sequence number is persistent over both reboots and factory state operations.

When the system log is exported see [Retrieving Signed Logs](#) the sequence number is incremented and a new log is started with the new sequence number. The previous log is not deleted by the `export` command but remains stored internally within the HSM.

To prevent unintentional duplication of logs, signed logs can normally only be exported once. If it is required to export the same log for a second time the `--saved` option must be used with the `export` command.

Once a log has been successfully exported it can be cleared from the HSM as described in [Expiring Signed Logs](#). The use of the sequence number in this command ensures that no logs are deleted that have not been exported.



It is possible to export and expire logs in a single command but Entrust do not recommend doing this because if the command fails for any reason there is a risk that logs may be lost. The recommended procedure is to export the logs first and then expire the logs as a separate procedure only once the export has completed successfully.

Each exported log begins with one entry showing its own sequence number and another entry showing the sequence number of the preceding log. This allows logs to be chained together to identify any missing gaps.

9.4.2. Retrieving the system log

9.4.2.1. Retrieving unsigned logs

The `system` log can be retrieved with the following command:

```
hsmadmin logs get --esn <ESN> --log system [--json | --out <OUTFILE>]
```

This command will retrieve the logs in unsigned format and will work for HSMs running any firmware.

9.4.2.2. Retrieving signed logs

For HSMs running firmware version 13.5 or later logs can be retrieved in signed format with the following command:

```
hsmadmin logs export --esn <ESN> [--json | --outdir <OUTDIR>] [--saved] [--expire]
```

This command automatically verifies the logs as part of the export process



Logs should normally be exported only once. If an attempt is made to export a log that has already been exported the command will fail with a warning. If you wish to export a log that has previously been exported you should use the `--saved` option with the above command.



When upgrading to Firmware version 13.5 or later from a firmware version lower than 13.5, there may initially be a saved log in the system created by the previous firmware. You should export this log using the `--saved` option and then expire it.



It is possible to automatically expire logs whilst exporting them by use of the `--expire` option but this is not recommended as it may result in loss of logs should the command fail for any reason.

9.4.3. Clearing the system log

The procedures for clearing the system log differ depending on whether the logs are produced in signed or unsigned format. If your HSM is running firmware version 13.5 or later your logs are produced in signed format. If your HSM is running a firmware version earlier than 13.5 your logs are produced in unsigned format.

The process of clearing a signed log is referred to as expiring.

9.4.3.1. Clearing unsigned logs

For HSMs running firmware versions earlier than 13.5 the log can be cleared with the following command:

```
hsmadmin logs clear [--esn <ESN>] --log system
```



For HSMs running firmware versions later than 13.5 this command will fail. You must follow the procedures described in [Expiring Signed Logs](#).



It is possible to use this command to clear signed logs if the HSM is in recovery mode. See [Recovery Mode](#). Entrust do not recommend this procedure unless instructed to do do by Entrust Support.

9.4.3.2. Expiring signed logs

If your HSM is running a firmware version of 13.5 or later the logs will be in signed format. To prevent unintentional loss of logs, signed logs must be exported before they can be cleared. You can export signed logs by following the procedures at [Retrieving Signed Logs](#).

Logs that have previously been exported can be expired with the following command:

```
hsmadmin logs expire --esn <ESN> --seq SEQ_NO
```

This will expire the log with the sequence number **SEQ_NO**. The sequence number is included as the first line of any exported log.



It is possible to automatically expire logs whilst exporting them by use of the **--expire** option on the export command but this is not recommended as it may result in loss of logs should the command fail for any reason.

9.4.3.3. Verifying signed logs

Signed logs are automatically verified as part of the export process.

It is also possible to verify exported logs at any time in the future should you wish to do so, provided that you have access to the verification key. This verification can be conducted without access to the HSM.

The verification key is persistent over reboots but will be changed by a factory state operation. Therefore it is recommended that you record the verification key as soon as possible after any factory state operation. See [Return to Factory State](#) for information about factory state.



The system log will contain an entry for the factory state event. This log entry will contain the value of the verification key so it will always be possible to obtain the verification key if you forget to record it.

The current log verification key can be obtained with the following command:

```
hsmadmin logs getkey --esn <ESN> [--json | --out <OUTFILE>]
```

This command automatically validates the certificate protecting the verification key and produces a warning if it fails to find a certificate for the key. This warning is expected if the HSM:

- is in recovery mode
- has been upgraded to a firmware version of 13.5 or later but has not performed a factory state operation since the upgrade.

If you receive this warning in any other circumstance you should contact Entrust support.

10. Client Software and module configuration

This chapter describes software and module configuration tasks that you can choose to perform after the initial installation of Security World Software and hardware. See the Installation Guide for more information about hardware and software installation.

You must determine whether particular configuration options are necessary or appropriate for your installation. The additional configuration options described in this chapter can be performed either before or after the creation of a Security World (as described in [Creating a Security World](#)) and an OCS (as described in ["Creating Operator Card Sets \(OCSs\)](#)).

10.1. 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.

10.2. Setting up client cooperation

You can allow an nShield HSM to automatically access the remote file system (RFS) belonging to another nShield HSM and share the Security World and key data stored in the Key Management Data directory. Client hardware security modules that access data in this way are described as *cooperating clients*.

To configure client cooperation for hardware security modules that are not nShield HSMs:

1. Configure the RFS used by your nShield HSM 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>
```



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 deleting 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=0000000000000000000000000000000000000000000000000000000000000000
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=0000000000000000000000000000000000000000000000000000000000000000
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=0000000000000000000000000000000000000000000000000000000000000000
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=0000000000000000000000000000000000000000000000000000000000000000
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

```

10.2.1. Useful utilities

10.2.1.1. 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

10.2.1.2. 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.

10.2.1.2.1. Usage

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

10.2.1.2.2. 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 `KNETI` key to use for nToken or software-based authentication of the RFS.

`--rfs-esn=ESN`

This options specifies 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.

10.2.2. 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**.

10.2.3. Logging and debugging

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. If you are setting up the client to develop software that uses it, you should configure debugging before commencing software development.

10.2.4. Configuring Java support for KeySafe

To use KeySafe, follow the instructions in [Using KeySafe](#).

10.3. Configuring the hardserver

The hardserver handles secure transactions between the HSMs connected to the host

computer and applications that run on the host computer. In addition, the hardserver, for example:

- Controls any Remote Operator slots that the HSM uses
- Loads any SEE (Secure Execution Engine) machines that are to run on the HSM
- Enables Remote Administration and provides the communication channel between the Remote Administration Service and the HSM

The hardserver can handle transactions for multiple HSMs. This does not require configuration of the hardserver. For more information, see [Using multiple modules](#).

The hardserver must be configured to control:

- The way the hardserver communicates with remote HSMs
- The way the hardserver communicates with local HSMs
- The import and export of Remote Operator slots
- The loading of SEE machines on to the HSM when the hardserver starts up
- The number of Dynamic Slots available on the HSM
- The port used to connect to the Remote Administration Service
- Whether a Dynamic Slot needs to be exchanged with slot 0 of an HSM
- Timeout values for nShield Remote Administration Card presence assurance
- Configuring the audit logging destination.

The hardserver configuration file defines the configuration of the hardserver. By default, it is stored in `/opt/nfast/kmdata/config/` (**Linux**) or `%NFAST_KMDATA%\config` (**Windows**), and a default version of this file is created when the Security World Software is installed. See [Overview of hardserver configuration file sections](#) for an overview of the hardserver configuration file, and see "Hardserver configuration files" for detailed information about the various options available through it.



In some previous releases of the Security World Software, hardserver configuration was controlled by environment variables. The use of these variables has been deprecated. If any of these environment variables are still set, they override the settings in the configuration file.

You must load the configuration file for the changes to the configuration to take effect.

To configure the hardserver, follow these steps:

1. Save a copy of the configuration file, `/opt/nfast/kmdata/config/` (**Linux**) or `%NFAST_KMDATA%\config\config` (**Windows**), so that the configuration can be restored if necessary.

2. Edit the configuration file to contain the required configuration. (See "Hardserver configuration files" for descriptions of the options in the configuration file.)
3. Run the `cfg-reread` command-line utility to load the new configuration.



If you changed the `server_startup` section of the hardserver configuration file, you must restart the hardserver instead of running `cfg-reread`. For more information, see [Stopping and restarting the hardserver](#).

4. Test that the hardserver is configured correctly by running the `enquiry` command-line utility.

Check that an HSM with the correct characteristics appears in the output.

5. Test that the client has access to the Security World data by running the `nfkminfo` command-line utility.

Check that an HSM with the correct ESN appears in the output and has the state `0x2 Usable`.

10.3.1. Overview of hardserver configuration file sections

10.3.1.1. Configuring remote HSM connections

A *remote HSM* is an HSM that is not connected directly to the host computer but with which the hardserver can communicate. It can be one of the following:

- A network-connected nShield HSM that is configured to use the host computer as a client computer
- An HSM to which an attended Remote Operator slot is imported for the hardserver's unattended local HSM

(Remote Operator feature only).

You configure the hardserver's communications with remote HSMs in the `server_remotecomms` section of the hardserver configuration file. This section defines the port on which the hardserver listens for communications from remote HSMs. You need to edit this section only if the default port (9004) is not available.

For detailed descriptions of the options in this section, see [server_remotecomms](#).

For information about configuring the Remote Operator feature (Remote Operator slots), as opposed to remote HSMs, see [Remote Operator](#).

10.3.1.2. Hardserver settings

You configure the hardserver's settings in the `server_settings` section of the configuration file.

This section defines how connections and hardserver logging are handled. These settings can be changed while the hardserver is running.

For detailed descriptions of the options in this section, see [server_settings](#).

10.3.1.3. Hardserver performance settings

You configure the hardserver performance settings in the `server_performance` section of the configuration file.

This section determines whether multi-threaded performance scaling is enabled or not. By default, scaling is not enabled. Any changes you make to the settings in this section do not take effect until after you restart the hardserver.

For detailed descriptions of the options in this section, see [server_performance](#).

10.3.1.4. HSM settings

You configure the HSM's settings in the `module_settings` section of the configuration file.

This section defines the settings for the HSM that can be changed while the hardserver is running.

For detailed descriptions of the options in this section, see [module_settings](#).

10.3.1.5. Hardserver start-up settings

You configure the hardserver's start-up settings in the `server_startup` section of the configuration file.

This section defines the sockets and ports used by the hardserver. You need to change this section only if the default ports for privileged or unprivileged connections (9000 and 9001) are not available.



Windows only

You should use the `nt_privpipe_users` option to define the name of the user who is allowed to carry out privileged operations, for example, using the `nopclearfail` utility. See [nt_privpipe_users](#) for more

information.

For detailed descriptions of the options in this section, see [server_startup](#).

10.3.1.6. SEE machines

You configure the hardserver to load SEE machines on start-up in the `load_seemachine` section of the configuration file. The SEE Activation feature must be enabled on the HSM, as described in [Enabling optional features](#).

This section defines the SEE machines and optional user data to be loaded, as well any other applications to be run in order to initialize the machine after it is loaded.

For detailed descriptions of the options in this section, see [load_seemachine](#).

For information about SEE machines, see [CodeSafe applications](#).

10.3.1.7. Remote Operator slots

You configure Remote Operator slots in the `slot_imports` and `slot_exports` sections of the configuration file. These sections define the slots that are imported to or exported from the HSM. This applies to the Remote Operator feature only.

For detailed descriptions of the options in these sections, see [slot_imports](#) and [slot_exports](#).

The Remote Operator feature must be enabled on the HSM, as described in [Enabling optional features](#).

10.3.1.8. Remote file system

Each client's remote file system is defined separately in the `remote_file_system` section of the configuration file with a list of HSMs that are allowed to access the file system on the given client. For information about setting up client cooperation, see [Setting up client cooperation](#).



The `remote_file_system` section is updated automatically when the `rfs-setup` utility is run. Do not edit the `remote_file_system` section manually.

As a reference, for detailed descriptions of the options in this section, see [remote_file_system](#).

10.3.1.9. Audit logging

You configure the hardserver's audit logging in the `auditlog_settings` section of the configuration file.

This section defines the host IP address and port used as the destination for the syslog output of the audit logging capability. Optionally, the audit logging messages can be copied to the hardserver's log file.

For further details see [Audit Logging](#).



The hardserver needs to be restarted for these settings to take effect.

10.3.2. 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.

10.3.2.1. 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.

10.3.2.2. 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.

10.3.2.2.1. 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.

10.3.2.3. 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](#).

10.3.2.4. 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.

10.4. 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.



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

10.5. 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 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
```



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:



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**.

11. 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 Feature Enable Tool.

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 are non persistent and, if the HSM is factory stated, must be enabled again using the features file or activator card.



Dynamic features are identified in [Available optional features](#). If a feature is not identified as dynamic it is a static feature.

Features that have been set in the factory will persist if the module is returned to factory state as described in [return to factory state](#). Features that have been set with the Feature Enable tool will be lost if the unit is returned to factory state and must be enabled again when the HSM is returned to service.

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.

After you have enabled features on a module, you must clear the module to make them available. Clear the module by running the command `nopclearfail --clear --all`



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](#).

11.1. Available optional features

This section lists the features that can be added to the HSM. For details of all available features, contact Sales.

11.1.1. 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.

11.1.2. Elliptic Curve activation

Prior to V13.5 firmware, 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](#)

From firmware V13.5, elliptic curve support is always enabled.

Contact Sales if you require an EC activation.

nShield modules with elliptic curve activation support *MQV (Menezes-Qu-Vanstone)* modes.

11.1.3. 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 Brainpool curves are accelerated ⁴ . | Yes |
| 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 |

¹Accessed via nCore, PKCS#11 and JCE APIs.

²Both Prime and Binary named curves are supported. Refer to [Named Curves](#), below, which lists the most commonly supported elliptic curves.

³Offload acceleration refers to offloading the elliptic curve operation from the main CPU for dedicated EC hardware acceleration.

⁴Binary curves are supported, but are not hardware offload accelerated.

⁵Brainpool curves are supported as named curves via nCore, PKCS#11 and JCE only.

11.1.4. 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 . |

¹Java 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.

To demonstrate the accelerated performance of elliptic signing and verify operations, run

the `perfcheck` utility. See *perfcheck: performance measurement checking tool*.

11.1.5. 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 | |
| | BrainpoolP192t1 | NISTP384 | |
| | BrainpoolP224r1 | NISTP521 | |
| | BrainpoolP224t1 | NISTB163 | |
| | BrainpoolP256r1 | NISTB233 | |
| | BrainpoolP256t1 | NISTB283 | |
| | BrainpoolP320r1 | NISTB409 | |
| | BrainpoolP320t1 | NISTB571 | |
| | BrainpoolP384r1 | NISTK163 | |
| | BrainpoolP384t1 | NISTK233 | |
| | BrainpoolP512r1 | NISTK283 | |
| | BrainpoolP512t1 | NISTK409 | |
| | | NISTK571 | |

11.1.6. 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.

11.1.7. 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.

11.1.8. Secure Execution Engine (SEE)

The SEE is a unique secure execution environment. The SEE features available to you are:

| | |
|----------------------------|--|
| SEE Activation, CodeSafe 5 | <p>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.</p> <p>For more information about how to develop SEE applications, see the <i>CodeSafe 5 Developer Guide</i>.</p> |
|----------------------------|--|

11.1.9. 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.

11.1.10. 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.

11.1.11. 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.

11.1.12. 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).

11.2. 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 serial number. This can be found on the circuit board of

the nShield module.

You can also get the serial number of the nShield HSM with `enquiry`:

```
physical serial 12-A30045
hardware part no PCA10005-02 revision 02
hardware status OK
```

- Run the `enquiry` command and note the Electronic Serial Number of the module.

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. You can use the Feature Enable Tool to save the ESN details to a file. For more information about using the Feature Enable Tool, see [Enabling features](#).

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>.

11.3. Enabling features

11.3.1. Viewing enabled features

The **Feature Enable Tool** can be used to view the status of modules connected to the host or to confirm that a feature has been successfully enabled on all modules connected to the host. To view the status of features, run the tool without a smart card.



Some features do not appear in the default output from the **Feature Enable Tool** because they are no longer sold. To see the status of all features, run `fet --show-all`.

11.3.2. Enabling features with a smart card

When it is launched, the **Feature Enable Tool** automatically scans the smart card readers of all modules attached to a host computer for any Feature Enabling smart cards present in the smart card readers, including imported Remote Operator slots and Dynamic Slots. However, feature enable smart cards do not work in Dynamic Slots.

To enable a new feature with a Feature Enabling smart card from Entrust:

1. Insert the Feature Enabling card from Entrust into a slot available to the module to be updated, excluding any Dynamic Slots.
2. Run the `fet` command-line utility to start the **Feature Enable Tool**.

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](#).

11.3.3. Enabling features without a smart card

The **Feature Enable Tool** can also obtain the Feature Enabling Certificate information supplied by Entrust from a file or from the keyboard.

When you run the **Feature Enable Tool** without a Feature Enabling smart card in an HSM slot, a message similar to the following is displayed. There is a line for the features on each module, and a list of options.

In this example, only one module (`ESN 14BD-B089-E078`) is attached to the host.

```

Feature Enable Tool
=====

ISO Smart Card Support
| Remote Operator
| | Korean Algorithms
| | | SEE Activation (EU+10)
| | | | SEE Activation (Restricted)
| | | | | SEE Activation, CodeSafe 5
| | | | | | Elliptic Curve algorithms
| | | | | | | Elliptic Curve MQV
| | | | | | | | Fast RNG for ECDSA
| | | | | | | | | HSM Speed Rating
Mod Electronic | | | | | | | | |
No. Serial Number
1 14BD-B089-E078 -- Y Y Y N N Y Y Y Y High Speed

0. Exit Feature Enable Tool.
1. Read FEM certificate(s) from a smart card or cards.
2. Read FEM certificate from a file.
3. Read FEM certificate from keyboard.
4. Write table to file.

Enter option :
```



When using the option to read the FEM certificate from a file you must

either enter a fully qualified filename or alternatively run the `fet` command from the directory in which the FEM certificate is stored. You may also use `fet -c [FILENAME]` to specify the filename directly on the command line.

12. 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.

12.1. Creating a Security World

You can use the following to create Security Worlds:

- The `new-world` command line utility
See [Creating a Security World using new-world](#).

12.1.1. 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 on hard disk of the host computer
 - 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.



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.

12.1.2. Security World Files

The Security World infrastructure stores encrypted key material and related data in files on the host. For multiple hosts to use the same Security World, the system administrator must ensure that these files are copied to all the hosts and updated when required.

12.1.2.1. 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 `${NFAST_KMDATA}/local` on Linux, `%NFAST_KMDATA%\local` on Windows.
3. Otherwise, if `NFAST_HOME` is set, use `${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.



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 | <code>world</code> (for each module in the Security World) <code>module_ESN</code> |
| Load a Security World | creates or modifies | (for each module in the Security World) <code>module_ESN</code> |
| Replace an ACS | modifies | <code>world</code> |
| Create an OCS | creates | <code>card_HASH</code> <code>cards_HASH_NUMBER</code> |
| Create a softcard | creates | <code>softcard_HASH</code> |
| Generate a key | creates | <code>key_APPNAME__IDENT</code> |
| Recover a key | modifies | <code>key_APPNAME</code> (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.

12.1.2.2. 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 host 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.

12.1.3. 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.

12.1.3.1. Security World basic options

When you create a Security World, you must always configure the basic options described in this section.

12.1.3.1.1. Cipher suite

Only one Cipher suite is supported and this is SP800-131 compliant.

12.1.3.1.2. 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.



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.



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.

12.1.3.1.3. 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.



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.

12.1.3.1.4. 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.

12.1.3.1.5. 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.

12.1.3.2. 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 or 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](#).

12.1.3.3. 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](#).

12.1.3.4. 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.



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.

12.1.3.5. 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.

12.1.3.5.1. 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.

12.1.3.5.2. 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.

12.1.3.6. 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_KMDATA%\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_KMDATA%\local_N` (**Windows**) directory from the old Security World, you must delete it manually.

12.1.4. Creating a Security World using new-world

12.1.4.1. Before you start

Before you start to create a Security World:

- The HSM must be in pre-initialization mode. See [nShield 5s modes of operation](#) for more about changing the mode.
- You must be logged in to the host computer as one of the following users:
 - `root` (**Linux**).
 - A user who is permitted to create privileged connections (**Windows**).
 - A user in the group `nfast`.
- (**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`.

Follow the directions in this section to create a Security World from the command line with the `new-world` utility.


12.1.4.2. 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 |
|-------------------------------------|---|
| <code>--initialize</code> | <p>This option creates a new Security World, replacing any existing <code>/opt/nfast/kmdata/local/</code> (Linux) or <code>%NFAST_KMDATA%\local</code> (Windows) directory.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>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 <code>/opt/nfast/kmdata/local/</code> (Linux) or <code>%NFAST_KMDATA%\local</code> (Windows) directory in which these reside as <code>%NFAST_KMDATA%\localN</code> (Linux) or <code>/opt/nfast/kmdata/localN</code> (Windows) (where <i>N</i> is an integer assigned depending on how many Security Worlds have been previously saved during overwrites).</p> </div> |
| <code>--factory</code> | This option erases an HSM, restoring it to factory state. |
| <code>--no-remoteshare-cert</code> | This option prevents making the HSM from becoming a target for remote shares. |
| <code>--no-strict-rsa-keygen</code> | If you have not specified a mode parameter you can use the <code>-no-strict-rsa-keygen</code> flag to disable the <code>UseStrongPrimes</code> setting. Otherwise it will be enabled by default. See UseStrongPrimes Security World setting . |
| <code>--mode=MODE</code> | <p><code>FIPS-140-level-3</code> creates a Security World compliant with FIPS 140 Level 3.</p> <p><code>common-criteria-cmts</code> creates a Security World supporting <i>Common Criteria PP 419 221-5</i>.</p> <p>Omitting this option will create a default Security World compliant with FIPS 140 Level 2.</p> |

| Option | Description |
|--|---|
| <code>--no-recovery</code> | <p>This option disables the ability to recovery or replace OCSs and softcard (which is otherwise enabled by default). This is equivalent to setting <code>!r</code>, where the <code>!</code> operator instructs the system to turn off the specified feature (<code>r</code>).</p> <p>By default, <code>new-world</code> 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.</p> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="margin-right: 10px; text-align: center;">  </div> <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>We recommend that you do not disable the recovery and replacement option.</p> </div> </div> <div style="margin-top: 10px;"> <div style="display: flex; align-items: flex-start; margin-bottom: 10px;"> <div style="margin-right: 10px; text-align: center;">  </div> <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>If you set the <code>--no-recovery</code> 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).</p> </div> </div> <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px; text-align: center;">  </div> <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>OCS and softcard replacement cannot be enabled after Security World creation without reinitializing the Security World and discarding all the existing keys within it.</p> </div> </div> </div> |
| <code>--cipher-suite=<CIPHER-SUITE></code> | <p>This option specifies the Cipher suite and type of key that is used to protect the new Security World. <code><CIPHER-SUITE></code> should be set to <code>DLf3072s256mAEScSP800131Ar1</code>.</p> |
| <code>--nso-timeout=<TIMEOUT></code> | <p>This option allows you to specify the time-out (<code><TIMEOUT></code>) for new Security Worlds. By default, an integer given for <code>TIMEOUT</code> is interpreted in seconds, but you can supply values for <code>TIMEOUT</code> in the form <code>N s</code>, <code>N h</code>, or <code>N d</code> where <code>N</code> is an integer and <code>s</code> specifies second, <code>h</code> specifies hours, and <code>d</code> specifies days.</p> |
| <code>--module=<MODULE></code> | <p>This option specifies the module to use (by its <code>ModuleID</code>). If you have multiple modules, <code>new-world</code> initializes them all together.</p> |

| Option | Description |
|--|---|
| <p><code>--acs-quorum=<K>/<N></code></p> | <p>In this option, <code><K></code> specifies the minimum number of smart cards needed from the ACS to authorize a feature. You can specify lower <code>K</code> values for a particular feature. All the <code>K</code> values must be less than or equal to the total number of cards in the set. If a value for <code>K</code> is not specified, <code>new-world</code> creates an ACS that requires a single card for authorization.</p> <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;"></div> <div>When the Security World is created in <code>common-criteria-cmts</code> mode, <code>new-world</code> requires a minimum <code>K</code> of 2.</div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="margin-right: 10px;"></div> <div>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 <code>K=1</code>.</div> </div> <p><code><N></code> 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, <code>new-world</code> creates an ACS that contains a single card.</p> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="margin-right: 10px;"></div> <div>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.</div> </div> <p>This option only takes effect if you are creating a new Security World.</p> |
| <p><code>--reduced-features</code></p> | <p>This option instructs <code>new-world</code> to use a reduced default feature set when creating a Security World. A Security World created with the <code>--reduced-features</code> option has no passphrase recovery; no NVRAM, RTC, or FTO; and no NSO delegate keys. However, such a reduced-features Security World can perform many operations faster than more fully featured Security Worlds.</p> |
| <p><code>--disablepkcs1pad</code></p> | <p>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.</p> <p>PKCS#1 v1.5 signature operations are not affected.</p> <p>PSS and OAEP are not affected.</p> |

| Option | Description |
|--------------------------------------|--|
| <p><code>--pp-min=LENGTH</code></p> | <p>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:</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  Warning: short passphrase. </div> <p>However, the passphrase can still be used.</p> <p>Example:</p> <pre style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">new-world --initialize --acs-quorum=K/N --pp-min=14</pre> <p>If <code>--pp-min=<length></code> is not used, the minimum passphrase length is set to the default value (0).</p> |
| <p><code>--pp-strength</code></p> | <p>This option enables passphrases to have at least one uppercase, lowercase, number, and symbol.</p> <p>If the <code>--pp-strength</code> argument is omitted, the complexity requirements are not enforced.</p> |
| <p><code>--audit-logging</code></p> | <p>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.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  The log destination must have already been set in the hardserver configuration file. See Audit Logging. </div> <p>Audit logging is automatically enabled when the Security World is created in <code>common-criteria-cmts</code> mode.</p> |
| <p><code>--max-keyusage</code></p> | <p>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 <code>--max-keyusage</code> limit is applied to Assigned keys on creation.</p> |
| <p><code>--max-keytimeout</code></p> | <p>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 <i>Ns</i>, <i>Nh</i>, or <i>Nd</i> where <i>N</i> is an integer and <i>s</i> specifies second, <i>h</i> specifies hours, and <i>d</i> 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.</p> |



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.

12.1.4.3. 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 a 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 |
|------------------|---|
| <code>m</code> | This feature makes it possible to add new HSMs into the Security World. This feature cannot be disabled. |
| <code>r</code> | This feature enables OCS and softcard replacement; see Replacing Operator Card Sets . |
| <code>p</code> | This feature enables passphrase replacement; see passphrase replacement and Changing card and softcard passphrase . |
| <code>nv</code> | This feature specifies that ACS authorization is needed to enable nonvolatile memory (NVRAM) allocation. |
| <code>rtc</code> | This feature specifies that ACS authorization is needed to set the real-time clock (RTC); (see Real-time clock (RTC) options). |

| Feature name | Description |
|----------------------|---|
| <code>dsee</code> | This feature specifies that that ACS authorization is needed to enable SEE World debugging. |
| <code>dseeall</code> | This feature enables SEE World debugging for <code>all</code> users. |
| <code>fto</code> | 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:

- `nvr`
- `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.



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.



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).

12.1.4.4. 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 [nShield 5s modes of operation](#) for more about changing the mode.



If the HSM is in the pre-initialization mode, `new-world` prompts you for smart cards and passphrases as required.

12.1.5. 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 host data is stored in the directory to which the `NFAST_KMLOCAL` environment variable points (see [Security World Files](#)). The data in this directory is encrypted. You should:

- Ensure that this directory is backed up regularly.
- Check the file permissions for this directory.

- Ensure that the nFast Administrator role, and any user that you want to be able to create Operator Cards or keys, have write permission for this directory.
- All other valid users must have read permission.



Installation of Security World Software must be performed by a user with Administrator rights that allow read and write operations, and applications to be started and stopped.

The HSM can now be used to create Operator Cards and keys for the new Security World.

12.2. Displaying information about your Security World

To display information about the status of your Security World:

- 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](#).

You can also use KeySafe to view a summarized description of the Security World.

12.2.1. 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 |
|-------------------------------------|---|
| <code>-r --repeat</code> | 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 . |
| <code>-p --preload-client-id</code> | 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 the 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 |

12.2.2. 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.

12.3. Adding or restoring an HSM to the Security World

When you have created a Security World, you can add additional HSMs to it. These additional HSMs can be on the same host computer as the original HSM or on any other host. The HSMs may have previously been removed from the same Security World, that is, the Security World can be restored on an HSM by adding the HSM to the Security World again.

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*.
- After installing additional HSM hardware and restarting host machine, you must stop and then restart the hardserver as described in [Stopping and restarting the client hardserver](#). This ensures that the added HSM is recognized and accessible.
- Have a copy of the Security World data on this host. This is the host data written by Keysafe, the nShield HSM CSP wizard (**Windows**), or `new-world` when you created the Security World. This data is stored in the `local` directory within the Key Management Data directory.



If the Key Management Data directory is not in the default location, ensure that the `NFAST_KMDATA` environment variable is set with the correct location for your installation.

- Be logged in to the host computer as `root` (**Linux**) or as a user who is permitted to create privileged connections (**Windows**). See [Hardserver start-up settings](#) and `server_startup`.
- Have started the HSM in pre-initialization mode. See [nShield 5s modes of operation](#) for more about changing the mode.
- Possess a sufficient number of cards from the ACS and the appropriate passphrases.

Adding or restoring an HSM to a Security World:

- Erases the HSM
- 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 computer's hard disk
- 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 contained 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`.



It is not possible to program an HSM into two separate Security Worlds simultaneously.

Initialization removes any data stored in an HSM's nonvolatile memory (for example, data for an SEE program or NVRAM-stored keys). To preserve this data, you must back it up before initializing the HSM and restore it after the HSM has been reprogrammed. We provide the `nvr-am-backup` utility to enable data stored in nonvolatile memory to be backed up and restored.

In order to continue using existing keys and Operator Cards, you must reprogram the HSM:

- After you upgrade the firmware
- If you replace the HSM
- If you need to add an HSM to an existing Security World.

12.3.1. 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 shortcut 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.

12.3.2. 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 [nShield 5s modes of operation](#) 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 [nShield 5s modes of operation](#) for more about changing the mode.

3. Store the ACS in a safe place.



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.

12.4. 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 Algorithms 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.



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.



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.

12.4.1. 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.

12.4.2. 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.

12.4.3. 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.

12.4.3.1. 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... |
|--|--|
| <code>-k <KEYS> --keys-at-once=<KEYS></code> | 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. |
| <code>-h --help</code> | Obtain information about the options you can use with the utility. |
| <code>-c <CARDSETS> --cardsets-at-once=<CARDSETS></code> | 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). |
| <code>--version</code> | View the version number of the utility. |
| <code>--src-warrant=<src-warrantfile></code> | Specify the location of the warrant file of the source module. |
| <code>--src-module=<MODULE></code> | Specify which module ID to use as the source module. |
| <code>--source=<SOURCE></code> | Specify the path to the folder that contains the source World data. |
| <code>--plan</code> | View the list of steps that will be carried out. |
| <code>--perform</code> | Migrate keys interactively. |
| <code>--dst-warrant=<dst-warrantfile></code> | Specify the location of the warrant file of the destination module. |
| <code>--dst-module=<ModuleId></code> | Specify which module ID to use as the destination module. |
| <code>--debug</code> | 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. |
| <code>--key-logging</code> | 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. |
| <code>--src-prots=<list of source protections></code> | 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 <code>--dst-prots</code> . |
| <code>--dst-prots=<list of destination protections></code> | 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 <code>--src-prots</code> in the source security world. |
| <code>--prots-config=<PATH></code> | Specify a configuration file that lists the source and destination protection pairs for migration. The file must contain pairs of tab-separated protection names <code>src_prot dst_prot</code> , 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.

12.4.4. Migrating keys

12.4.4.1. 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](#)



You must enable all your features on the destination module before migration. Otherwise, the migration will fail.

12.4.5. Migrating keys process



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.



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.



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 procedure. See [nShield 5s modes of operation](#) for more about changing the mode.

12.4.6. Verifying the integrity of the migrated keys

To verify the integrity of the migrated keys, run the `nfmverify` 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 `nfmverify` utility with the **-R** option in combination with the preload utility.

Example:

```
preload --admin=RE nfmverify -R -m1 <application> <key-ident>
```



Do not use the `nfmverify` 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.



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.

12.4.7. 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 |
| card set | name2 (duplicate name and type) hash: XXXXXXXX1 | ocstarget1 |
| card set | name2 (duplicate name and type) hash: XXXXXXXX2 | 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,XXXXXXXX1,XXXXXXXX2" \
--dst-prots "ocstarget1,softcardtarget,ocstarget1,softcardtarget,ocstarget1,ocstarget2"
```

By using a configuration file specified with the `--prot-config` option:

```
migrate-world [OPTIONS] --prot-config=migration.cfg

--- migration.cfg ---
ocs 1    ocstarget1
softcard 1    softcardtarget
c:name1    ocstarget1
s:name1    softcardtarget
XXXXXXXX1    ocstarget1
XXXXXXXX2    ocstarget2
-----
```

12.4.8. 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. | <p>Any migrate-able keys found in the source World already exist in the destination World. The migration utility returns this error if:</p> <ul style="list-style-type: none"> • The keys have already been migrated • All keys are non-recoverable and therefore cannot be migrated. | None. |
| <p>Source module must be specified.</p> <p>Destination module must be specified.</p> <p>Source and Destination modules must be different.</p> <p>Module is not usable.</p> | 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. |
| <p>Source World has indistinguishable cardsets or softcards.</p> <p>Destination World has indistinguishable keys.</p> | There are irregularities in one of the Worlds, but these irregularities do not affect the migration process. | None. |
| <p>Destination World has indistinguishable cardsets or softcards.</p> <p>Source World has indistinguishable keys.</p> <p>Cannot determine protection of keys.</p> | 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. | <p>If the source World is not recoverable, you cannot use the migration utility to migrate keys.</p> <p>Contact Support.</p> |
| <p>Missing security World at PATH.</p> <p>Source world must be specified.</p> | <p>The path for the source World is wrong.</p> <p>There is no World data at the location that was specified when running the migration utility.</p> | 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. | <p>An incorrect path was supplied for the source World data when running the utility.</p> <p>The destination World data does not exist in the default location defined by the environment variable <code>NFAST_KMLOCAL</code> or <code>NFAST_KMDATA</code>.</p> | <p>Run the utility with the correct path to the source World data.</p> <p>Move the source World data to a different location and then copy the destination World data to the default location.</p> <p>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.</p> |
| <p>Cannot find <NAME> utility, needed by this utility.</p> <p><NAME> utility is too old, need at least version <VERSION NUMBER>.</p> | 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. | <p>Reinstall the software.</p> <p>Ensure that the path points to the latest version of the software.</p> |
| 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 <code>--key -logging</code> 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. |

12.5. 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**.

12.6. 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 **nvr-am-backup** utility to enable data stored in nonvolatile memory to be backed up and restored.

In order to erase a module, you must:

- Be logged into your computer as **root (Linux)** or as a user who is permitted to create privileged connections (**Windows**).
- Have started the module in the pre-initialization mode. See [nShield 5s modes of operation](#) for more about changing the mode.



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}).



If you are physically removing the module from the machine where it is installed, run `hsmadmin enroll` after removing the module. This refreshes the list of nShield 5s modules currently installed on the machine.

12.6.1. 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 |
|---|---|
| <code>-e --factory</code> | These options restore a module to its factory state. |
| <code>-m --module=<MODULE></code> | These options specify the <code>ModuleID</code> to use. <code>new-world</code> erases only one module at a time. To erase multiple modules, you must run <code>new-world</code> once for every module that you want to erase. |

12.6.1.1. 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.

12.6.2. 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.

12.6.3. 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.

12.6.3.1. 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.

12.7. 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 softcards](#).



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.



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*.

13. 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](#)).

If you want to use the Remote Operator feature to configure smart cards for use with a remote module, see [Remote Operator](#).

13.1. 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.



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.

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 \(CNG\)](#).

13.1.1. 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](#).

13.1.2. 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.

13.1.3. 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.


13.1.4. 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=<MODULE> | 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> | <p>In this option, <K> is the minimum required number of cards. If you do not specify the value <K>, the default is 1.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <div> <p>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.</p> </div> </div> <p><N> is the total number of cards. If you do not specify the value <N>, the default is 1.</p> |

| Option | Description |
|--|--|
| <code>-N --name=<NAME></code> | 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 <code>-M --name-cards=<NAME></code> options. |
| <code>-M --name-cards</code> | 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 <code>--name='NAME' option. 'createocs</code> prompts for the names of the cards as they are created. Not all applications can display individual card names. |
| <code>-p --persist</code> | This option creates a persistent card set. |
| <code>-P --no-persist</code> | This option creates a non-persistent card set. |
| <code>-R --no-pp-recovery</code> | 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 <code>--pp-recovery</code> option. |
| <code>-q --remotely-readable</code> | This option allows this card set to be read remotely. For information on configuring Remote OCSs, see Remote Operator .  Not required for Remote Administration. |
| <code>-T --timeout=<TIME></code> | This option sets the time-out for the card set. Use the suffix <code>s</code> to specify seconds, <code>m</code> for minutes, <code>h</code> for hours, and <code>d</code> 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 <code>TIME</code> has passed since the OCS was loaded. |
| <code>-e --erase</code> | Specifying this option erases a card (instead of creating a card set). You can specify this option twice in the form <code>-ee</code> to repeatedly erase cards. |



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.



The nShield PKCS #11 library requires Operator Cards with passphrases.



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.

13.1.5. 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.

13.1.6. 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.



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](#).

13.2. 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).



Softcards are not supported for use with the nCipherKM JCA/JCE CSP in Security Worlds that are compliant with FIPS 140 Level 3.



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](#))

13.2.1. 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. If you are working within a FIPS 140 Level 3 compliant Security World, you must provide authorization to create new softcards. The `ppmk` utility prompts you to insert a card that contains this authorization. Insert any card from the ACS. If you insert an Administrator Card from another Security World, `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 prompts you to type a passphrase.

3. 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).

4. 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.

13.2.2. 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 softcard or choose a different operation by clicking one of the menu buttons.

13.2.3. 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 available 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 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 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.



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](#).

13.3. 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 --ignoreauth`. 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 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).



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_KMDATA\local%` (**Windows**) directory after you have erased the last card.



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](#).

13.3.1. 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.

13.3.2. 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.

- 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.

13.3.3. 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 |
|---|---|
| <code>-m --module=<MODULE></code> | These options specify the module number of the module. If you only have one module, <i>MODULE</i> is 1. |
| <code>-e --erase</code> | 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 authorization 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`.

13.3.4. 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.

13.3.4.1. 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.

13.3.4.2. 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 authorization. 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.

13.4. 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 KeySafe or the **nfkminfo** command-line utility. To check which passphrase is associated with a card, you can use 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.

13.4.1. 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.

13.4.1.1. 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.

13.4.1.2. 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](#))

- Replace an Administrator Card Set (see [Replacing an ACS with KeySafe](#))
- Discard a card set (see [Erasing cards with KeySafe](#)).

13.4.2. 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
```



Not all software can give names to individual cards.

13.4.3. Viewing softcards

To view softcards, use KeySafe or the command line. The command line provides several options for viewing softcard information.

13.4.3.1. 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](#).

13.4.3.2. 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"
hk1tu    7fb95888ea2850d4e3ffcc8f0c22100937344308
Keys protected by softcard 7fb95888ea2850d4e3ffcc8f0c22100937344308:
AppName simple      Ident mykey
AppName simple      Ident myotherkey
```

13.4.3.3. 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`).

13.4.4. Verifying the passphrase of a card or softcard

13.4.4.1. 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 |
|--------------------------------------|---|
| <code>--check</code> | This option checks the passphrase. |
| <code>--module=<MODULE></code> | This option specifies the number of the module to use. If you only have one module, <code><MODULE></code> is 1. If you do not specify a module number, <code>cardpp</code> 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.

13.4.4.2. 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 `nfkminfo --softcard-list`).

`ppmk` prompts you to enter the passphrase and then tells you whether the passphrase you entered is correct for the specified softcard.

13.5. 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 softcard and the existing passphrase. Known card passphrase can be changed using 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 currently 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.

13.5.1. 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.



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.

13.5.1.1. 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**.

13.5.1.1.1. 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.



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.

13.5.1.2. 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.



With Security World Software v11.72 and later, passphrases are limited to a maximum length of 254 characters, when using **cardpp**. See [Maximum 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.

13.5.1.3. Changing known softcard passphrase with ppmk



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.

13.5.2. Changing unknown or lost passphrase

13.5.2.1. 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.

13.5.2.2. 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 `nfkm info` 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.

13.6. Replacing Operator Card Sets



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.



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 host 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 token

- KeySafe includes an option to replace OCSs on the Card Operations panel (click the **Replace OCS** navigation button).
- The **rocs** command-line utility provides an interactive method or a command-line only method to replace OCSs.

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 host 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, or if you created the Security World with an early version of the `pkcs-init` command-line utility that did not support OCS and softcard replacement, you cannot recover keys from lost or damaged smart cards or softcards.

- Have created the original OCS using `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 `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 host 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](#).

13.6.1. 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](#).)

- Click the **Card Sets** menu button, or select **Card Sets** from the **Manage** menu. KeySafe takes you to the **List Operator Card Sets** panel.
- 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.

- Select an OCS that you want to replace, and click **Replace card set**.



If an OCS does not have any recoverable keys, it cannot be replaced.

- 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.



Only insert your ACS into a module that is connected to a trusted server.

- 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 protected 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.

13.6.2. 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.

13.6.2.1. 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>**



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**.

13.6.2.1.1. 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**.

13.6.2.12. 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 <code>rocs</code> 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 <i>K</i> of <i>N</i> 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. |

13.6.2.13. 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 <code>mark</code> and <code>unmark</code> commands. |
| Name | The key name. |
| App | The application with which the key is associated. |

| Output | Description |
|--------------|---|
| 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 <i>name1</i> marked for recovery to OCS or softcard <i>name2</i> . |
| 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. |

13.6.2.1.4. 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.



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.

13.6.2.1.5. 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.

13.6.2.16. 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.

13.6.2.17. 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.

13.6.2.18. rescan

This command updates the card set and key information.

13.6.2.19. 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.

13.6.2.110. 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.

13.6.2.111. status

This command lists the currently selected hardware security module and target OCS or softcard.

13.6.2.112. 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.

13.6.2.113. unmark <key-spec>

This command unmarks the listed keys. Unmarked keys are not recovered.

13.6.2.2. 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 necessary.

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` commands to transfer the keys.

13.6.2.3. 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 |
|--------------------------------------|---|
| <code>[number] cardset-number</code> | A value of this form selects the OCS or softcard with the given number from the list produced by the <code>list cardsets</code> command. |
| <code>[name] cardset-name</code> | 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). |
| <code>hash cardset-hash</code> | 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.



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.

13.6.2.4. Specifying keys

The `--keys=<KEYS-SPEC>` option identifies one or more keys. It may have any of the following forms:

| Value | Description |
|------------------------------|--|
| <code>mark key-number</code> | <p>A value of this form selects the key with the given number from the list produced by the <code>list keys</code> command. Examples of usage are:</p> <pre>rocs -t <target_OCS> -k <key_number></pre> <p>and</p> <pre>rocs -t <target_OCS> -k "mark 56"</pre> |

| Value | Description |
|-------------------------------------|---|
| <code>appname_keyident</code> | A value of this form selects keys by their internal application name and <code>ident</code> . You must supply at least one of <code>appname</code> or <code>keyident</code> , but you can use wildcard patterns for either or both in order to select all matching keys. An example of usage is: <pre>rocs -t <target_OCS> --keys="simple:simplekey"</pre> |
| <code>hash keyhash</code> | A value of this form selects the key with the given key hash. An example of usage is: <pre>rocs -t <target_OCS> --keys="hash e364[...]"</pre> |
| <code>--cardset cardset-spec</code> | A value of this form selects all keys protected by a given card set. |

13.7. 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 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.



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.

13.7.1. 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.



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.



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, erase the old Administrator Cards; for more information, see [Erasing cards and softcards](#).

13.7.2. Replacing an Administrator Card Set using `racs`

The `racs` utility creates a new ACS to replace a set that was created with the `new-word` 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. When you have finished replacing the ACS, erase the old Administrator Cards. For more information, see [Erasing cards and softcards](#).

14. 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.



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.



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.

14.1. 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.

14.1.1. 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) | <code>\$JAVA_HOME/jre/lib/ext</code> | <code>%JAVA_HOME%\jre\lib\ext</code> |
| Java Runtime Environment (JRE) | <code>\$JAVA_HOME/lib/ext</code> | <code>%JAVA_HOME%\lib\ext</code> |

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) | <code>\$JAVA_HOME/jre/lib/security</code> | <code>%JAVA_HOME%\jre\lib\security</code> |
| Java Runtime Environment (JRE) | <code>\$JAVA_HOME/lib/security</code> | <code>%JAVA_HOME%\lib\security</code> |

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.

- 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](#).



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.

14.1.1.1. 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
```



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](#).

14.1.2. 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 ServiceLoader 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
```

14.1.3. keytool

Use the Java `keytool` utility to read and edit an nShield KeyStore. You must specify the correct `nCipher.world` 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.world -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.

14.1.4. Using keys

Only the `nCipherKM` provider can use keys stored in an `nShield` KeyStore because the underlying key material is held separately in the Security World.

You can always store `nShield` keys in an `nShield` KeyStore. You can also store keys generated 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 associated 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 presented; for example, the ACL can specify that signing key cannot be used for encryption.

14.1.5. 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 |
|--|--|
| <code>JCECSP_DEBUG</code> | This property is a bit mask for which different values specify different debugging functions; the default value is <code>0</code> . For details about the effects of setting different values for this property, see JCECSP_DEBUG property values . |
| <code>JCECSP_DEBUGFILE</code> | This property specifies a path to the file to which logging output is to be written. Set this property if the <code>JCECSP_DEBUG</code> property is set to a value other than the default of <code>0</code> . 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. |
| <code>protect</code> | 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 <code>module</code> , <code><SOFTCARD_NAME:SOFTCARD_IDENT></code> , or <code>cardset</code> . OCS protection (<code>cardset</code>) uses the card from the first slot of the first usable hardware security module. To find the logical token hash <code><IDENT></code> of a softcard, run the command <code>nfkminfo --softcard-list</code> . |
| <code>module</code> | 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. |
| <code>slot</code> | 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. |
| <code>ignorePassphrase</code> | If the value of this property is set to <code>true</code> , the nCipherKM provider ignores the passphrase provided in its KeyStore implementation. This feature is included to allow the Oracle or IBM <code>keytool</code> utilities to be used with module-protected keys. The <code>keytool</code> utilities require a passphrase be provided; setting this property allows a dummy passphrase to be used. |
| <code>seeintegname</code> | 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. |
| <code>com.ncipher.provider.announcemode</code> | The default value for this property is <code>auto</code> , 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 <code>on</code> forces the provider to advertise all mechanisms at start-up. Setting the value of this property to <code>off</code> forces the provider to advertise no mechanisms at start-up. |

| Property | Function for different values |
|---|---|
| <code>com.ncipher.provider.enable</code> | 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. |
| <code>com.ncipher.provider.disable</code> | 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 <code>com.ncipher.provider.disable</code> property overrides the same mechanism if it is supplied in the value for the <code>com.ncipher.provider.enable</code> property. |

14.1.5.1. 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, <code>debugFunc</code> and <code>debugFuncEnd</code> generate debugging information for functions that call them. |
| 16 | If this property has the bit 5 set, <code>debugFunc</code> and <code>debugFuncEnd</code> 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 <code>ThreadID</code> 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. |
| 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`.

14.1.6. 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.

14.2. 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.



If you are using the `preload` command-line utility in conjunction with the nShield PKCS #11 library, you can create *K/N* OCSs.

14.2.1. 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.

14.2.1.1. 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.

14.2.1.2. 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.

14.2.1.3. 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.

14.2.1.4. 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.

14.2.1.5. 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 throughput, the library performs these operations on the host computer.

14.2.1.6. Mechanisms

The mechanisms currently supported by the nShield PKCS #11 library, including some vendor-supplied mechanisms, are listed in the *Cryptographic API Integration Guide*.

14.2.1.7. 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*.

14.2.2. 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 applications 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 `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` environment variable.

The `CKNFAST_OVERRIDE_SECURITY_ASSURANCES` environment variable uses a semicolon separated 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_ASSURANCES`, the PKCS #11 library fails to perform the operation in question, and issues a warning, 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_OVERRIDE_SECURITY_ASSURANCES](#).

14.2.2.1. 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.

14.2.3. 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 `CKNFAST_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](#).

14.2.3.1. 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 `CKNFAST_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_ACCELERATOR_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.

14.2.3.2. 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.

14.2.3.3. 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.

14.2.3.4. 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 preloaded, 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.

14.2.4. 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.



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>
```



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**.



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.

14.2.4.1. 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.

14.2.4.2. CKNFAST_ASSURANCE_LOG

This variable is used to direct all warnings from the Security Assurance Mechanism to a specific log file.

14.2.4.3. 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 `CKNFAST_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 `hk1tu`. For more information about using `nfkminfo`, see [nfkminfo: information utility](#).

14.2.4.4. 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.

14.2.4.5. 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 |

14.2.4.6. 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.

14.2.4.7. CKNFAST_DEBUGFILE

You can use this variable to write the output for `CKNFAST_DEBUG` (`Path name > file name`).

14.2.4.8. 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.

14.2.4.9. 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`.

14.2.4.10. 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.

14.2.4.11. 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.

14.2.4.12. 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).

14.2.4.13. 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.

14.2.4.14. CKNFAST_NO_SYMMETRIC

If this variable is set, the nShield PKCS #11 library does not advertise any symmetric key operations.

14.2.4.15. 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 handshake 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.

14.2.4.16. 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.

14.2.4.17. CKNFAST_NVRAM_KEY_STORAGE

When this environment variable is set, the PKCS #11 library generates only keys in nonvolatile memory (NVRAM). You must also ensure this environment variable is set in order to delete NVRAM-stored keys.

14.2.4.18. 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`
- `wrapping_crypt`
- `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.

14.2.4.18.1. 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.

14.2.4.18.2. 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.

14.2.4.18.3. 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. Therefore, 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 automatically. 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 insecure 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_ASSURANCES`, attempts to create, generate, or unwrap insecure keys with `CKA_TOKEN=true` are allowed.

14.2.4.18.4. `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.



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 functions 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_PERMITTED`, 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_EncryptInit`, and `C_DecryptInit` check the `CKA_CREATION_DATE` against the current time.

14.2.4.18.5. `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.



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.

14.2.4.18.6. `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.



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.

14.2.4.18.7. `wrapping_crypt`

The `wrapping_crypt` parameter allows you to create keys with insecure combinations of wrap/unwrap and encrypt/decrypt operations.

By default, when `wrapping_crypt` is not supplied as a parameter for `CKNFAST_OVERRIDE_SECURITY_ASSURANCES`, trying to create a key with either `CKA_UNWRAP=true` or `CKA_WRAP=true` and `CKA_DECRYPT=true` or `CKA_ENCRYPT=true` will fail with `CKR_TEMPLATE_INCONSISTENT`.

Combinations such as wrap+encrypt or unwrap+encrypt are prohibited because for some

mechanisms (e.g. counter mode), encrypt and decrypt are the same operation, so allowing encrypt is functionally the same as allowing decrypt.

14.2.4.18.8. unwrap_kek

When a key is transferred into the HSM in encrypted form, the key is usually treated as insecure 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).

14.2.4.18.9. 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).

14.2.4.18.10. 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).

14.2.4.18.11. 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_MECHANISMS` with `CKM_CONCATENATE_BASE_AND_KEY`, this `C_DeriveKey` operation

will not be allowed.

14.2.4.18.12. `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.

14.2.4.18.13. `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.

14.2.4.18.14. `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.

14.2.4.18.15. 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.

14.2.4.18.16. 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 PKCS #11 library Security Assurance Mechanism log information is sent to the specified file. These variables must be set whenever `generatekey` or `KeySafe` are used.



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`.

14.2.4.19. 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.

14.2.4.20. 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.

14.2.4.21. CKNFAST_SHARE_SESSION_KEYS

This variable can take a list of one or more semi-colon (;) 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.

14.2.4.22. 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.

14.2.4.23. CKNFAST_USE_THREAD_UPCALLS

If this variable is set and `CKF_OS_LOCKING_OK` is passed to `C_Initialize`, `NFastApp_SetThreadUpcalls` is called by means of `nfast_usencthread` 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_LOCKING_OK` is not passed, `C_Initialize` fails with `CKR_FUNCTION_FAILED`. (`NFastApp_SetThreadUpcalls` 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.

14.2.4.24. `CKNFAST_LOAD_KEYS`

This variable will load private objects at `C_Login` time, rather than at the first cryptographic operation.

14.2.4.25. `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.

14.2.4.26. `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 `CKNFAST_LOADSHARING` is not set.

14.2.5. 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 initialization 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.

14.2.6. 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 |
| 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.

14.3. 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.



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.

14.4. Microsoft CAPI CSP

We provide a Cryptographic Service Provider (CSP) that implements the Cryptographic API (CAPI) supported in Windows 2003 and later.

14.4.1. 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.



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`, `nfkcheck` 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.

14.4.2. 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.

14.4.3. Supported algorithms

The nShield CSPs support a similar range of algorithms to the Microsoft CSP.

14.4.3.1. Symmetric algorithms

- `CALG_DES`
- `CALG_3DES_112` (double-DES)
- `CALG_3DES`
- `CALG_RC4`

- `CALG_AES_128`
- `CALG_AES_192`
- `CALG_AES_256`

14.4.3.2. 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)

14.4.3.3. 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.

14.4.4. Container storage format



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.

14.4.5. 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 |
|-----------------------|---|
| <code>cspcheck</code> | This utility checks that CSP container files are intact and uncorrupted, and also that referenced key files exist. Use <code>cspcheck</code> in conjunction with <code>nfmcheck</code> , but run <code>nfmcheck</code> first in order to test the integrity of your Security World files. |

| Utility | Description |
|-------------------------------------|--|
| <code>cspimport</code> | 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. |
| <code>cspmigrate</code> | 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 <code>-q</code> to keep the new keys. Enter <code>-f</code> to overwrite new keys with old keys. |
| <code>cspnvfix</code> | Regenerate the NVRAM key counter area for a specified nShield CSP key. |
| <code>csptest</code> | Test the installed Cryptographic Service Providers. |
| <code>csputils</code> | 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. |
| <code>configure-csp-poolmode</code> | The <code>--mscapi</code> option allows HSM Pool mode to be enabled or disabled for the nShield CAPI CSP without using the CSP wizard. |
| <code>keytst</code> | This utility displays information about existing CSP key containers by using the Microsoft CryptoAPI. If you have the appropriate permissions, <code>keytst</code> also allows you to create containers and their keys, as well as delete containers. |



Each of these commands has an `-h` option that displays the usage message for the command.

14.4.6. 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](#).

14.5. 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.

14.5.1. 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](#).

14.5.1.1. Registering the CNG CSP

You can register the nShield CNG CSP with:

- CNG Configuration Wizard
- The `engregister` command-line utility

To register the nShield CNG CSP, the hardware 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. See the Installation Guide for more information.

14.5.1.1.1. 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. 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.

14.5.1.1.2. 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](#).

14.5.1.2. 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 reregister 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
```



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](#).

14.5.1.3. 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
```



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:

```
enginstall32 -U
```

- On 64-bit versions of Windows, run the command:

```
enginstall -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:

```
enginstall32 -i
```

- On 64-bit versions of Windows, run the command:

```
enginstall -i
```

2. Reregister the nShield CNG CSP on your system by running the command:

```
engregister
```

For more information about these command-line utilities, see [Utilities for CNG](#)

14.5.2. 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.

14.5.2.1. Signature interfaces (key signing)

| Interface name | Type of support |
|----------------|-----------------|
| RSA PKCS#1 v1 | Hardware |
| RSA PSS | Hardware |
| DSA | Hardware |

| Interface name | Type of support |
|----------------|-----------------|
| ECDSA_P224 | Hardware |
| ECDSA_P256 | Hardware |
| ECDSA_P384 | Hardware |
| ECDSA_P521 | Hardware |



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.

14.5.2.2. 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 |



MD5 is not supported in FIPS 140 mode.

14.5.2.3. 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 |

14.5.2.4. 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 |

14.5.2.5. Key exchange

| Protocol name | Type of support |
|---------------|-----------------|
| DH | Hardware |
| ECDH_P224 | Hardware |
| ECDH_P256 | Hardware |
| ECDH_P348 | Hardware |
| ECDH_P521 | Hardware |



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.

14.5.2.6. Random Number Generation

| Name | Type of support |
|------|-----------------|
| RNG | Hardware |

14.5.3. 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
```



To list the keys available from the *Security World Key Storage Provider*, run the command `cnglist --list-keys` (without specifying the `--provider` option).

14.5.3.1. 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 --esp="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 --esp="Microsoft Strong Cryptographic Provider"
--key="EXAMPLE_CAPICONTAINER"
  "EXAMPLE_IMPORTED_SIGNATURE_ONLY_CAPICONTAINER"
```

Run the `cnplist` command with the `--list-keys` option to check the success of the key import:

```
cnplist --list-keys
EXAMPLE_IMPORTED_SIGNATURE_ONLY_CAPICONTAINER: RSA
```



The `cnimport` 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](#).

14.5.3.2. 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 `cnimport` command as shown in the following example:

```
cnimport -m
-k "EXAMPLE_RSA_1024"
"IMPORTED_RSA_1024"
```

Run the `cnplist` command with the `--list-keys` option to check the success of the key import:

```
cnplist --list-keys
IMPORTED_RSA_1024: RSA
```

The original key is not deleted from the provider from which it was imported:

```
cnplist --list-keys --provider="Microsoft Software Key Storage Provider"
EXAMPLE_RSA_1024
```



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:

```
cnimport -m -k WIN-KQ1Z6JMCUTB-CA WIN-ncipher-CA
Unable to continue.
This key can not be exported from Microsoft Software Key Storage Provider.
```

14.5.3.3. 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

`cngimport` 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 ncspp         * *
3a2b082a8    CAPICONTAINER       SYWELL\Administrato ncspp         * *
2 containers and 4 keys found.
```



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:  ncspp.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 `ce51a0ee0ea164b993d1edcbf639f2be62c53222` 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.

14.5.4. 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 Counting properties.

The CNG application has to be written such that it calls `NCryptOpenKey` to open a CAPI key explicitly.

14.5.5. Utilities for CNG

Use the `nfmverify` command-line utility to check the security of all stored keys in the Security World. Use `nfminfo`, `nfmcheck`, 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 |
|--------------------------------|------------------------------|---|
| <code>cngimport32.exe</code> | <code>cngimport.exe</code> | This key migration utility is used to migrate Security World, CAPI, and CNG keys to the Security World Key Storage Provider. |
| <code>cnginstall32.exe</code> | <code>cnginstall.exe</code> | This utility is the nShield CNG CSP installer. Only use this utility to remove or reinstall the provider DLLs and associated registry entries manually. |
| <code>cnglist32.exe</code> | <code>cnglist.exe</code> | This utility lists information about CNG CSP. |
| <code>cngregister32.exe</code> | <code>cngregister.exe</code> | This is the nShield CNG CSP registration utility. You can use it to unregister and re-register the |
| <code>ncsvcdep32.exe</code> | <code>ncsvcdep.exe</code> | 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. The |

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).

14.5.5.1. `cngimport`

Use `cngimport` to migrate keys to the Security World Key Storage Provider. For more information, see [Migrating keys for CNG](#).

14.5.5.2. `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:

```
ncksppt.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 uninstalled 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](#).

14.5.5.3. 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
```



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.

14.5.5.4. 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.

14.5.5.5. 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"
```



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
```



Microsoft Certificate Services require that the `certsvc` service is made dependent on the `hardserver`.



Microsoft Internet Information Services require that the `http` service is made dependent on the `hardserver`.

14.5.5.6. `cnlist`

Use `cnlist` to display details of CNG providers, keys, and algorithms.

To list details of the CNG providers, run the `cnlist` command with the `--list-providers`

option:

```
cnplist --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 `cnplist` command with the `--list-algorithms` option:

```
cnplist --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 `cnplist` command with the `--list-algorithms`, `--keystorage`, and `--nc` options:

```
cnplist --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 `englist` command with the `--list-algorithms` and `--provider="ProviderName"` options:

```
englist --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

```

14.5.5.6.1. 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
```

15. Working with CodeSafe

15.1. 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.

15.2. The csadmin utility tool

The `csadmin` tool is used to manage CodeSafe 5 applications.



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.



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
```

15.2.1. 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
```

15.2.1.1. Generate loadable images

Only applicable to CodeSafe developers. See the *CodeSafe 5 Developer Guide*

15.2.1.2. 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
```

15.2.1.3. Sign images

Only applicable to CodeSafe developers. See the *CodeSafe 5 Developer Guide*

15.2.2. 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"
    }
  }
}
```

15.2.3. Start a CodeSafe application



Although multiple SEE machines can be loaded and started at the same time, only one can talk to the `ncoreapi` service.



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](#).

15.2.4. Stop a CodeSafe application

Stopping an SEE machine previously started on the HSM will gracefully terminate its programs 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 example:

```
$ 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
  }
}
```

15.2.5. 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 `csadmin 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
```

15.2.6. 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
```

15.2.7. 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.

15.2.7.1. 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"
    }
  }
}
```


15.2.7.2. 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"
    }
  }
}
```

15.2.7.3. 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
```

15.2.7.3.1. 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"
    }
  }
}
```

15.2.7.3.2. 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"
    }
  }
}
```

15.2.8. 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.

15.2.8.1. 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.

15.2.8.2. 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"
    }
  }
}

```

15.2.9. 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 `--uuiid|-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 | 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 | 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"
          }
        }
      }
    }
  }
}

```

15.2.10. 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.

```

15.2.10.1. 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
```

15.2.10.1.1. 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="
      }
    }
  }
}
```

15.2.10.1.2. 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
```

15.2.10.2. 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
```

15.2.10.2.1. 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 you enable the SSHD while the container is not running, it will start when the container starts, otherwise it starts immediately.

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"
    }
  }
}
```

15.2.10.2.2. 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
  --json                Show response in json format
```

If the container sshd is running, the disable command will also stop it. It will not start again until it has been enabled.

15.2.10.2.3. 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`.

15.2.10.3. 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%lxcb0]:[CONTAINER_PORT] -f -N -p [SSHD_PORT]
launcher@[LISTENING_ADDRESS]%nshieId0
```

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:

```
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.

15.2.11. 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.

15.2.11.1. Create a developer ID certificate

Only applicable to CodeSafe developers. See the *CodeSafe 5 Developer Guide*

15.2.11.2. 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
```

15.2.11.3. 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
```

15.2.11.4. Remove identities

Remove identity certificates using `csadmin ids remove`:



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
```

15.2.11.5. 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
```

15.2.11.6. 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
```

16. Remote Operator

This chapter explains:

- The concept of Remote Operator
- How to configure Remote Operator.



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.

16.1. 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.

16.2. Configuring Remote Operator

This section explains how to configure Remote Operator.

16.2.1. 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 [nShield 5s modes of operation](#) for more about changing the mode.

2. Configure the HSM hardservers on their respective host machines 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.

16.2.2. 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 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
hkm1 391eb12cf98c112094c1d3ca06c54bfe3c07a103
```

16.2.3. 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 `cfg-remoteslots` utility.
- Update the HSM configuration file, see [Configuring hardservers for Remote Operator using the HSM configuration file](#).

16.2.3.1. 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. Edit the `slot_exports` section of the hardserver configuration file by adding lines of the form:

```
local_esn=AAAA-AAAA-AAAA
local_slotid=0
remote_ip=222.222.222.222
```

```
remote_esn=BBBB-BBBB-BBBB
```

- b. Run the `cfg-reread` command-line utility to prompt the hardserver to read the configuration changes.
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-BBBB-BBBB).
 - a. Edit the `slot_imports` section of the hardserver configuration file by adding lines of the form:

```
local_esn=BBBB-BBBB-BBBB
local_slotid=2
remote_ip=111.111.111.111
remote_esn=AAAA-AAAA-AAAA
remote_slotid=0
```

This example assigns the imported slot to ID **2**.

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.

16.2.4. 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 the following:
 - a. Use the `slot_imports` section in the hardserver configuration file to import remote slots from HSMs in the same Security World for each relevant HSM.
 - b. Use the `slot_mapping` section in the hardserver configuration file to define the remote slot which is to be swapped with slot #0 for each relevant HSM.

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.

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.

16.2.5. 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.

16.2.6. 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. Re-read the hardserver configuration file by running `cfg-reread`.
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
```

16.2.7. 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.

16.3. Creating OCSs and keys for Remote Operator

When you have configured the HSMs and hardservers for Remote Operator, 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.

16.3.1. 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, 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.



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.



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.



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.

16.3.2. 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](#).



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.

16.3.3. 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](#).



If you generate keys protected by a Remote OCS on the attended module, then you must copy the files in the Key Management Data directory on the attended machine to the unattended module.



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.

16.3.4. Configuring the application

After you have configured the HSMs and hardservers for Remote Operator, 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.

17. 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

17.1. 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.



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.



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.

17.1.1. 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, lowercase ASCII letters, and hyphens (-).



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`), see [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.

17.1.1.1. 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.

17.1.2. 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.

17.1.3. Generating NVRAM-stored keys

NVRAM key storage provides a mechanism for generating keys stored in a module's nonvolatile 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 explicitly 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 `preload` 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 `nvr-am-backup` utility to enable the copying of files, including NVRAM-stored keys, between a module's nonvolatile memory and a smart card.

17.2. 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 possible. 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.

17.2.1. 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 |
|---|--|
| <code>--import</code> | This option specifies key importation. |
| <code><OPTIONS></code> | You can specify particular options when running <code>generatekey</code> that control details of key importation. |
| <code><APPNAME></code> | This option specifies the name of the application for which the key is to be imported. This must be an application for which <code>generatekey</code> can generate keys. |
| <code><NAME>=<VALUE></code> | 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.

17.2.1.1. 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.

17.2.2. 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.

17.3. Listing supported applications with generatekey

To list supported applications, use the command:

```
generatekey --list-apps
```

17.4. Retargeting keys with generatekey

The **--retarget** option 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 |
|---|---|
| <code>from-application=<appname></code> | This option specifies the name of the application with which the key is currently associated. |
| <code>from-ident=<keyident></code> | This option specifies the identifier of the key to be retargeted. You can find this identifier by using the <code>nfkminfo</code> 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.

17.5. Viewing keys

You can view existing keys in the Security World using KeySafe, or the `nfkminfo` command-line utility.

17.5.1. Viewing keys with KeySafe

In order to view a list of keys 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 lists all the keys in the Security World. 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](#))

17.5.2. 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 `nfk` 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 '

```

17.6. Verifying Key Generation Certificates with `nfmverify`

The `nfmverify` command-line utility verifies key generation certificates. You can use `nfmverify` to confirm how a particular Security World and key are protected. It also returns some information about the Security World and key.

The `nfmverify` 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.

17.6.1. 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 |
|---------------------------------|--|
| <code>-h --help</code> | This option displays help for <code>nfkmverify</code> . |
| <code>-V --version</code> | This option displays the version number for <code>nfkmverify</code> . |
| <code>-u --usage</code> | This option displays a brief usage summary for <code>nfkmverify</code> . |
| <code>-m --module=MODULE</code> | This option performs checks with module <code>MODULE</code> . |
| <code>-f --force</code> | This option forces display of an output report that might be wrong. |
| <code>-U --unverifiable</code> | This option permits operations to proceed even if the Security World is unverifiable. <div data-bbox="707 1563 775 1628" data-label="Image"> </div> <div data-bbox="860 1538 1329 1646" data-label="Text"> <p>If you need the <code>-U --unverifiable</code> option, there may be some serious problems with your Security World.</p> </div> |
| <code>-v --verbose</code> | This option prints full public keys and generation parameters. |
| <code>-C --certificate</code> | This option checks the original ACL for the key using the key generation certificate. This is the default. |
| <code>-L --loaded</code> | These options check the ACL of a loaded key instead of the generation certificate. |
| <code>-R --recov</code> | This option checks the ACL of the key loaded from the recovery blob. |

| Option | Description |
|--|--|
| <code>--allow-dh-unknown-sg-group</code> | This option allows an operation to proceed even if a Diffie-Hellman key is using an unrecognized Sophie-Germain group. |

17.7. 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.

17.8. 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.

18. 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.



By default, KeySafe uses the same mechanisms and supports the same features and applications as the **generatekey** utility.

18.1. 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.



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_KMDATA%\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.

18.2. 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.

18.3. 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.

18.3.1. 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).



The options listed below are also available from the **Manage** menu.

18.3.2. 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 | <p>Clicking the World menu button opens the World Operations panel, from which you can:</p> <ul style="list-style-type: none"> • Add modules to a Security World • Remove modules from a Security World. <p>You cannot perform these operations on a module that is not in the pre-initialization mode.</p> <p>Do not use the Initialize option. Creating a Security World from KeySafe is deprecated.</p> |
| Card Sets | <p>Clicking the Card Sets menu button opens the List Operator Card Sets panel, from which you can:</p> <ul style="list-style-type: none"> • 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 | <p>Clicking the Softcards menu button opens the List Softcards panel, from which you can:</p> <ul style="list-style-type: none"> • Create a new softcard • Change or recover the passphrase on a softcard • Discard a softcard |

| Menu button | Description |
|-------------|--|
| Keys | <p>Clicking the Keys menu button opens the Keys panel, from which you can:</p> <ul style="list-style-type: none"> • 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.


18.3.3. 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.

18.3.4. 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.



At the top of the Module Status tree is an icon representing the computer on which the running 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 representing the current Security World and your module(s). To the left of each icon is an expand/collapse 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.

18.3.4.1. 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)
- **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)

18.3.4.2. 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. |
| Operational:AccelOnly | The module is an acceleration-only module. |

| Mode | Description |
|------------------------------|---|
| 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).



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** —

18.3.5. 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.

18.3.5.1. 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.

18.3.5.2. 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).

18.4. Errors

If KeySafe detects an error from which it cannot recover, it may display a Fatal Error message.

18.4.1. 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.

18.4.2. 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
3. Restart KeySafe

The firmware upgrade process destroys all persistent data held in a key-management module. If your security system requires that the persistent data held in a key-management module 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.

19. Warrant Management

19.1. 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.

19.1.1. Warranting steps for nShield Solo and nShield Edge



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.



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.

19.1.2. 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 |
|-------------------------------------|---|
| <code>-h --help</code> | Displays the options you can use with the utility. |
| <code>--list</code> | List ESNs of installed warrants |
| <code>--check</code> | List ESNs of known modules and their warrant state |
| <code>--warrant</code> | Perform warrant operations |
| <code>--csr</code> | Perform CSR operations |
| <code>--details=<FILE></code> | Display the module ESN found in the CSR/warrant <i><file></i> |
| <code>--install=<FILE></code> | Install the warrant from <i><file></i> |
| <code>--req=<MODULE></code> | Request a warrant CSR for the given module number/ESN |
| <code>--out=<FILE></code> | Save the new requested CSR to <i><file></i> |
| <code>--verbose</code> | Print extra information about CSR and warrant files |
| <code>--version</code> | Print the version number of the nfwarrant tool |

19.1.2.1. 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.

19.1.2.2. Generate a warrant upgrade request for nShield Solo

Run the following command:

```
$ nf warrant --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.

19.1.2.3. Validate the warrant you receive from Entrust

1. Run the following command:

```
$ nf warrant --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:

```
$ nf warrant --csr --details <file>
```

The following is an example output:

```
XXXX-XXXX-CF11
```

19.1.2.4. Install a warrant for nShield Solo

Run the following command:

```
$ nfwarrant --warrant --install <file>
```

<file> is the signed warrant provided by Entrust.

19.2. 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.

19.3. Warrant management for nShield 5s and nShield 5c

You do not need to manage the warrants for nShield 5s or nShield 5c. 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.

20. 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.

20.1. 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.

20.1.1. `hsmadmin`

Manages the administration of the HSM using different subcommands. See *Administration of platform services* in the *Install Guide* for your nShield HSM.

20.1.2. `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 the hardware status of nShield PCIe HSMs

See the Installation Guide for more information.

20.1.3. `checkmod`

Check modulo exponentiations performed on the module against the test data located in

`opt/nfast/testdata` (**Linux**) or `%NFAST_HOME%\testdata` (**Windows**).

20.1.4. `cfg-mkdefault`

Create a default client configuration file for the hardserver configuration sections.

20.1.5. `cfg-remoteslots`

Configures Remote Operator slot imports and exports. See [Remote Operator](#).

20.1.6. `cfg-reread`

Load the hardserver configuration from the configuration file.

20.1.7. `fet`

- Activate features on an nShield module connected to the host
- View the status of features on a connected module
- 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.



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](#)

20.1.8. `ncdate`

View, set, and update the time on a module's real-time clock.

20.1.9. `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

20.1.9.1. `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](#).



Run this utility only if requested to do so by Support.

`a|hsmdiagnose a|`Run automated health check against connected nShield HSMs. This diagnostic utility saves information to an XML file.

20.1.10. `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`)

For information about changing the nShield HSM mode, see [nShield 5s modes of operation](#).

20.1.11. `nvrām-backup`

Copy files between a module's NVRAM and a smart card, allowing files to be backed up and restored.

20.1.12. `nvrām-sw`

View and modify information about NVRAM areas.

20.1.13. `pubkey-find`

Obtain information of the public key from a certificate or certificate request (in a Base-64 encoded PEM file).

20.2. randchk

Run a universal statistical test on random numbers returned by the module.

20.2.1. retrievewarrants

Retrieves warrants stored within the HSM and writes them to a file.

20.2.2. rtc

View and set the module's real-time clock.

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.

20.2.3. slotinfo

- Obtain information about tokens in a module
- Format a smart card

20.2.4. snmpbulkwalk snmpget snmpgetnext snmptable snmpset snmpstat snmptranslate snmpwalk

Obtain system, module, connection and software information from the SNMP agent.

For more information, see [Using the SNMP command-line utilities](#).

20.2.5. 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.

20.3. 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... |
|----------------------------|--|
| <code>cryptest</code> | Test all defined symmetric cryptographic mechanisms. |
| <code>des_kat</code> | Perform DES known-answer tests. This utility indicates if any of them fail. |
| <code>floodtest</code> | Perform hardware speed-testing by using modular exponentiation. |
| <code>kptest</code> | Test the consistency of encryption and decryption, or of signature and verification, with the RSA and DSA algorithms. |
| <code>ncthread-test</code> | Stress test modules and test nCore API concurrent connection support. |
| <code>perfcheck</code> | Run various tests to measure the cryptographic performance of a module. For more information, see perfcheck: performance measurement checking tool . |
| <code>sigtest</code> | Measure module speed using RSA or DSA signatures or signature verifications. |


| Utility | Enables you to... |
|--------------------------|--|
| <code>ncperfctest</code> | Test the performance of various crypto commands using attached nShield hardware. Available since v12.10 it contains all the functionality in <code>sigtest</code> and <code>floodtest</code> as well as several new features and greater accuracy and throughput capability in performance management. |


20.4. 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... |
|-------------------------|---|
| <code>bulk erase</code> | <p>Erase multiple smart cards including Administrator Cards, Operator Cards, and FEM activation cards, in the same session.</p> <div style="display: flex; align-items: center;">  <div> <p>Do not use the <code>bulk erase</code> utility to erase Administrator Cards from the current Security World.</p> </div> </div> |
| <code>cardpp</code> | <p>Change, verify, and recover a passphrase of an Operator Card. For more information, see:</p> <ul style="list-style-type: none"> • Verifying the passphrase of a card with cardpp. • Changing known passphrases with cardpp. • Changing unknown or lost passphrases. |
| <code>createocs</code> | <p>Create and erase an OCS. For more information, see:</p> <ul style="list-style-type: none"> • Creating an Operator Card Set from the command line. • Erasing cards from the command line. |
| <code>initunit</code> | <p>Initialize an nShield module.</p> <p>For more information, see Erasing a module with initunit.</p> |

| Utility | Enables you to... |
|----------------------------|--|
| <code>generatekey</code> | <p>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:</p> <ul style="list-style-type: none"> • 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. |
| <code>kmfile-dump</code> | Obtain key management information from a Security World's key management data file. |
| <code>migrate-world</code> | Migrate existing keys to a destination Security World. For more information, see Security World migration . |
| <code>mkac1x</code> | <p>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 <code>generate-key</code> utility.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>Ensure that you run the <code>mkac1x</code> 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.</p> </div> </div> |
| <code>new-world</code> | <p>Create and manage Security Worlds on nShield modules.</p> <p>You must use a privileged connection to use this utility with the following parameter:</p> <ul style="list-style-type: none"> • Initialize the HSM (<code>new-world -e/i/l</code>) <p>For more information, see:</p> <ul style="list-style-type: none"> • Creating a Security World using new-world. • Adding a module to a Security World with new-world. • Erasing a module with new-world. |
| <code>nfkmcheck</code> | Check Security World data for consistency. |

| Utility | Enables you to... |
|------------------|---|
| nfkminfo | <p>Obtain information about a Security World and its associated cards and keys. For more information, see:</p> <ul style="list-style-type: none"> • 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. |
| nfkverify | <p>Perform Security World verification.</p> <p>For more information, see Verifying Key Generation Certificates with nfkverify.</p> |
| postrocs | <p>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.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>PKCS #11 keys either have keys_pkcs_um or key_pkcs_uc as the prefix.</p> </div> </div> |
| ppmk | <ul style="list-style-type: none"> • 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 <p>For more information, see:</p> <ul style="list-style-type: none"> • 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 | <p>Load keys into a module before an application is run in another session.</p> |
| racs | <p>Create a new ACS to replace an existing ACS.</p> <p>For more information, see Replacing an Administrator Card Set using racs.</p> |

| Utility | Enables you to... |
|-------------------|--|
| <code>rocs</code> | <ul style="list-style-type: none"> • Restore an OCS from a quorum of its cards • Restore softcards <p>For more information, see:</p> <ul style="list-style-type: none"> • Replacing OCSs or softcards with rocs. • Using rocs from the command line. |

20.5. 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... |
|---|---|
| <code>elftool</code> | Convert ELF format executables into a format suitable for loading as an SEE machine. |
| <code>hsc_loadseemachine</code> | Load an SEE machine into each module that is configured to receive one, then publishes a newly created SEE World, if appropriate. |
| <code>loadsee-setup</code> | Set up the configuration of auto-loaded SEE machines. |
| <code>modstate</code> | View the signed module state. |
| <code>see-sock-serv</code> <code>see-stdioe-serv</code> <code>see-stdioesock-serv</code> <code>see-stdoe-serv</code> | Activate or enable standard IO and socket connections for SEE machines using the <code>bsdlib</code> architecture. |
| <code>tct2</code> (Trusted Code Tool) | Sign, pack, and encrypt file archives so that they can be loaded onto an SEE-ready nShield module. |

20.6. 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... |
|------------------------------|---|
| <code>ckcerttool</code> | Import a certificate as a PKCS #11 <code>CKO_CERTIFICATE</code> object of type <code>CKC_X_509</code> , and optionally, associate it with the corresponding private key. |
| <code>ckcheckinst</code> | Verify the installation of the nShield PKCS #11 libraries. For more information, see Checking the installation of the nCipher PKCS #11 library . |
| <code>ckimportbackend</code> | <p>Generate keys for use with PKCS #11 applications. When you run the <code>generatekey</code> utility to generate PKCS #11 keys, the <code>ckimportbackend</code> utility is executed in the background.</p> <div style="display: flex; align-items: center;">  <div style="border-left: 1px solid black; padding-left: 10px;"> <p>Do not run this utility unless directed to do so by Support.</p> </div> </div> |
| <code>cknfkmid</code> | View values of attributes of PKCS #11 objects. |
| <code>ckshahmac</code> | Perform a PKCS #11 test for vendor-defined <code>SHA1_HMAC</code> key signing and verification capabilities. |
| <code>cksigtest</code> | 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... |
|-------------------------|--|
| <code>ckinfo</code> | View PKCS #11 library, slot, and token information. Use this utility to verify that the library is functioning correctly. |
| <code>cklist</code> | <p>View details of objects on all slots. If invoked with a PIN argument, the utility lists public and private objects. If invoked with the <code>-n (--nopin)</code> option, the utility lists only the public objects.</p> <p>This utility does not output any potentially sensitive attributes, even if the object has <code>CKA_SENSITIVE</code> set to <code>FALSE</code>.</p> |
| <code>ckmechinfo</code> | View details of the supported PKCS #11 mechanisms provided by the module. |
| <code>ckrsagen</code> | Test RSA key generation. You can use specific PKCS #11 attributes for generating RSA keys. |
| <code>cksotool</code> | Create a PKCS #11 Security Officer role, and manage its PIN. |

20.7. 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](#).



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... |
|--|--|
| <code>cspcheck</code> <code>cspcheck64</code> | Check that CSP container files and keys in the <code>%NFAST_KMDATA%</code> directory are intact and uncorrupted and that the referenced key files exist. |
| <code>cspimport</code> <code>cspimport64</code> | Insert keys manually into existing CSP containers. For more information, see Installing the CAPI CSP . |
| <code>cspmigrate</code> <code>cspmigrate64</code> | Move CSP container information for an existing Security World from the registry into the Security World. |
| <code>cspnvfix</code> <code>cspnvfix64</code> | Regenerate the NVRAM key counter area for a specified nShield CSP key. |
| <code>csptest</code> <code>csptest64</code> | Test the installed Cryptographic Service Providers. |
| <code>csputils</code> <code>csputils64</code> | Obtain detailed information about CSP containers. <div style="display: flex; align-items: center;"> <p>You must have Administrator privileges to view or delete machine containers or containers that belong to other users.</p> </div> |
| <code>keytst</code> <code>keytst64</code> | Create, test, and display information about keys and CSP key containers. |
| <code>configure-csp-poolmode</code> <code>configure-csp-poolmode64</code> | Configure HSM Pool mode for the nShield CAPI CSP. |

20.8. 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](#).




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... |
|---|--|
| <p>cngimport</p> | <p>Migrate Security World, CAPI and CNG keys to the Security World Key Storage Provider.</p> <p>For more information, see:</p> <ul style="list-style-type: none"> • 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. |
| <p>cnginstall32</p> <p>cnginstall</p> <p>(nShield CNG provider installer utility)</p> | <p>Remove or reinstall the provider DLLs and associated registry entries manually.</p> <p>For more information, see cnginstall.</p> |
| <p>cnglist32</p> <p>cnglist</p> | <p>View information about CNG providers.</p> <p>For more information, see:</p> <ul style="list-style-type: none"> • 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. |
| <p>cngregister (nShield CNG provider registration utility)</p> | <p>Unregister and re-register the nShield providers manually.</p> <p>For more information, see:</p> <ul style="list-style-type: none"> • Registering the CNG CSP. • Unregistering or reregistering the CNG CSP. • Uninstalling or reinstalling the CNG CSP. • cngregister. |
| <p>cngsoak</p> <p>cngsoak64</p> <p>(nShield CNG soak tool)</p> | <p>Evaluate the performance of signing, key exchange, and key generation by using a user-defined number of threads.</p> <p>For more information, see cngsoak.</p> |

| Utility | Enables you to... |
|--|--|
| <code>ncsvcdep</code> (nShield Service dependency tool) | <p>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.</p> <p>For more information, see:</p> <ul style="list-style-type: none"> • Uninstalling or reinstalling the CNG CSP. • ncsvcdep. |
| <code>configure-csp-poolmode</code> <code>configure-csp-poolmode64</code> | <p>Configure HSM Pool mode for the nShield CNG CSP. For more information, see configure-csp-poolmode.</p> |

20.9. 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... |
|-----------------------|---|
| <code>pollbare</code> | <p>Obtain information about state changes. The functionality of this test utility depends on whether the server or an HSM supports nCore API poll commands.</p> <div style="display: flex; align-items: center;">  <p>To know if your server or HSM supports nCore API poll commands, run the <code>enquiry</code> utility.</p> </div> |
| <code>trial</code> | <p>Test the nCore API commands. You can use this utility interactively or from a script file.</p> |

20.10. 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 |
|----------------------------------|----------------------------|
| <code>nopclearfail -I/M/0</code> | Change the mode of the HSM |
| <code>nopclearfail -c</code> | Clear the module |
| <code>new-world -e/i/l</code> | Initialize the HSM |

21. 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.



See the Installation Guide for more information about installing Security World Software.

21.1. 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, `%PROGRAMFILES%\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.

21.2. 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.

22. 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.

22.1. 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 semi-colon separated list of paths in the `NC_PS_ADDITIONAL_DIRECTORIES` environment variable.

```
$env:NC_PS_ADDITIONAL_DIRECTORIES="C:\Path1;C:\Path2"
```

- 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.

- 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 |
|-----------------------|--|
| <code>AllHosts</code> | (Recommended) Profile is available for all PowerShell hosts, for example both ConsoleHost and ISE, rather just the currently running host. |
| <code>AllUsers</code> | Profile is available for all users on the local machine rather than just the current user. |

Example:

```
Add-nShieldToProfile -AllHosts
```

22.2. 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.

22.3. 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 PowerShell 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 suitable 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 therefore 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` environment variable:

```
$env:NC_PS_INTERACTIVE=0
```

Mode change commands provided in the `nShieldTools.psd1` module:

| Command | Notes |
|---|--|
| <code>Set-nShieldBatchMode</code> | |
| <code>Set-nShieldInteractiveMode</code> | |
| <code>Reset-nShieldCommandMode</code> | 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`.

22.4. 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...
```

22.5. 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 PowerShell, 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:

```
ppmk --new --newpp $passphrase newsftcard
```

23. Preload Utility

23.1. Overview

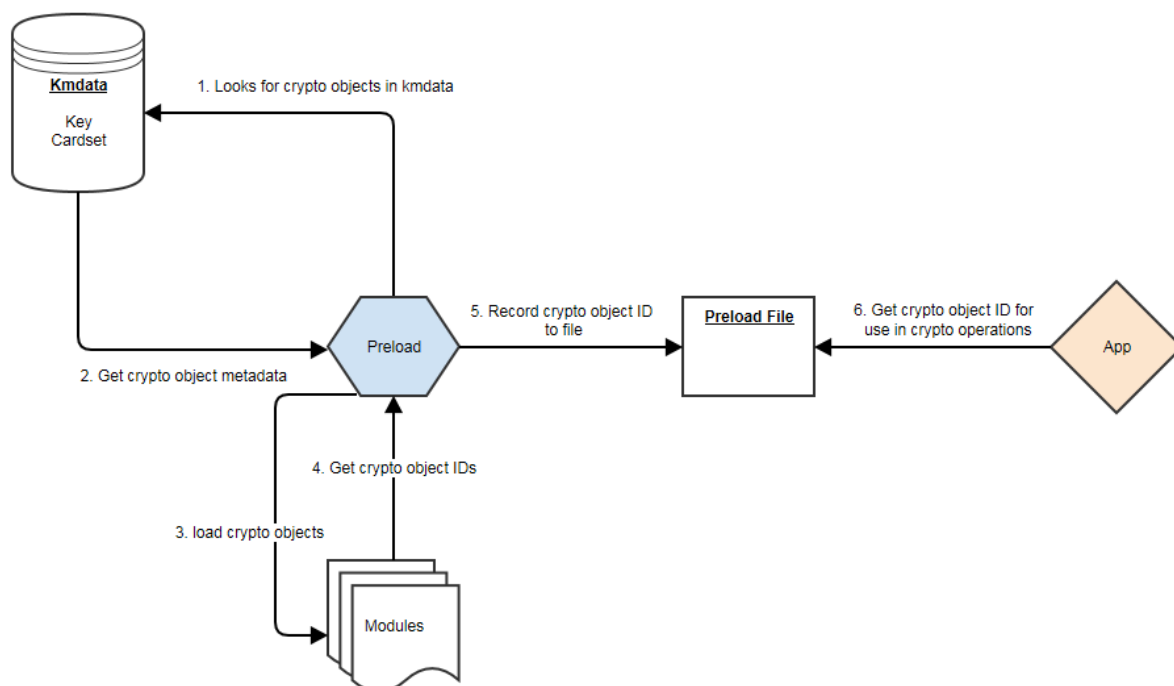
The preload utility loads persistent cryptographic objects (keys/OCS/softcards) onto a chosen set of modules, then makes those objects available for use by applications. This removes the need for applications to load keys/cards themselves, and allows for easy sharing of keys/cards between multiple applications. 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:



23.2. Using Preload

23.2.1. 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
```

23.2.2. 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.

23.2.3. Preload Command Line Arguments

| Argument | Effect |
|------------------|--|
| --version | show program's version number and exit |

| Argument | Effect |
|---|---|
| <code>-h --help</code> | show help message and exit |
| <code>-m MODULE_NUMBER --module=MODULE_NUMBER</code> | Load on specified module (may be repeated; default = all). |
| <code>-c IDENT --cardset=IDENT</code> | Load all cardsets matching <code>IDENT</code> . If <code>IDENT</code> 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 <code>--cardset-name</code> . |
| <code>--cardset-name=NAME</code> | Load cardset(s) named <code>NAME</code> . |
| <code>-s IDENT --softcard=IDENT</code> | Load all softcards matching <code>IDENT</code> . If <code>IDENT</code> looks like a hash it will be interpreted as that, otherwise it will be interpreted as a name. |
| <code>--softcard-name=NAME</code> | Load softcard(s) named <code>NAME</code> . |
| <code>-o --any-one</code> | Load a single cardset. |
| <code>-i --interactive</code> | Load cardsets interactively until told to stop. |
| <code>-A APP --appname=APP</code> | Choose the <code>appname</code> for subsequent <code>-K</code> options. |
| <code>-K IDENT --key-ident=IDENT</code> | Load keys with ident matching <code>IDENT</code> . |
| <code>-n PATTERN --name-pattern=PATTERN</code> | Load keys with name matching <code>PATTERN</code> . Use <code>*</code> for wildcard. |
| <code>--name-exact=NAME</code> | Load keys with name <code>NAME</code> . |
| <code>-M --module-prot</code> | Load all module protected keys, in addition to any others requested. |
| <code>--no-cardset-keys</code> | Do not automatically load keys protected by requested cardsets. Deprecated. |
| <code>--no-token-keys</code> | Do not automatically load keys protected by requested tokens. |
| <code>--admin=KEYS</code> | Load admin keys (separate with commas, or use <code>all</code>). |
| <code>--list-admin</code> | List available admin key names (for <code>--admin</code>). |
| <code>-F --require-fips</code> | Require FIPS-auth to be loaded. |
| <code>-N --no-fips</code> | Do not record FIPS auth, even if available. (overrides <code>-F</code>). |
| <code>-H --high-availability</code> | High availability mode. |
| <code>--polling-interval=POLLING_INTERVAL</code> | Interval (s) between polls for changes to the module list (default=60). High availability mode only. |
| <code>-f PRELOAD_FILE --preload -file=PRELOAD_FILE</code> | Use specified preloaded objects file, instead of the default. |
| <code>-R --reload-everything</code> | Reload keys and tokens that are already loaded. |
| <code>--show-key-info</code> | Display key information for keys as they are loaded. |
| <code>-l --file-logging</code> | Log to file. |

| Argument | Effect |
|-------------------------------------|---|
| <code>-S --no-stderr-logging</code> | Do not log to <code>stderr</code> , this is independent of file logging. |
| <code>--log-file=LOG_FILE</code> | The file destination for the log, defaults to <code>preload_%pid.log</code> in the <code>nfast</code> log directory. |
| <code>--log-level=LOG_LEVEL</code> | The log level to log, options: <code>DEBUG</code> , <code>INFO</code> , <code>WARNING</code> , <code>ERROR</code> . Default is <code>INFO</code> , if unrecognized option it will fall back to default. |

23.2.4. Pattern Matching

Options to preload that use pattern matching, namely `--name-exact` and `--key-ident`, can accept the following wildcards:

| Wildcard | Definition |
|---------------------|----------------------------------|
| <code>*</code> | matches everything |
| <code>?</code> | matches any single character |
| <code>[seq]</code> | matches any character in seq |
| <code>[!seq]</code> | 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
```

23.3. Preload File

The IDs of preloaded crypto objects are persistently stored in a preload file.

Each entry has the following format:

| Element | Description |
|-------------------------|--|
| <code>Hash</code> | The <code>sha1</code> hash of the crypto object. |
| <code>module</code> | The module which this object is present. |
| <code>objectid</code> | The id reference as a <code>M_KeyID</code> . |
| <code>generation</code> | 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`.

23.4. 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 protected by the softcard has been loaded on both modules.

23.4.1. 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 protected 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
```

23.5. 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.

23.6. Admin Keys

23.6.1. 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`.

23.6.2. 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.

23.7. 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.

23.7.1. 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.

23.7.2. 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.

23.7.3. 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.

23.7.4. 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`.



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.

23.7.5. 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=SECONDS`

By default the polling interval is 60 seconds.

23.7.6. 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.

23.7.7. 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
```

23.7.7.1. 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
600bcc26336c13f2371b0bb54b1cde293ded9a15 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 (**Ox44313d41**).

23.7.8. 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
```

23.7.9. 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.

23.7.9.1. Use PKCS #11 for key reloading

PKCS #11 key reloading requires preload to be run in high availability mode, with the

following 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.



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.



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*.

23.7.10. 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`

23.8. 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`.

23.9. 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:

```
import nfk

conn = nfk.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 = nfk.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
```

24. Environment variables

This appendix describes the environmental variables used by Security World Software.



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 hardserver and 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 and Hardserver debugging . | 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 <code>local</code> subdirectory of the Key Management Data directory. See <code>NFAST_KMDATA</code> . | 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 <code>NFAST_HOME</code> , <code>NFAST_KMDATA</code> , and <code>NFAST_CERTDIR</code> . | 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 <code>nshieldlogs</code> 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 <code>ClientID</code> and <code>KeyIDs</code> are stored by the <code>preload</code> command-line utility. | y | y |
| NFAST_SEE_MACHINEENCKEY_DEFAULT | This variable is the name of the <code>SEEConf</code> key needed to decrypt SEE-machine images. Running the command <code>loadmache --encryptionkey='IDENT</code> (or <code>loadmache --unencrypted</code>) overrides any value set by this variable. | y | y |
| NFAST_SEE_MACHINEENCKEY_<module> | This variable is the name of the <code>SEEConf</code> key needed to decrypt the SEE-machine image targeted for the specified module. It overrides <code>NFAST_SEE_MACHINEENCKEY_DEFAULT</code> for the specified module. Running the command <code>loadmache --encryptionkey=<IDENT></code> (or <code>loadmache --unencrypted</code>) 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 <code>machine-filename</code> parameter when running the <code>loadmache</code> command-line utility overrides this variable. This variable is not affected when running the <code>loadsee-setup</code> or <code>hsc_loadseemachine</code> utilities. | y | y |

| Variable | Description | Win | Lnx |
|--|---|-----|-----|
| <code>NFAST_SEE_MACHINEIMAGE_<module></code> | 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 <code>NFAST_SEE_MACHINEIMAGE_DEFAULT</code> for the specified module. Supplying the <code><machine-filename></code> parameter when running the <code>loadmache</code> command-line utility overrides the <code>NFAST_SEE_MACHINEIMAGE_<module></code> variable. This variable is not affected when running the <code>loadsee-setup</code> or <code>hsc_loadseemachine</code> utilities. | y | y |
| <code>NFAST_SEE_MACHINESIGHASH_DEFAULT</code> | This variable is the default key hash of the vendor signing key (<code>seeinteg</code>) 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 <code>loadmache --sighash=<HASH></code> any value set in this variable. | y | y |
| <code>NFAST_SEE_MACHINESIGHASH_<module></code> | This variable is the key hash of the vendor signing key (<code>seeinteg</code>) that signs SEE machine images for the specified module. It overrides <code>NFAST_SEE_MACHINESIGHASH_DEFAULT</code> for the specified module. This variable is only required if you are using a dynamic SEE feature with an encrypted SEE machine. Running the command <code>loadmache --sighash=<HASH></code> any value set in this variable. | y | y |
| <code>NFAST_SERVER</code> <code>NFAST_PRIVSERVER</code> | If these variables are set in the hardserver's environment, 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 hardserver. These variables are available for this purpose for backward compatibility only; you should configure sockets in the hardserver configuration file, see server_startup | y | y |

| Variable | Description | Win | Lnx |
|--|--|-----|-----|
| NFAST_SERVER_PORT NFAST_SERVER_PRIVPORT | <p>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.</p> <p>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.</p> | y | y |
| NFLOG_CATEGORIES | <p>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.</p> | y | y |
| NFLOG_SEVERITY | <p>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.</p> | y | y |
| NFLOG_DETAIL | <p>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.</p> | y | y |
| NFLOG_FILE | <p>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).</p> | y | y |

25. 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.

25.1. Logging and debugging

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.

| | |
|-----------------------------|---|
| <code>CKNFAST_DEBUG</code> | All severities from <code>DL_Call</code> to <code>DL_DetailMutex</code> may result in sensitive data being logged |
| <code>JCECSP_DEBUG</code> | No further guidance |
| <code>NFJAVA_DEBUG</code> | No further guidance |
| <code>NFLOG_DETAIL</code> | The <code>0x00010000</code> flag may result in sensitive data being logged |
| <code>NFLOG_SEVERITY</code> | All the <code>DEBUGn</code> settings may result in sensitive data being logged |

25.1.1. 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.



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 |
| 0x00000010 | This flag shows the external <code>time_t</code> (that is, the time in machine clock ticks rather than local time) with the log entry. | external_time_t |

| Hexadecimal flag | Function | logdetail flags |
|------------------|---|------------------|
| 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 <code>StackPush</code> and <code>StackPop</code> 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 <code>dbg_stub</code> 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 <code>dbg_stubbignum</code> 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 <code>dbg_stubinit</code> 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 <code>dbg_dumpenv</code> flag with the logging functionality found in previous Security World Software releases.) |
| nfkm-getinfo | Logs <code>nfkm-getinfo</code> messages. |
| nfkm-newworld | Logs messages about world generation. |
| nfkm-admin | Logs operations using the Administrator Card Set. |
| nfkm-kmdata | Logs file operations in the <code>kmdata</code> (Linux) or <code>%NFAST_KMDATA%</code> (Windows) directory. |
| nfkm-general | Logs general NFKM library messages. |
| nfkm-keys | Logs key loading operations. |
| nfkm-preload | Logs <code>preLoad</code> 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. |
| serv-startup | Logs fatal startup errors. |
| servdbg-stub | Logs all generic stub debugging messages. |

| Category | Description |
|-------------------|--|
| 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](#).



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.

25.1.2. 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.



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. |



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.



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`
- `nspdddebug.txt`
- `neckspwdebug.txt`

25.1.3. 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 | "> " |
| 10 | DL_Detail | DEBUG5 | "D " |
| 11 | DL_DetailMutex | DEBUG6 | "DM " |

25.1.4. Hardserver debugging

Hardserver debugging is controlled by specifying one or more `servdbg-*` categories (from the `NFLOG_CATEGORIES` environment variable) in the hardserver configuration file; see [server_settings](#). However, unless you also set the `NFAST_DEBUG` environment variable to a value in the range 1 – 7, no debugging is produced (regardless of whether or not you specify `servdbg-*` categories in the hardserver configuration file). This behavior helps guard against the additional load debugging places on the CPU usage. You can set the desired `servdbg-*` categories in the hardserver configuration file, and then enable or disable debugging by setting the `NFAST_DEBUG` environment variable.

The `NFAST_DEBUG` environment variable controls debugging for the general stub or hardserver. The value is an octal number, in the range 1 – 7. It refers bitwise to a number of flags:

| Flag | Result |
|------|--|
| 1 | Generic stub debugging value. |
| 2 | Show bignum values. |
| 4 | Show initial <code>NewClient</code> or <code>ExistingClient</code> command and response. |

For example, if the `NFAST_DEBUG` environment variable is set to 6, flags 2 and 4 are used.



If the `NFAST_DEBUG` environment variable value includes flag 1 (Generic stub debugging value), the `logdetail` value in the hardserver configuration file (one of the values for the `NFLOG_DETAIL` environment variable) controls the level of detail printed.

Do not set the `NFAST_DEBUG` environment variable to a value outside the range 1 – 7. If you set it to any other value, the hardserver does not start.

25.1.5. Debugging information for Java

This section describes how you can specify the debugging information generated by Java.

25.1.5.1. 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**)

25.1.5.2. Setting Java debugging with the command line

You can set the Java debug options by immediately preceding them with a **-D**. Use the **NFJAVA_DEBUGFILE** property to direct output to a given file name, for example:

```
java -DNFJAVA_DEBUGFILE=myfile -DNFJAVA_DEBUG=129 -classpath .... classname
```



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.



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.

25.1.6. appliance-cli: system logs utility

The `appliance-cli` command-line utility is a diagnostics tool that enables certain *appliance* administration or status commands to be executed from a privileged client.

This supports remote administration of nShield Connect and nShield 5c devices, in addition to enabling certain administration or status commands to be made available to privileged clients running in guest VMs that have access to nShield SoloXC or nShield 5s local HSMs running on the host machine over nCipher Secure Transport/Impath.

25.1.6.1. Usage



Run `appliance-cli` with no arguments to display information about the general usage of the command, with `-m 0` to display locally available commands, or with the `-m <MODULE>` option to display the available appliance commands for the specified remote module.

| Appliance command | Action |
|---------------------------|--------------------------------------|
| <code>expirehsmlog</code> | Expire signed nShield 5 HSM syslogs |
| <code>exporthsmlog</code> | Get signed nShield 5 HSM syslogs |
| <code>gethsmlog</code> | Get nShield 5 HSM syslogs |
| <code>gethsmlogkey</code> | Get nShield 5 HSM syslog signing key |
| <code>getservcfg</code> | Get server config file |
| <code>getservlog</code> | Get server logs |



All operations require a privileged client (by default, that means an elevated command-prompt on Windows or membership of the `nfast` group on Linux). Operations run against a remote module additionally require that the module has been enrolled as privileged in the local hardserver.

25.2. Diagnostics and system information



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.

25.2.1. 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.

25.2.1.1. Usage

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 when you submit the `nfdiag` output to Entrust, run:

```
nfdiag -e|--extrainfo <your-plaintext-file>
```

By default, `nfdiag` runs in verbose mode, providing feedback on each command that it executes 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.

25.2.1.2. Output

`nfdiag` generates a plain text output file and displays its file name. It does NOT capture any passphrases in the output file.

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:

- 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`

25.2.2. `nfkminfo`: information utility

The `nfkminfo` utility displays information about the Security World and the keys and card sets associated with it.

25.2.2.1. 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>...]
```

25.2.2.1.1. 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.

25.2.2.1.2. 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.

25.2.2.2. 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
hkns0     hash_kns0
hkm       hash_km
hkmsk     hash_kmsk
hkrc      hash_krc
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_ktu
Module #1 Slot #1 IC 0
generation 1
phystype   SmartCard
slotlistflags 0x2 SupportsAuthentication
state      0x4 Admin
flags      0x10000 passphrase
shareno    1
shares     LTN50(PIN) LTM(PIN) LTR(PIN) LTNV(PIN) LTRTC(PIN) LTDSEE(PIN)
LFTFO(PIN)
error      OK
No Cardset

```


No Pre-Loaded Objects

25.2.2.2.1. 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:

| | |
|-------------------------------------|--|
| <code>Initialised</code> | This indicates that the Security World has been initialized. |
| <code>Usable</code> | This indicates that there is at least one usable HSM in this Security World on this host. |
| <code>!Usable</code> | This indicates that there are no usable HSMs in this Security World on this host. |
| <code>Recovery</code> | This indicates that the Security World has the OCS and softcard replacement and the key recovery features enabled. |
| <code>!Recovery</code> | This indicates that the Security World has the OCS and softcard replacement and the key recovery features disabled. |
| <code>AdminAuthRequired</code> | This indicates that additional authorization is required for the following operations: <ul style="list-style-type: none"> • 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: |
| <code>ExistingClient</code> | This indicates that there is a Client ID set, for example, by <code>preload</code> . This Client ID is given in the <code>ex.client</code> output if the <code>--preload-client-id</code> flag was supplied. |
| <code>!ExistingClient</code> | This indicates that no Client ID is set. The <code>ex.client</code> output will be empty. |
| <code>AlwaysUseStrongPrimes</code> | This indicates that the Security World always generates RSA keys in a manner compliant with FIPS 186-3. |
| <code>!AlwaysUseStrongPrimes</code> | This indicates that the Security World leaves the choice of RSA key generation algorithm to individual clients. |
| <code>SEEDebug</code> | This indicates that the Security World has an SEE Debugging delegation key. |

| | |
|-----------------------------|---|
| <code>!SEEDebug</code> | This indicates the Security World has no SEE Debugging delegation key. |
| <code>SEEDebugForAll</code> | This indicates no authorization is required for SEE Debugging. |
| <code>PINRecovery</code> | This indicates that the Security World has the passphrase replacement feature enabled. |
| <code>!PINRecovery</code> | This indicates that the Security World has the passphrase replacement feature disabled. |
| <code>FTO</code> | This indicates that the Security World has an FTO delegation key. |
| <code>!FTO</code> | This indicates that the Security World has no FTO delegation key. |
| <code>NVRAM</code> | This indicates that the Security World has an NVRAM delegation key. |
| <code>!NVRAM</code> | This indicates that the Security World has no NVRAM delegation key. |
| <code>RTC</code> | This indicates that the Security World has an RTC delegation key. |
| <code>!RTC</code> | This indicates that the Security World has no RTC delegation key. |
| <code>AuditLogging</code> | This indicates that Audit Logging is enabled for this Security World. |
| <code>!AuditLogging</code> | 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. |
| commoncriteria-cmts | 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). |

25.2.2.2.2. 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. |
| Acce1Only | This indicates that the HSM is acceleration only. |
| Unchecked | This indicates that, although the HSM appears to be in the current Security World, <code>nfkminfo</code> could not find a module initialization certificate (a <code>module_<ESN></code> file) for this HSM. |
| Failed | This indicates that the HSM has failed. For nShield PCIe HSMs running firmware 2.61.2 and above, use the <code>enquiry</code> utility for further information about the failure reason. |
| 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).

25.2.2.2.3. 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.

25.2.2.2.4. 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:

| | |
|---|--|
| <code>NotPersistent</code> | This indicates that the Operator Card is not persistent. |
| <code>Persistent</code> | This indicates that the Operator Card is persistent. |
| <code>NotRemoteEnabled</code> | This indicates that the card in the slot is not from a Remote Operator Card Set. |
| <code>RemoteEnabled</code> | This indicates that the card in the slot is from a Remote Operator Card Set. |
| <code>PINRecoveryForbidden(disabled)</code> | 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. |
| <code>PINRecoveryRequired(enabled)</code> | 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.

25.2.3. 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 RSAPKCS1 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 respectively. 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 | <p>This value is the number of outstanding jobs in the queue when the test was run.</p> <p>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 <code>--queue</code> option, in which case only that queue length will be run with, except for some <code>misc</code> suite tests which set their own queue.</p> |
| Rate (Units/s) | <p>This value is a measure of throughput. It is calculated by dividing the number of repetitions by total time.</p> <p>If a test has been rerun to improve accuracy, as is the case by default, then this is the mean across all the runs.</p> <p>Some tests, for example <code>enc</code>, set the Unit to something other than an operation, for example KB, to indicate the amount of data that can be encrypted.</p> |
| Min latency (ms) | <p>This value is the time in milliseconds that the quickest individual job across all the test runs took to round-trip.</p> |
| Mean latency (ms) | <p>This value is the mean time in milliseconds that jobs took to round-trip.</p> <p>If a test has been rerun, this is the mean of the mean latency values from each run.</p> |
| Max latency (ms) | <p>This value is the time in milliseconds that the slowest individual job across all the test runs took to round-trip.</p> |

| Value | Description |
|-----------------|---|
| CV (%) | <p>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.</p> <p>If a test has been rerun, this is the mean of the CV (%) values from each run.</p> |
| Min rate (tps) | <p>This is the estimated lower bound of the throughput for this queue size in transactions per second.</p> <p>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.</p> |
| Mean rate (tps) | <p>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.</p> <p>Mean rate (tps) is included for comparison against the Min rate (tps) and Max rate (tps) figures.</p> |
| Max rate (tps) | <p>This is the estimated upper bound of the throughput for this queue size in transactions per second.</p> <p>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.</p> |
| Reps | <p>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.</p> <p>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 <code>--repetitions</code> option but note that in most cases more repetitions will be run depending on the <code>--accuracy</code> setting provided that the timeout is not reached. It is recommended to set <code>--accuracy</code> rather than <code>--repetitions</code> to control the accuracy of the test instead of adjusting the repetitions.</p> |

25.2.3.1. 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 hardserver.
2. The client hardserver encrypts the command.
3. When the HSM is free, the command is submitted from the hardserver queue.

4. The command is executed by the HSM, and the reply is given to the hardserver.
5. The unit hardserver queues the reply.
6. The unit hardserver sends the command back to the client hardserver over the network.
7. 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 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.

25.2.4. `stattree`: information utility

The `stattree` utility returns the statistics gathered by the hardserver and HSMs.

25.2.4.1. Usage

```
stattree [<node> [<node> [...]]]
```

25.2.4.2. 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 IDs. 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 |

25.2.4.2.1. 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. |

| Category | Contains |
|-----------------------|--|
| 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. |
| 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. |

25.2.4.2.2. 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. |

| ID | Value |
|---------------------------|---|
| 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. |
| 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 . |

| ID | Value |
|---------------------------|--|
| 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 . |
| 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 ModulePCISStats 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. |

| ID | Value |
|--------------------------|---|
| 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. |
| CPULoadPercent | The current processing load on the HSM, represented as a number between 0 and 100. Because an HSM typically contains a number of different types of processing resources (for example, main CPU, and RSA acceleration), this figure is hard to interpret precisely. In general, HSMs report 100% CPU load when all RSA processing capacity is occupied; when performing non-RSA tasks the main CPU or another resource (such as the random number generator) can be saturated without this statistic reaching 100%. |
| HostIRQs | On PCI HSMs, the total number of interrupts received from the host. On current HSMs, approximately equal to the total of HostReadCount and HostWriteCount . |
| 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. |

| ID | Value |
|---------------|---|
| 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. |
| 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. |

25.3. 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.



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.

26. Hardserver configuration files

The default location of the hardserver configuration file is `/opt/nfast/kmdata/config/config` (**Linux**) or `%NFAST_KMDATA%\config\config` (**Windows**).

The hardserver configuration file has the following sections that you can update to configure the hardserver on an nShield module. If a section is not present, it is assumed to have no entries.

26.1. Hardserver configuration files

Hardserver 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 a `#` character 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 hardserver 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 hardserver behavior. Each section begins with its name in square brackets, as in this example:

```
[slot_imports]
```

You can update the parameters defined in most of these sections to configure the way that the hardserver handles secure transactions between modules connected to the host computer and applications that run on the host computer.



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.

If a particular section is not present in the configuration file, it is assumed to have no entries.

26.2. General hardserver configuration settings

26.2.1. server_settings

The `server_settings` section defines the settings for the client hardserver you can modify while the hardserver is running.



These flags are used by the `NFLOG_DETAIL` environment variable (see [Environment variables to control logging](#)).

The section contains the following fields:

| Field | Description |
|-------------------------------|---|
| <code>loglevel</code> | This field specifies the level of logging performed by the hardserver. See hardserver loglevel and Logging, debugging, and diagnostics . |
| <code>logdetail</code> | 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. |
| <code>connect_retry</code> | This field specifies the number of seconds to wait before retrying a remote connection to a client hardserver. The default is 10. |
| <code>connect_maxqueue</code> | 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 <code>connect_maxqueue</code> to a high value allows high throughput, but may cause long latency if the hardserver goes down. |
| <code>connect_broken</code> | This field specifies the number of seconds of inactivity allowed before a connection to a client hardserver is declared broken. The default is 90. |

| Field | Description |
|------------------------------------|---|
| <code>connect_keepalive</code> | This field specifies the number of seconds between <code>keepalive</code> packets for remote connections to a client hardserver. The default is 10. |
| <code>accept_keepidle</code> | This field specifies the number of seconds before the first <code>keepalive</code> packet for remote incoming connections. The default is 30. Ideally, <code>accept_keepalive</code> should be at least twice the value of the <code>connect_keepalive</code> setting on the unattended machines. |
| <code>accept_keepalive</code> | This field specifies the number of seconds between <code>keepalive</code> packets for remote incoming connections. The socket will be closed after up to ten consecutive probe failures. The default is 10. Ideally, <code>accept_keepalive</code> should be a value such that $(10 * \text{accept_keepalive}) > \text{connect_broken}$ on the unattended machine. Using the default values for both these fields will fulfil this requirement. |
| <code>connect_command_block</code> | When the module has failed, this field specifies the number of seconds the hardserver should wait before failing commands directed to that module with a <code>NetworkError</code> message. For commands to have a chance of succeeding after the module has failed this value should be greater than that of <code>connect_retry</code> . If it is set to 0, commands to a module are failed with <code>NetworkError</code> immediately, as soon as the module. The default is 35. |
| <code>max_pci_if_vers</code> | This field specifies the maximum PCI interface version number. If <code>max_pci_if_vers</code> is set to 0 (the default), there is no limit. |
| <code>enable_remote_mode</code> | <p>If this field is set to <code>yes</code> (the default value) in the module configuration file, nShield HSM mode changing using the <code>nopclearnfail</code> utility is enabled. If set to <code>no</code>, mode changing using <code>nopclearnfail</code> is disabled.</p> <div style="display: flex; align-items: center;">  <div style="border-left: 1px solid black; padding-left: 10px;"> <p>Do not set <code>enable_remote_mode</code> in the client configuration file.</p> </div> </div> |
| <code>enable_remote_reboot</code> | If this field is set to <code>yes</code> (the default value) in the module configuration file, the nShield HSM remote reboot using the <code>nopclearnfail</code> is enabled. If set to <code>no</code> , remote reboot using <code>nopclearnfail</code> is disabled. Run <code>cfg-pushnethsm</code> to push the new config file to the module. |
| <code>enable_remote_upgrade</code> | If this field is set to <code>yes</code> (the default value) in the module configuration file, the nShield HSM remote upgrade using the <code>nopclearnfail</code> is enabled. If set to <code>no</code> , remote upgrade using <code>nopclearnfail</code> is disabled. Run <code>cfg-pushnethsm</code> 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 |
|---------------------------------------|--|
| <code>external_time</code> | This flag specifies the external time (that is, the time according to your machine's local clock) with the log entry. |
| <code>external_date</code> | This flag specifies the external date (that is, the date according to your machine's local clock) with the log entry. |
| <code>external_tid</code> | This flag specifies the external thread ID with the log entry. |
| <code>external_time_t</code> | This flag specifies the external <code>time_t</code> (that is, the time in machine clock ticks rather than local time) with the log entry. |
| <code>stack_backtrace</code> | This flag specifies the stack backtrace with the log entry. |
| <code>stack_file</code> | This flag specifies the stack file with the log entry. |
| <code>stack_line</code> | 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. |
| <code>msg_severity</code> | This flag specifies the message severity (a severity level as used by the <code>NFLOG_SEVERITY</code> environment variable) with the log entry. |
| <code>msg_categories</code> | This flag specifies the message category (a category as used by the <code>NFLOG_CATEGORIES</code> environment variable) with the log entry. |
| <code>msg_writeable</code> | This flag specifies message writeables, and extra information that can be written to the log entry, if any such exist. |
| <code>msg_file</code> | 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. |
| <code>msg_line</code> | 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. |
| <code>options_utc</code> | This flag showing the date and time in UTC (Coordinated Universal Time) instead of local time. |
| <code>unix_file_descriptor_max</code> | 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 <code>unix_file_descriptor_max</code> is set to <code>0</code> (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 ensure that the hardserver is capable of satisfying the maximum number of hardserver connections that applications may make use of. |

26.2.2. 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.

| <code>[server_settings]/loglevel</code> value | Appears in the hardserver log as | Description |
|---|--|--|
| <code>info</code> | Information | Report about the hardserver start-up configuration, connections that have been established or closed. General information for nShield Support for debugging. |
| <code>notice</code> | Notice | Report about certain start-up events, some non-fatal and routine errors that the hardserver can handle internally. |
| <code>client</code> | Detected error in client behaviour | Malformed or invalid messages are received from a client, typically from a local client. |
| <code>remoteserver</code> | Remote server error | Malformed messages or protocol errors are received while communicating with remote peers over the nCipher Secure Transport/Impath protocol. |
| <code>error</code> | Nonfatal error | Unexpected but handled errors from system calls, for example for device or TCP I/O. |
| <code>serious</code> | 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. |
| <code>internal</code> | 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. |
| <code>startup</code> | 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. |
| <code>fatal</code> | 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. |
| <code>fatalinternal</code> | Fatal internal error | A non-recoverable failure occurred, for example certain internal self-consistency checks to detect program logic errors. The hardserver will abort. |

26.2.3. 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 |
|---------------------------------|--|
| <code>enable_scaling</code> | 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 . |
| <code>target_concurrency</code> | 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 <code>target_concurrency</code> 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. |

26.2.4. module_settings

The `module_settings` section defines the settings for the module that can be changed while the hardserver is running.

The section contains the following fields:

| Field | Description |
|-----------------------|---|
| <code>esn</code> | This field specifies the electronic serial number of the module. |
| <code>priority</code> | This field specifies the priority of the module. The value for this field can be an integer from 1 (highest) to 100 (lowest). The default is 100. |

26.2.5. server_remotecomms

The `server_remotecomms` section defines the remote communication settings for the client hardserver. These are read only at hardserver start-up.

This section contains the following fields:

| Field | Description |
|--------------------------|---|
| <code>impath_port</code> | This field specifies the port on which the hardserver listens for incoming <code>impath</code> connections. The default is 9004. Setting this field to <code>0</code> specifies that the hardserver does not listen for incoming connections. Ensure that firewall settings are consistent with port settings. See the Installation Guide for more information about firewall settings. |

26.2.6. `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 hardserver](#).

The section contains the following fields:

| Field | Description (Linux) | Description (Windows) |
|-----------------------------------|---|--|
| <code>unix_socket_name</code> | This field specifies the name of the socket to use for non-privileged connections on Linux. The default is <code>/dev/nfast/nserver</code> . If the <code>NFAST_SERVER</code> environment variable is set, it overrides any value set for <code>unix_socket_name</code> in the hardserver configuration file. For more information about environment variables, see Environment variables . | This field is not used on Windows. |
| <code>unix_privsocket_name</code> | This field specifies the name of the socket to use for privileged connections on Linux. The default is <code>/dev/nfast/privnserver</code> . If the <code>NFAST_PRIVSERVER</code> environment variable is set, it overrides any value set for <code>unix_privsocket_name</code> in the hardserver configuration file. | This field is not used on Windows. |
| <code>nt_pipe_name</code> | This field is not used on Linux. | This field specifies the name of the pipe to use for non-privileged connections on Windows. An empty string specifies none. The default is <code>\\.pipe\crypto</code> . If the <code>NFAST_SERVER</code> environment variable is set, it overrides any value set for <code>nt_pipe_name</code> in the hardserver configuration file. |

| Field | Description (Linux) | Description (Windows) |
|--------------------------------|---|---|
| <code>nt_pipe_users</code> | This field is not used on Linux. | 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, <i>bob</i> not <i>MYDOMAIN\bob</i> or <i>bob@MYDOMAIN</i>). |
| <code>nt_privpipe_name</code> | This field is not used on Linux. | This field specifies the name of the pipe to use for privileged connections on Windows. An empty string specifies none. The default is <code>\\.pipe\privcrypto</code> . If the <code>NFAST_PRIVSERVER</code> environment variable is set, it overrides any value set for <code>nt_privpipe_name</code> in the hardserver configuration file. |
| <code>nt_privpipe_users</code> | This field is not used on Linux. | 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, <i>bob</i> not <i>MYDOMAIN\bob</i> or <i>bob@MYDOMAIN</i>). |
| <code>nonpriv_port</code> | This field specifies the port on which the hardserver listens for local non-privileged TCP connections. The value <code>0</code> (which is the default) specifies none. Java clients default to connecting to port 9000. Ensure that your network firewall settings are correct. See the Installation Guide for more information about firewall settings. If the <code>NFAST_SERVER_PORT</code> environment variable is set, it overrides any value set for <code>nonpriv_port</code> in the hardserver configuration file. | |
| <code>priv_port</code> | This field specifies the port on which the hardserver listens for local privileged TCP connections. The value <code>0</code> (which is the default) specifies none. Java clients default to connecting to port 9001. If the <code>NFAST_SERVER_PRIVPORT</code> environment variable is set, it overrides any value set for <code>priv_port</code> in the hardserver configuration file. | |

26.2.7. load_seemachine

The `load_seemachine` section of the hardserver configuration file defines SEE machines that the module should load and, if required, start for use by other clients. Each SEE machine is defined by the following fields:

| Field | Description |
|------------------------------|--|
| <code>module</code> | This field specifies the module on to which to load the SEE machine. The value must be an integer. A module with this ID must be configured on the client computer. |
| <code>machine_file</code> | This field specifies the file name of the SEE machine. |
| <code>userdata</code> | This field specifies the <code>userdata</code> 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. |
| <code>worldid_pubname</code> | This field specifies the <code>PublishedObject</code> name to use for publishing the <code>KeyID</code> of the started SEE machine. If this field is blank (" "), the <code>KeyID</code> is not published. This field is ignored if the value of the <code>userdata</code> field is blank. |
| <code>postload_prog</code> | 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 <code>-m module#</code> . To run <code>see-sock-serv</code> directly on the nShield HSM, set this field to <code>sockserv</code> . |
| <code>postload_args</code> | This field specifies arguments to pass to the program specified by the <code>postload_prog</code> field. The argument <code>-m module#</code> is automatically passed as the first argument. The <code>postload_args</code> field is ignored if <code>postload_prog</code> is not specified or is blank. To run <code>see-sock-serv</code> directly on the nShield HSM, set this field to <code>`-p `pubname</code> . |
| <code>pull_rfs</code> | This field specifies whether the SEE machine name and userdata should be pulled from the RFS. The default is <code>0</code> : set to <code>1</code> to pull the SEE machine and userdata from the RFS before loading on the remote module. This field will be ignored if set on client machine configurations. <div style="border-left: 1px solid #000; padding-left: 10px; margin-left: 20px;">  <p>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 <code>load_seemachine</code> section of your existing configuration file.</p> </div> |

26.2.8. slot_imports

The `slot_imports` section defines slots from remote modules that will be available to the local computer. Each slot is defined by the following fields:

| Field | Description |
|------------------------|--|
| <code>local_esn</code> | This field specifies the ESN of the local module importing the slot. |

| Field | Description |
|----------------------------|--|
| <code>local_slotid</code> | This field specifies the <code>SlotID</code> to use to refer to the slot when it is imported on the local module. The default is 0, and provides automatic assignment to the lowest available slotID after any configured dynamic slots. |
| <code>remote_ip</code> | This field specifies the IP address of the machine that hosts the slot to import. |
| <code>remote_port</code> | This field specifies the port for connecting to the nShield HSM. |
| <code>remote_esn</code> | This field specifies the ESN of the remote module from which to import the slot. |
| <code>remote_slotid</code> | This field specifies the <code>SlotID</code> of the slot to import on the remote module. The value of this field must be an integer. The default is 0. |

26.2.9. slot_exports

The `slot_exports` section defines the slots on local modules that the local hardserver should allow network modules to import. Each local slot has an entry for each remote module that can import it, consisting of the following fields:

| Field | Description |
|---------------------------|---|
| <code>local_esn</code> | This field specifies the ESN of the local module whose slot can be imported by a network module. |
| <code>local_slotid</code> | This field specifies the <code>SlotID</code> of the slot that is to be imported. The value must be an integer. The default is 0. |
| <code>remote_ip</code> | This field specifies the IP address of the module that is allowed to import the slot. Use <code>0.0.0.0</code> to allow all machines. The default is <code>0.0.0.0</code> |
| <code>remote_esn</code> | This field specifies the ESN of the module allowed to import the slot. Leave the value blank to allow all permitted modules in the Security World. The default is blank. |

26.2.10. dynamic_slots

The `dynamic_slots` section defines the number of Dynamic Slots that each HSM is to support for the Remote Administration Service.

| Field | Description |
|------------------------|---|
| <code>esn</code> | ESN of the HSM to be configured with Dynamic Slots. |
| <code>slotcount</code> | The number of Dynamic Slots that the HSM is to support. If set to 0 (default) the HSM does not support the Remote Administration Service. |

26.2.11. 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 and/or Remote Slot and the local slot becomes the specified slot number. This enables applications and utilities that only support slot 0 to use Remote Administration and Remote Operator.

| Field | Description |
|-------------------|--|
| <code>esn</code> | ESN of the HSM to which the mapping is applied. |
| <code>slot</code> | The slot number to be swapped with slot 0, so that: <ul style="list-style-type: none"> • Slot 0 refers to a Dynamic and/or Remote Slot • The specified slot number refers to the local slot of the HSM. If <code>slot</code> is set to 0 (default) there is no slot mapping. |

26.2.12. 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 |
|--|--|
| <code>round_trip_time_limit</code> | <code>round_trip_time_limit > 5s + network latency time</code> 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. |
| <code>card_remove_detect_time_limit</code> | <code>card_remove_detect_time_limit >= round_trip_time_limit * 2</code> Maximum number of seconds that can pass without a response from the smartcard, 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. |

26.2.13. audit_logging

The `audit_logging` section defines the settings for for the syslog infrastructure used by the audit logging capability. These values require a restart of the hardserver to be recognized.

| Field | Description |
|----------------------------------|---|
| <code>auditlog_port</code> | This field specifies the UDP port to which audit log syslog messages should be delivered. The default is 514. |
| <code>auditlog_addr</code> | This field specifies the IP address of the machine that hosts the syslog server to which audit logging syslog messages should be sent. |
| <code>auditlog_copy_hslog</code> | This field specifies if audit logging syslog entries should be copied into the hardserver log as well as being transmitted to the syslog server. The default is off. It can be turned by setting to yes , true , on or 1 . Care should be taken when setting this field as this can cause the hardserver log to grow significantly. |

26.3. Sections only in client configuration files

26.3.1. nethsm_imports

The `nethsm_imports` section defines the network modules that the client imports. It can also be set up by the `nethsmenroll` utility. Each module is defined by the following fields:

| Field | Description |
|---------------------------|---|
| <code>local_module</code> | This field specifies the ModuleID to assign to the imported module. The value must be an integer. A module with this ID must not be already configured on the client computer. |
| <code>remote_ip</code> | This field specifies the IP address of the module to import. |
| <code>remote_port</code> | This field specifies the port for connecting to the nShield HSM. |
| <code>remote_esn</code> | This field specifies the ESN of the imported module. |
| <code>keyhash</code> | This field specifies the hash of the key that the module should use to authenticate itself. |
| <code>privileged</code> | The value in this field specifies whether the client can make a privileged connection to the module. The default is <code>0</code> , which specifies no privileged connections. Any other value specifies privileged connections. |
| <code>ntoken_esn</code> | 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.

26.3.2. rfs_sync_client

This section defines which remote file system the client should use to synchronize its key

management data:

| Field | Description |
|-----------------------------|--|
| <code>remote_ip</code> | The IP address of the RFS against which to synchronize. |
| <code>remote_port</code> | This field specifies the port for connecting to the RFS. |
| <code>use_kneti</code> | 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. |
| <code>local_esn</code> | This is only required if <code>use_kneti</code> is set to yes . It is the ESN of the local module used for authentication. |
| <code>remote_keyhash</code> | Software or module KNETI hash used to authenticate the RFS, or 40 zeroes to indicate no authentication required (default is 40 zeroes). |
| <code>remote_esn</code> | ESN of the remote module used to authenticate the RFS, or empty when using software KNETI authentication or no authentication required (default is empty). |

26.3.3. 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 modules allowed to access the file system on this client. Each module is defined by an entry consisting of the following fields:

| Field | Description |
|--------------------------|---|
| <code>remote_ip</code> | This field specifies the IP address of the remote module that is allowed to access the file system on this client. |
| <code>remote_esn</code> | This field specifies the ESN of the remote module allowed to access the file system on this client. |
| <code>keyhash</code> | This field specifies the hash of the key with which the client must authenticate itself to the module. The default is 40 zeros, which means that no key authentication is required. |
| <code>native_path</code> | This field specifies the local file name for the volume to which this entry corresponds. |
| <code>volume</code> | This field specifies the volume that the remote host would access to use this entry. |
| <code>allow_read</code> | 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 . |

| Field | Description |
|---------------------------|--|
| <code>allow_write</code> | If this field is set to yes , it means that a remote server is allowed to write to the file. The default is no . |
| <code>allow_list</code> | 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 . |
| <code>is_directory</code> | If this field is set to yes , it means that this entry represents a directory. The default is no . |
| <code>is_text</code> | If this field is set to yes , it means that line endings should be converted to and from the Linux convention for transfers. |



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.

26.3.4. `remote_administration_service_slot_server_startup`

The `remote_administration_service_slot_server_startup` section defines the communication settings that are applied at start-up to the Remote Administration Service.

| Field | Description |
|-------------------|---|
| <code>port</code> | Which port to use to connect to the Dynamic Slot Server. The default is 9005. |

27. Cryptographic algorithms

27.1. 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 |

¹ Not FIPS 140 approved for encryption operations, but available for decryption operations.

27.2. 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 DHE _x | 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 | Y |
| Ed448 | N | N | Ed448 | N |
| ElGamal | Y | Y | DH | Y |
| KCDSA | N | N | KCDSA | N |
| RSA | Y | Y | RSA | Y |
| X25519 | N | N | X25519 | Y |

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.

27.3. 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.

27.4. 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.

28. Audit Logging

28.1. Aims of audit logging

Audit Logging on nShield HSMs provides the means to log administrative operations and key usage events across your estate of HSMs.

Some events are always logged and some events are logged only if certain conditions are met. Details of which events are logged, the conditions for logging the event, and the main information in those logs is described in [Commands audited](#).

Logs are accurately time-stamped to allow correlation of these logs with other systems that may form part of your security or network environments.

Logs are signed and can be verified by the tools provided so that any attempts to alter the contents of the logs can be detected.

Audit logging is a compliance requirement of some certification schemes such as Common Criteria but can optionally be used with any Security World. The option must be specified when the World is created and cannot be added later.

28.1.1. Audit logging and CEF audit logging

There are two different audit log formats which differ depending on the firmware loaded on the HSM.

HSMs loaded with firmware versions prior to 13.5 will produce audit logs in CEF format and this guide refers to the production of logs in that format as 'CEF audit logging'.

HSMs loaded with firmware versions of 13.5 or later will produce audit logs in a new format and this guide refers to the production of logs in this format as 'audit logging'.

For a description of CEF audit logging and how to manage it please refer to the 13.4 version of the User Guide. Older versions of the User Guides are available on the Entrust website. Since CEF was the only format of audit logs available prior to the release of Security World 13.5 older versions of this User Guide refer to CEF audit logging simply as 'audit logging'.

If your Security World uses a mix of HSMs with firmware before and after 13.5 you will receive logs in both formats. These logs are essentially independent and you must manage them separately.

28.2. Configuring audit logging

28.2.1. Enabling audit logging

Audit Logging is enabled on an HSM when it is added to a Security World that is configured for audit logging.

Since accurate time-stamps are important for audit logging Entrust recommends that you ensure that the system clock has been accurately set before enabling audit logging. See, [Manage the system clock of an nShield 5s](#).

The Security World is configured for audit logging when it is created. This can be done in one of two ways:

- Specifying the `--audit-logging` or `-G` option in the `new-world` command
- Specifying the `--mode=common-criteria-cmts` option in the `new-world` command

For the overall procedure, see [Creating a Security World using new-world](#).



The `--mode=common-criteria-cmts` option will create a Security World supporting Common Criteria PP 419 221-5 which imposes other restrictions not related to audit logging. Only use this option if you require compliance to Common Criteria

There are some differences in the conditions that must be met for particular events to be audited depending on which of the two options above has been used to create the Security World. These differences are explained in [Commands audited](#).

28.2.2. Disabling audit logging

Audit Logging is set for the lifetime of the Security World.

To disable Audit Logging on an HSM:

1. Reinitialize the HSM using `initunit`. See [Erasing a module with initunit](#)



Before removing the HSM from the Security World you should check that you have exported and verified all the logs that you require. Once the HSM has been erased it may not be possible to obtain and verify any logs that had not yet been exported.

28.2.3. Key usage logging

By default Audit Logging does not log usage of keys for cryptographic operations such as sign, verify, encrypt and decrypt or their usage in channels for these purposes. The capability to log these operations 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 Developer 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
```

The table in [Commands audited](#) shows which commands implement conditional key usage logging.

28.3. Commands audited

The table below shows which nCore commands may produce audit logs. The column titled 'Logged?' specifies if the logging is conditional or not. An entry in that column of 'ALWAYS' means that the command is unconditionally logged. An entry in that column of 'CONDITIONAL' means that logging is conditional according to the criteria below:

Conditional items:

(a) means the command is logged if an ACL check on a supplied key involves the **LogKeyUsage** bit, see [Key usage logging](#)

(b) means the command may need NSO permissions or (for files/shares with their own ACLs) some other certifier. If the command is certified using a **CertType_SigningKey** entry, a log record may be generated as a result of an ACL check (for **UseAsCertificate**

permission) on that key.

(c) means the command is logged if the object involved is "loggable". Note that keys are always loggable if the Security World was created with `--mode=common-criteria-cmts`, see [Enabling audit logging](#)

(d) means this command becomes 'Always Logged' if the Security World was created with `--mode=common-criteria-cmts`, see [Enabling audit logging](#)

(e) means this command is logged if it accesses a (local or remote) smartcard

Commands marked "(a),(c)" will therefore generate audit-log records if either one of the "input" keys requires logging as part of an ACL check, or one of the "output" keys is a loggable object.

The other two columns in the table refer to the information that is logged for that command. See [Audit log contents](#) for more information.

| Command | Logged? | CmdAuditInfo contents | Extra AuditInfo entries |
|--------------------------------------|--------------------------|--|----------------------------------|
| Cmd_ChangeShareGroup PIN | ALWAYS | SlotID slot table: ShareInfo shares | ObjectUse(cert) TokenID |
| Cmd_ChangeSharePIN | ALWAYS | SlotID slot ShareInfo share | ObjectUse(cert) TokenID |
| Cmd_ChannelOpen | CONDITIONAL (a) | ChannelMode mode | ObjectUse(key,cert) |
| Cmd_CheckUserAction | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_CreateSEEConection | ALWAYS | UUID containerid | - |
| Cmd_Decrypt | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_DeriveKey | CONDITIONAL (a),(c),(d) | - | ObjectUse(key,cert) ObjectNew |
| Cmd_Destroy | CONDITIONAL (c) | AuditObjectID objid Word refcount | - |
| Cmd_Duplicate | CONDITIONAL (a),(c), (d) | - | ObjectUse(key,cert) ObjectNew |
| Cmd_DynamicSlotCreate Association | ALWAYS | SlotID slot | - |
| Cmd_DynamicSlotsConfig ure | ALWAYS | Word slot | - |
| Cmd_Encrypt | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_EraseFile | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |

| Command | Logged? | CmdAuditInfo contents | Extra AuditInfo entries |
|--------------------------|-------------------------|---|-------------------------------------|
| Cmd_EraseShare | ALWAYS | SlotID slot Word i ShortHash hkt | ObjectUse(cert) TokenID |
| Cmd_Export | CONDITIONAL (a),(d) | - | ObjectUse(key,cert) |
| Cmd_FileCopy | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |
| Cmd_FileCreate | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |
| Cmd_FileErase | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |
| Cmd_FileOp | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |
| Cmd_FormatToken | ALWAYS | SlotID slot optional: KeyHashEx auth_key | ObjectUse(cert) TokenID |
| Cmd_GenerateKey | CONDITIONAL (c),(d) | - | ObjectUse(cert) ObjectNew |
| Cmd_GenerateKeyPair | CONDITIONAL (c),(d) | - | ObjectUse(cert) ObjectNew (x 2) |
| Cmd_GenerateLogicalToken | ALWAYS | - | ObjectUse(key,cert) ObjectNew |
| Cmd_GetACL | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_GetAppData | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_GetRTC | CONDITIONAL (b) | - | ObjectUse(cert) |
| Cmd_GetTicket | CONDITIONAL (c) | AuditObjectID objid | - |
| Cmd_ImpathKXBegin | ALWAYS | optional: ASCIIString esn | ObjectNew |
| Cmd_ImpathKXFinish | ALWAYS | - | ObjectUse(imp) |
| Cmd_Import | CONDITIONAL (c),(d) | - | ObjectNew |
| Cmd_InitialiseUnit | ALWAYS | =InitialiseUnitEx | - |
| Cmd_InitialiseUnitEx | ALWAYS | InitModeFlags initmodeflags KMLType kmltype | Startup |
| Cmd_InsertSoftToken | ALWAYS | SlotID slot | - |
| Cmd_LoadBlob | CONDITIONAL (a),(c),(d) | - | ObjectUse(lt,key,cert) ObjectNew |
| Cmd_LoadLogicalToken | ALWAYS | - | ObjectUse(lt,cert) ObjectNew |
| Cmd_MakeBlob | CONDITIONAL (a) | - | ObjectUse(key,cert) |

| Command | Logged? | CmdAuditInfo contents | Extra AuditInfo entries |
|-------------------------|---------------------|--|-----------------------------|
| Cmd_NVMemAlloc | CONDITIONAL (b) | - | ObjectUse(cert) |
| Cmd_NVMemFree | CONDITIONAL (b) | - | ObjectUse(cert) |
| Cmd_NVMemOp | CONDITIONAL (b) | - | ObjectUse(cert) |
| Cmd_ReadFile | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |
| Cmd_ReadShare | ALWAYS | SlotID slot Word i Word sharesleft | ObjectUse(cert,It) TokenID |
| Cmd_ReceiveKey | ALWAYS | - | ObjectNew ObjectUse(imp) |
| Cmd_ReceiveShare | ALWAYS | Word i Word sharesleft | ObjectUse(imp,It) |
| Cmd_RedeemTicket | CONDITIONAL (c) | AuditObjectID objid Word refcount | - |
| Cmd_RemoveKM | ALWAYS | =RemoveKMEx | ObjectUse(cert) |
| Cmd_RemoveKMEx | ALWAYS | KeyHashEx hkm | ObjectUse(cert) |
| Cmd_RemoveSoftToken | ALWAYS | SlotID slot | - |
| Cmd_SendKey | ALWAYS | - | ObjectUse(imp,key,cert) |
| Cmd_SendShare | ALWAYS | SlotId slot Word i KeyHashEx hkm TokenHash hkt | ObjectUse(imp,cert) TokenID |
| Cmd_SetACL | CONDITIONAL (a),(d) | - | ObjectUse(key,cert) |
| Cmd_SetAppData | CONDITIONAL (a),(d) | - | ObjectUse(key,cert) |
| Cmd_SetKM | CONDITIONAL (a),(d) | - | ObjectUse(key,cert) |
| Cmd_SetNSOPerms | ALWAYS | =SetNSOPermsEx | - |
| Cmd_SetNSOPermsEx | ALWAYS | NSOPermsModeFlags nsflags KeyHashEx hknso NSOPerms publicperms | - |
| Cmd_SetRTC | ALWAYS | RTCTime oldtime RTCTime newtime | - |
| Cmd_Sign | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_SignModuleState | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_StartAuditLogging | ALWAYS | ByteBlock nonce | Startup |
| Cmd_StaticFeatureEnable | ALWAYS | FeatureInfo info | - |
| Cmd_StopAuditLogging | ALWAYS | - | Shutdown |

| Command | Logged? | CmdAuditInfo contents | Extra AuditInfo entries |
|-------------------------------------|---------------------|-----------------------|-------------------------------|
| Cmd_Verify | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_WriteFile | CONDITIONAL (b),(e) | - | ObjectUse(cert) TokenID |
| Cmd_WriteShare | ALWAYS | SlotID slot Word i | ObjectUse(lt,cert) TokenID |
| Cmd_CreateBuffer (Solo-XC only) | CONDITIONAL (a) | - | ObjectUse(key,cert) |
| Cmd_CreateSEWorld (Solo-XC only) | ALWAYS | - | - |
| Cmd_ForeignTokenOpen (Solo-XC only) | ALWAYS | SlotID slot | - |
| Cmd_SetSEEMachine (Solo-XC only) | ALWAYS | - | - |

28.4. Audit log contents

Audit log information is recorded in audit records. Audit records are sent from the HSM within a data structure called an audit segment. Each audit segment starts with an [AuditSegmentHeader](#) and ends with a signature block signed by a key [KML](#). See [Audit log verification](#) for further information on this signature and how to verify it.

The number of audit records contained within an audit segment is variable and depends upon factors such as the processing load of the HSM and the number of auditable events being triggered.



Audit logs exported by the system are not in human readable format unless converted using a suitable tool such as [ncore-audit-verify.py](#). Examples shown below have been converted to JSON format.

28.4.1. Audit segment header contents

The [AuditSegmentHeader](#) contains the following information:

- The [ESN](#) of the HSM that produced the audit segment
- The [logID](#) (KML hash) of the HSM that produced the audit segment
- The [runID](#). This changes each time the system boots
- The start and end audit indices
- A timestamp of when the audit segment was started in milliseconds since January 1

1970

- A timestamp of when the audit segment was signed in milliseconds since January 1 1970
- A hash of the audit logs contained within the segment

An example `AuditSegmentHeader` is shown below:

```
"header": {
  "v": 0,
  "flags": "0x0",
  "esn": "14BD-B089-E078",
  "logID": "33133940 fb686dc5 4184a726 b670b7c1 508a8ea1",
  "runID": 4,
  "startIndex": 3458,
  "endIndex": 3458,
  "starttime": 1690801099591,
  "signtime": 1690801099591,
  "datahash": {
    "mech": "SHA256Hash",
    "data": {
      "h": "2dcd7912 b9f0edd3 82df65f7 9ccdf289 4470a569 078638bf 20158eb1 e4c6aa6f"
    }
  }
}
```

28.4.1.1. Audit index

Each audit log record has an audit log index which increments by one for each log created since the start of audit logging. The `AuditSegmentHeader` shows the index for the first and last audit log contained in the segment.

In normal operation no audit records are lost and there will be no gaps in the audit index sequence.

If an abnormal shutdown occurs there is a possibility that some records may have been lost. This will be indicated by a gap in the audit index sequence and also by an `AbnormalShutdown` entry associated with the `StartAuditLogging` command.



It is not possible to tell exactly how many audit records were lost due to an abnormal shutdown. The gap in the audit index sequence indicates the maximum number that could have been lost based on the load of the HSM at the time of the abnormal shutdown but the actual number lost could be fewer than this and could even be zero.

28.4.2. Audit log contents

Each audit log produced will contain information as shown in the `CmdAuditInfo contents` column of the table in `Commands audited`. All audit logs contain basic information to

indicate the name of the command that triggered the audit log and a status that indicates the outcome of that command, as in the extract below from an audit log produced by a `GenerateKeyPair` command.

```
{
  "type": "Command",
  "body": {
    "flags": "0x0",
    "cmd": "GenerateKeyPair",
    "info": {},
    "status": "OK"
  }
}
```

If the `CmdAuditInfo contents` column shows other information that will appear in the `info` section.

If the `Extra AuditInfo entries` column of the table in `Commands audited` contains entries these will appear as separate log entries, as in the extract below.

```
{
  "type": "ObjectUse",
  "body": {
    "v": 0,
    "objid": 137,
    "action": {
      "type": "MakeBlob",
      "details": {
        "flags": "0x0"
      }
    }
  }
}
```

28.4.2.1. Object identification

Objects are identified by an `Objid`. A new ID is assigned when the object is created and will appear in the `ObjectNew` record as shown in the example below.

An object identifier is unique whilst the system runs uninterrupted but values may be reused if the system reboots. The combination of `RunID` from the audit segment header and `ObjID` is always unique.

```
{
  "type": "ObjectNew",
  "body": {
    "v": 0,
    "objid": 135,
    "type": "Key",
    "details": {
      "type": "RSAPrivate",
      "hash": {
        "mech": "SHA1Hash",

```

```

    "data": {
      "hash": "bb6b1e71 1cc06123 853af0d6 ff781cb0 b7a4b515"
    }
  }
}
}
}
}

```

The same ID will appear in **ObjectUse** records and other command records relating to the same object allowing you to trace the object through the audit log. If an object is copied the new object will have its own ID so that you can keep track of each copy. When the copy is created you can identify the object being copied by its hash.

28.5. Audit log verification

As described in [Audit log contents](#) audit log records are protected by a signature across the audit segment in which they are delivered. The signature is made using **KML** which is a secret key unique to the HSM and protected by the HSM.

Each time the HSM is re-initialized (e.g. after upgrading firmware and re-loading the Security World) the previous **KML** is destroyed and a new **KML** is created. The hash of **KML** is used as the **logID** for the audit log.

You can print the **KML** hash of the current Security World using **nfkminfo** and looking for **hkml**. If you have more than one HSM connected **nfkminfo** will show the **hkml** for each HSM.

The signature on the logs can be verified using a tool such as **ncore-audit-verify.py**.

Verification tools require access to the public **KML**. Since **KML** changes when the HSM is initialized it is important to have verified all audit log records before initializing the HSM.

Alternatively keep a record of the public **KML** for all audit logs produced by each HSM.

28.6. Audit log administration

28.6.1. Audit logging and firmware upgrade

The audit log should be finalized before upgrading the firmware on an HSM enrolled in an audit logging Security World. To finalize the audit log, see [Disabling audit logging](#)

If you do not finalize the audit log before attempting to follow the procedures at [Upgrading firmware](#) you will receive an error message. You may either finalize the audit log and repeat the command or use the **hsmadmin upgrade** command with the **--force** option to over-ride the warning.



Upgrading the firmware without finalizing the audit log will result in there being no final audit log record. This means it will not be possible to prove completeness of the audit log.

28.6.2. Audit logging and the system clock

It is important that the timestamps in the audit logs are accurate so that events can be correlated across the whole network in which the HSM is operating.

If the system clock is lost, for instance due to the HSM running on battery power for an extended period of time, `ncoreapi` commands that would result in an audit log being generated will be inhibited until the system clock is restored:

See [Managing the system clock](#) for more information on setting the system clock.

29. 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.



For information about generating keys with the `generatekey` utility, see [Generating keys with the command line](#).

29.1. 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 |
|---------------------|---|
| <code>simple</code> | Specifying the <code>simple</code> application type generates an nShield-native key. No special action is taken after the key is generated. |
| <code>custom</code> | <p>Specifying the <code>custom</code> application type generates a key for custom applications that require the key blob to be saved in a separate file.</p> <p>Specifying <code>custom</code> also causes the generation of a certificate request and self-signed certificate. However, we recommend that you specify the <code>simple</code> (instead of <code>custom</code>) application type whenever possible.</p> |
| <code>pkcs11</code> | <p>Specifying the <code>pkcs11</code> 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:</p> <ul style="list-style-type: none"> • <code>DES3</code> • <code>DH</code> • <code>DSA</code> • <code>ECDH</code> • <code>ECDSA</code> • <code>Ed25519</code> • <code>HMACSHA1</code> • <code>RSA</code> • <code>Rijndae1 (AES)</code> • <code>X25519</code> <p>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 <code>--no-verify</code> option.</p> |

| Parameter | Description |
|-----------------------|--|
| <code>embed</code> | <p>Specifying the <code>embed</code> application type generates a PEM-format RSA/DSA key file that points to a key in <code>NFAST_KMDATA</code> so a software application can then use the HSM-protected key.</p> <p>In applications that use Security World software older than v12.60 and would use the legacy OpenSSL CHIL engine with <code>hwcrhk</code>:</p> <ul style="list-style-type: none"> • The <code>plainname</code> specified in the <code>generatekey</code> command is used as the prefix for all 3 generated files (<code>.key</code>, <code>_req</code>, <code>_selfcert</code>) • <code>.key</code> is appended to all 3 files • The <code>embedsavefile</code> specified in the <code>generatekey</code> command is the destination for all 3 files <p>In applications that use v12.60 or later Security World software :</p> <ul style="list-style-type: none"> • The <code>plainname</code> specified in the <code>generatekey</code> command is used as the prefix for only the <code>.key</code> file, the prefix for the <code>_req</code> and <code>_selfcert</code> file is <code>embed<hash></code> • <code>.key</code> is not appended to the <code>_req</code> and <code>_selfcert</code> files • The <code>embedsavefile</code> is the destination only for the <code>.key</code> file, <code>_req</code> and <code>_selfcert</code> are created in the directory from which <code>generatekey</code> was run from |
| <code>kpm</code> | <p>Specifying the <code>kpm</code> application type generates a key for delivery by an nForce Ultra key server. The <code>generatekey</code> utility automatically creates a special ACL entry that permits a <code>kpm</code> to be delivered to an nForce Ultra's enrolled internal hardware security module.</p> |
| <code>seeinteg</code> | <p>Specifying the <code>seeinteg</code> 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.</p> |
| <code>seeconf</code> | <p>Specifying the <code>seeconf</code> 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.</p> |

29.2. Key properties (NAME=VALUE)

The `NAME=VALUE` syntax is used to specify the properties of the key being generated.



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 |
|----------------------------|---|
| <code>alias</code> | The <i>VALUE</i> for <code>alias</code> specifies an alias to assign to the key. |
| <code>blobsavefile</code> | When using the <code>custom</code> application type, the <i>VALUE</i> for <code>blobsavefile</code> specifies a file name of the form <code>FILENAME_req.ext</code> 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 <code>ROOT_inf.txt</code> ; for asymmetric key types, the public key blob is also saved to a file whose name has the form <code>ROOT_pub.EXT</code> . |
| <code>cardset</code> | The <i>VALUE</i> for <code>cardset</code> specifies an OCS that is to protect the key (if <code>protect</code> is set to <code>token</code>). 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 <code>nfkminfo</code>). |
| <code>certreq</code> | <p>Setting <code>certreq</code> 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.</p> <p>To generate a certificate request you must set the <i>VALUE</i> for <code>certreq</code> to <code>yes</code>, which makes <code>generatekey</code> 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 <code>FILENAME_req.ext</code> (where <code>FILENAME</code> is a name of your choice).</p> <p>An extra file with a name of the form <code>FILENAME.ext</code> 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 <code>certreq</code> with the <code>--retarget</code> option to generate a self-signed certificate for an existing key.</p> |
| <code>checks</code> | For RSA key generation only, this specifies the number of checks to be performed. Normally, you should leave <i>VALUE</i> empty to let the module pick an appropriate default. |
| <code>curve</code> | For ECDH and ECDSA key generation only, the <i>VALUE</i> for <code>curve</code> specifies which curves from the supported range to use. Supported curves are: ANSIB163v1, 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 |
| <code>embedconvfile</code> | The <i>VALUE</i> for <code>embedconvfile</code> specifies the name of the PEM file that contains the RSA key to be converted. |

| Option | Description |
|-------------------------------|---|
| <code>embedsavefile</code> | <p>When using the <code>embed</code> application type, the <i>VALUE</i> for <code>embedsavefile</code> 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.</p> <p>A certificate request and a self-signed certificate are also written. If the filename is <code>ROOT.EXT</code> then the request is saved to <code>ROOT_req.EXT</code> and the self-signed certificate is saved to <code>ROOT_selfcert.EXT</code>.</p> |
| <code>from-application</code> | <p>When retargeting a key, the <i>VALUE</i> for <code>from-application</code> specifies the application name of the key to be retargeted. Only applications for which at least one key exists are acceptable.</p> |
| <code>from-ident</code> | <p>When retargeting a key, the <i>VALUE</i> for <code>from-ident</code> specifies the identifier of the key to be retargeted (as displayed by the <code>nfkminfo</code> command-line utility).</p> |
| <code>hexdata</code> | <p>The <i>VALUE</i> for <code>hexdata</code> 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.</p> |
| <code>ident</code> | <p>The <i>VALUE</i> for <code>ident</code> specifies a unique identifier for the key in the Security World. For applications of types <code>simple</code>, 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.</p> <p>The following characters are allowed in key IDs:</p> <ul style="list-style-type: none"> • digits 0-9 • lower-case letters a-z • hyphen (-) |
| <code>keystore</code> | <p>The <i>VALUE</i> for <code>keystore</code> specifies the file name of the key store to use. This must be an nShield key store.</p> |
| <code>keystorepass</code> | <p>The <i>VALUE</i> for <code>keystorepass</code> specifies the password to the key store to use.</p> |
| <code>logkeyusage</code> | <p>The <i>VALUE</i> for <code>logkeyusage</code> 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.</p> |
| <code>module</code> | <p>The <i>VALUE</i> for <code>module</code> 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).</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <div style="border-left: 1px solid #ccc; padding-left: 10px;"> <p>You can also specify a module by setting the <code>--module</code> option.</p> </div> </div> |

| Option | Description |
|-----------------------------|--|
| <code>paramsreadfile</code> | The <i>VALUE</i> for <code>paramsreadfile</code> 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 <i>VALUE</i> for <code>paramsreadfile</code> is not specified, the module uses a default file. |
| <code>pemreadfile</code> | The <i>VALUE</i> for <code>pemreadfile</code> 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. |
| <code>plainname</code> | The <i>VALUE</i> for <code>plainname</code> 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 <code>kmdata</code> (Linux) or <code>%NFAST_KMDATA%</code> (Windows) and not used otherwise. |
| <code>protect</code> | The <i>VALUE</i> for <code>protect</code> specifies the protection method, which can be <code>module</code> for security-world protection, <code>softcard</code> for softcard protection or <code>token</code> for Operator Card Set protection. The default is <code>token</code> , except for <code>seeconf</code> keys, where the default is <code>module</code> . <code>seeinteg</code> keys are always token-protected. The <code>softcard</code> option is only available when your system has at least one softcard present. |
| <code>pubexp</code> | For RSA key generation only, the <i>VALUE</i> for <code>pubexp</code> 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. |
| <code>recovery</code> | The <i>VALUE</i> for <code>recovery</code> enables recovery for this key and is only available for card-set protected keys in a recovery-enabled world. If set to <code>yes</code> , the key is recoverable. If set to <code>no</code> , key is not recoverable. The default is <code>yes</code> . Non-recoverable module-protected keys are not supported. |
| <code>seeintegname</code> | If present, the <i>VALUE</i> for <code>seeintegname</code> identifies a <code>seeinteg</code> key. The ACL of the newly generated private key is modified to require a certificate from the <code>seeinteg</code> key for its main operational permissions, such <code>Decrypt</code> and <code>Sign</code> (<code>DuplicateHandle</code> , <code>ReduceACL</code> , and <code>GetACL</code> are still permitted without certification.) If you use <code>seeintegname</code> to specify a key that has been recovered with the <code>rocs</code> utility, you must also use the <code>-N</code> option with <code>generatekey</code> . |
| <code>selfcert</code> | The <i>VALUE</i> for <code>selfcert</code> 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 <code>selfcert</code> to <code>yes</code> , which makes <code>generatekey</code> 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 <code>FILENAME.ext</code> . You can use this parameter with the <code>--retarget</code> option to generated a self-signed certificate for an existing key. |

| Option | Description |
|----------------------------|--|
| <code>size</code> | For key types with variable-sized keys, the <i>VALUE</i> for <code>size</code> specifies the key size in bits. The range of allowable sizes depends on the key type and whether the <code>--no-verify</code> 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 <code>curve</code> . |
| <code>strict</code> | For DSA key generation only, setting the <i>VALUE</i> for <code>strict</code> to <code>yes</code> enables strict verification, which also limits the size to 2048 or 3072 bits. The default is <code>no</code> . |
| <code>type</code> | The <i>VALUE</i> for <code>type</code> 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 <code>--verify</code> option limits the available key types. |
| <code>x509country</code> | The <i>VALUE</i> for <code>x509country</code> specifies a country code, which must be a valid 2-letter code, for the certificate request. |
| <code>x509dnscommon</code> | The <i>VALUE</i> for <code>x509dnscommon</code> specifies a site domain name, which can be any valid domain name, for the certificate request. |
| <code>x509email</code> | The <i>VALUE</i> for <code>x509email</code> specifies an email address for the certificate request. |
| <code>x509locality</code> | The <i>VALUE</i> for <code>x509locality</code> specifies a city or locality for the certificate request. |
| <code>x509org</code> | The <i>VALUE</i> for <code>x509org</code> specifies an organization for the certificate request. |
| <code>x509orgunit</code> | The <i>VALUE</i> for <code>x509orgunit</code> specifies an organizational unit for the certificate request. |
| <code>x509province</code> | The <i>VALUE</i> for <code>x509province</code> specifies a province for the certificate request. |
| <code>xsize</code> | The <i>VALUE</i> for <code>xsize</code> 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. |

29.3. 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 |
|--------------------|----------|--------|----------|
| <code>alias</code> | 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 | |
| nvrn | 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 |
| nvrnm | 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 | |

30. 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](#).

30.1. Voltage Monitoring for Battery Replacement

All of the voltage rails in the nShield HSM are monitored to protect against potential over- or under-voltage attacks. You can view the most recent measurement of the voltages using the `stattree` command.

These modules also contain a user-replaceable battery. The battery powers security functions on the module when the main module power is removed, for example when the host is turned off, so it is expected that the battery voltage will drop over time as the battery drains. To avoid module downtime due to battery replacement we recommend that the battery voltage is monitored regularly, especially if a module has had its main power removed for considerable time.

`CPUVoltage10` reported by `stattree` under the `ModuleEnvStats` node tag displays the current battery voltage:

```
+PerModule:
+#1:
+ModuleEnvStats:
...
-CPUVoltage10      3.16
...
```

The battery supplied with the nShield HSM has a nominal voltage of 3.0V. In the above example the battery is fully charged and has been measured at 3.16V, which is within the acceptable range of 2.46V - 3.55V. If the battery voltage is measured to be lower than 2.46V, the module will report an `S0S-B1` error. See [HSM status and error codes](#) for more information regarding error reporting.



Contact Support to request information regarding a replacement battery if `stattree` reports the battery voltage to be below 2.70V.

Consult the *Installation Guide* for instructions on replacing the battery in your module.

30.2. Temperature Monitoring for Airflow Validation

Temperatures within a module are monitored to protect against potential attacks, and to prevent overheating. The temperatures of the processors within an nShield 5s are reported as `CurrentCPUTemp1` and `CurrentCPUTemp2` under the `ModuleEnvStats` node tag of `stattree`.



As an nShield 5s has a passively-cooled heatsink, care must be taken to install it in an environment with forced airflow. Refer to the *Installation Guide* for airflow guidance.

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.



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.

| <code>stattree</code> Statistic | Description | Minimum expected in optimum environment | Maximum expected in optimum environment |
|---------------------------------|--|---|---|
| <code>CurrentCPUTemp1</code> | First processor temperature | 10°C | 75°C |
| <code>CurrentCPUTemp2</code> | Second processor temperature | 10°C | 78°C |
| <code>MaxTempC</code> | Maximum temperature measured on either processor | - | 78°C |
| <code>MinTempC</code> | Minimum temperature measured on either processor | 10°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.

31. System upgrade

31.1. Terminology

Within this section the term 'software' is used to mean Security World software running on the PC in which the HSM is installed and the term 'firmware' is used to mean the Security World firmware running on the HSM.

The software and firmware can be upgraded independently.

31.2. Software and firmware compatibility

In general, Entrust recommends that you use the software and firmware from the same version of Security World. The system is designed to be backwards compatible so that it will still operate with differing versions of software and firmware but some functionality may not be available and you may receive warnings during operation.

This user guide describes the behaviour of v13.5 software interacting with v13.5 firmware. Some areas where functionality differs depending on the version of firmware loaded are also described in this guide but it is not possible to describe all possible combinations of software and firmware.

Release notes and user guides for each Security World release are available from the Entrust website and these together with Entrust Support will help you should you experience any problems when operating with differing versions of software and firmware.

31.3. System upgrade procedure

When upgrading the whole system, Entrust recommends that you always upgrade the host software before upgrading the HSM firmware, however this is not mandatory and you may upgrade the firmware first should you wish to do so.



Always read the release notes accompanying the Security World release before upgrading any part of the system as these may include additional upgrade steps.



If the Version Security Number (VSN) of the firmware has been increased, it may not be possible to roll-back the firmware to the previous version after upgrade. See [Version Security Number](#) for more information.

31.3.1. Software upgrade procedure

For Security World software upgrades, you do not need to delete key data or any existing Security World. If you do delete Security World data, it cannot be restored unless you have an up-to-date backup and a quorum of the Administrator Card Set (ACS) available.

31.3.1.1. Before upgrading software

You must perform these steps if you are planning to re-install the Security World software, for example to re-install it on the same machine after an operating system update, or to install a newer Security World software version as part of an upgrade.

Performing these steps is useful even if you are not planning a re-install because it preserves data that you would otherwise irretrievably lose when you uninstall the Security World software.

1. For Linux installations make a backup of your Security World and nShield configuration files stored in `/opt/nfast/kmdata/` and `/opt/nfast/hardserver.d` by copying them to external media or to a location not within `/opt/nfast`.

When you are upgrading the Security World, you will also restore the backup to preserve your PKCS #11 and Soft KNETI authentication settings and any customizations. If you delete the `/opt/nfast` or `$NFAST_HOME` directory without making a copy of it, you will lose these configuration settings. When you are restoring a Security World from a backup, you will need to maintain permissions.

2. Back up your SSH keys, see [Making a backup of installed SSH keys](#).
 - If you are planning a clean reinstallation of the Security World software on the **same machine and same operating system**, back up your SSH keys in `/opt/nfast/services` using `hsmadmin keys backup`.
 - If you are planning to re-create the Security World on a **different machine or after re-installing the operating system**, use `hsmadmin keys backup --passphrase`. `hsmadmin keys backup` alone is only suitable for a local backup followed by a local restore on the same machine and same operating system.



If you erase your SSH keys without making a backup you will need to use recovery mode, see [Recovery mode](#) to restore communication with the HSM. This will return the HSM to factory state, see [Factory state](#).

31.3.1.2. Upgrading software

Software upgrade is performed by uninstalling the old software as described in [Uninstalling Security World software](#) and then installing the new software as described in [Installing the software](#) in the *nShield 5s User Guide*.

31.3.1.3. After upgrading software

1. Copy back any data that was manually backed-up as part of the procedures in [Before upgrading software](#) to the locations from which it was copied.
2. Restore communication with the HSM by following the procedures at [restoring SSH keys from backup](#).

31.3.2. Firmware upgrade procedure

See [Upgrading firmware](#)

32. Uninstalling Security World software

This section describes how to uninstall Security World software.

Entrust recommends that you only uninstall the Security World software in the following circumstances:

- You are certain it is no longer required.
- You intend to upgrade it.

In Windows environments, because the **hardserver** is installed as a named service (known as the nFast server), it is only possible to have one Security World software installation on any given computer.

In Linux environments, it is also not possible to have more than one Security World software installation on the same computer.

Entrust therefore recommends that, when you are upgrading, any older versions of software are uninstalled and deleted.

The automated Security World software installers do not delete user created components, key data, or Security World data. However:

- On Linux, a manual installation using **.tar** files **does** overwrite existing data and directories.
- On Windows, if the installer detects an existing Security World Software installation, it asks you if you want to install the new components. These components replace your existing installation.



Make sure you have completed the steps described in [before upgrading software](#) before you uninstall any software.

32.1. Uninstalling the Security World software on Windows

This procedure requires the privileges of the built-in local Administrators group.

1. Open the **Control Panel** and select **Programs and Features**.
2. For the following programs, select **Uninstall** and follow the on-screen instructions:
 - **nShield Software**.
 - **CyberJack Base Components**.
3. If the file **nCipherKM.jar** is present, it is located in the extensions folder of your local Java Virtual Machine. The uninstall process may not delete this file and you have to

remove it manually.

32.1.1. Removal of Security World data

The uninstaller removes only those files that were created during the installation.

If you wish to completely uninstall the Security World, including any data created during its operation, including key data, navigate to the installation directory and delete the files in the `%NFAST_KMDATA%` folder.



Do not remove Security World data as part of an upgrade procedure. Only remove Security World data if you wish to do a complete uninstallation.

32.2. Uninstalling the Security World software on Linux

This procedure requires `root` privileges.

1. To remove drivers, install fragments, and scripts and to stop services, run the command:

```
/opt/nfast/sbin/install -u
```

2. If you have installed packages with `rpm`, you must also remove them with `rpm`:
 - a. Check which `rpm` packages are installed on the system:

```
rpm -qa nShield*
```

Example output:

```
For example, the output of this command on a test system looks as so:
nShield-hwsp-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-csdref-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-csd-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-ctd-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-ctls-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-devref-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-javasp-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-jd-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-ncsnmp-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
nShield-raserv-237.a1f2123456789ffffffffd1d0d996f12272a98928-13.4.2.x86_64
```

- b. Follow the package list from the output **bottom up** to uninstall the packages one by one:

```
rpm -e <full-package-name>
```

For example:

```
rpm -e nShield-hwsp-237.a1f2123456789fffffffffd1d0d996f12272a98928-13.4.2.x86_64
```

From the output above, uninstall **raserv** first and uninstall **hwsp** last.



Files that were created after an **rpm** was installed are not deleted.

3. Delete all the files, including those in subdirectories, from **/opt/nfast** and **/dev/nfast/**:



The following commands delete both software and data. If you are uninstalling as part of an upgrade ensure that you have performed all the steps in [Before upgrading software](#) to ensure that the critical data is backed-up so that it can be restored after the upgrade.

```
rm -rf /opt/nfast
rm -rf /dev/nfast
```

4. If you are not re-installing the product, delete the configuration file **/etc/nfast.conf** if it exists.
5. Unless they are needed for a new installation, remove the user and group **nfast** and, if they exist, the user and group **ncsnmpd**.

```
sudo userdel ncsnmp
sudo userdel nfast
sudo groupdel ncsnmp
sudo groupdel nfast
```

6. If the file **nCipherKM.jar** is present, it is located in the extensions folder of your local Java Virtual Machine. The uninstall process may not delete this file and you may have to remove it manually.

33. Upgrading firmware

This section describes how to upgrade firmware on your nShield HSM hardware security module.

33.1. Primary, recovery and bootloader firmware

HSM firmware consists of three major components:

- Primary image firmware
- Recovery image firmware
- Bootloader firmware

Upgrade packages may contain updates for any of these components. The same upgrade method is used in all cases. The system will automatically detect which components are included in the update package and will load the firmware to the correct location.

If upgrade packages are available for both Primary and Recovery firmware it is not recommended to upgrade them both at the same time. The recommended procedure is to always upgrade the Primary firmware first. Test that the system performs as expected and then upgrade the Recovery firmware at a later date.

33.2. Firmware version control

The version of Primary and Recovery image firmware that can be installed on an HSM is controlled by the Version Security Number, see [Version Security Number](#).

Bootloader firmware version control is described in [Bootloader version](#).

33.2.1. Version Security Number (VSN)

Entrust supply several versions of the module firmware. Primary and Recovery image firmware includes a Version Security Number (VSN). This number is increased whenever Entrust improve the security of the firmware. Ensuring you use firmware with the highest available VSN allows you to benefit from these security improvements.

However, if you have a regulatory requirement to use certified firmware such as that approved by FIPS or Common Criteria, you should only install the latest available firmware that has been certified by the relevant certification authority. This firmware may not have the highest VSN available.

Every HSM records the minimum firmware VSN that it will accept. You can always upgrade to firmware with an equal or higher VSN than the minimum VSN set on your module, even if the firmware currently installed on the module has a higher VSN than the firmware to which you are upgrading.

You can upgrade to a firmware version with a higher VSN than the HSM's current firmware, without committing yourself to the upgrade, by installing the newer firmware without altering the HSM's minimum VSN requirement. The older firmware can be reinstalled at any time provided the HSM's minimum VSN has not been altered.



You can never load firmware with a lower VSN than the target HSM's minimum VSN requirement. For example, if the HSM has a minimum VSN requirement of 3 and the currently installed firmware has a VSN of 4, you can install firmware with a VSN of 3 or above to the HSM. You cannot install firmware with a VSN of 1 or 2 to this HSM.

33.2.1.1. Configuring the minimum VSN

To increase the HSM's minimum VSN requirement, use the command `hsmadmin setminvsn`. The new VSN must be greater than or equal to the HSM's current minimum required VSN, and cannot be greater than the VSN of the firmware currently installed on the HSM.

It is recommended that the `hsmadmin setminvsn` command always be used as soon as the decision has been made not to return to the older version of the firmware. This prevents future downgrades of the firmware that could potentially weaken security.

33.2.2. Bootloader version

Bootloader firmware does not have a VSN. The Bootloader version number is included in the filename of the upgrade package. Entrust recommend that you always install the latest version available.

For security reasons some Bootloader firmware upgrades are irreversible. These Bootloader upgrades revoke the signing key used to sign previous Bootloader firmware and thus it is not possible to revert back to previous Bootloader firmware after such an upgrade.

Refer to the release notes accompanying the firmware release to identify whether the Bootloader upgrade is reversible or not.



It is only possible to update the Bootloader when running firmware version 13.4 or later.

33.3. Firmware on the installation media

Your Firmware installation media may contain several sets of firmware for each supplied product. These can include the:

- latest FIPS approved firmware
- latest Common Criteria approved firmware
- latest firmware available

You should ensure you are using the latest firmware available, unless you have a regulatory requirement to use firmware that has been certified by a specific certification authority.

33.3.1. Recognising firmware files

The firmware files are stored in subdirectories within the `firmware` directory on the installation media. The subdirectories are named by product and then certification status, which can be `latest`, `fips-pending`, `fips`, or `cc`.

Firmware files for nShield HSM modules have a `.npkg` filename suffix.

The VSN of a Primary or Recovery image firmware file is incorporated into its filename and is denoted by a dash and the letters "vsn" followed by the digits of the VSN. For example, `-vsn24` means the VSN is 24.

To display information about a firmware file on the installation media, enter the following command:

Linux

```
hsmadmin npkginfo /disc-name/firmware/nShield5s/status/firmware_file.npkg
```

In this command, `disc-name` is the directory on which you mounted the installation media, `status` is the certification status, and `firmware_file` is the file name.

Windows

```
hsmadmin npkginfo E:\firmware\nShield5s\status\firmware_file.npkg
```

In this command, `E` is the drive letter of your installation media, `status` is the certification status, and `firmware_file` is the file name.

33.4. Firmware installation overview

Normal procedure is to install firmware when the HSM is running in primary mode. If the HSM is running in recovery mode, as described in [Recovery Mode](#) the procedure is identical except that the reboot caused by `hsmadmin upgrade` will cause the module to factory state and it will be necessary to run `hsmadmin enroll` before continuing with the rest of the installation.



If you are upgrading a module which has SEE program data or NVRAM-stored keys in its nonvolatile memory, use the `nvr-ram-backup` utility to backup your data first.



If the HSM to be upgraded is part of an audit logging Security World you will need to finalize the audit log before starting the upgrade. See [audit logging and firmware upgrade](#) for information on how to do this.



You should always check that the system clock is correct before upgrading the firmware and adjust it if necessary, see [Managing the system clock](#).

1. Put the module in Maintenance mode.

See [Check and change the mode of operation](#).

2. Check the version of the firmware currently loaded, see [hsmadmin status](#).
3. View information about the firmware in the upgrade file including the version and the VSN, see [hsmadmin npkginfo](#).
4. Optionally make a dry-run of the upgrade by using `hsmadmin upgrade` with the `--dry-run` option. This will check that everything is in place for the upgrade to succeed but will not upgrade the firmware. If any errors are reported fix these before continuing to the next step.
5. Upgrade the firmware, see [hsmadmin upgrade](#).

The current firmware version and the firmware version being loaded will be displayed automatically.

The module will be programmed with the new firmware and will be automatically rebooted.



If the installation is being run from recovery mode this reboot will factory state the HSM and `hsmadmin enroll` must be run before continuing.



The module will report which internal components of the firmware have been updated. These components are pre-determined by the

individual upgrade file and the internal names are intended for use by Entrust support staff only.

6. Check the version of the firmware now loaded, see [hsmadmin status](#).
7. Put the module in initialization mode.

See [Check and change the mode of operation](#).

8. Restore the HSM to the Security World, see [Adding or restoring an HSM to the Security World](#)



If the HSM is not part of a Security World you can use the command `initunit` instead of this step.

9. Put the module in Operational mode.

See [Check and change the mode of operation](#).

10. Run the `enquiry` command to verify the module is in operational state and has the correct firmware version.

In Operational mode, the `enquiry` command shows the version number of the firmware loaded. This is the `version` field listed per module.

34. SNMP monitoring agent

This appendix describes the nShield SNMP monitoring agent that you can add to your SNMP manager application. It is based on the open-source Net-SNMP project, version 5.7.3.

34.1. Install and activate the nShield SNMP monitoring 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.

34.1.1. 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**).

34.1.2. 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 in `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.

34.1.3. 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/nesnmpd 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"
```

34.2. Basic configuration

34.2.1. 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](#).

34.2.2. 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?](#)



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**).

34.2.3. Create the configuration files (Windows)

On Windows, the `snmp.conf` and `snmpd.conf` files are not created automatically by the installation. 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.



The sample `snmpd.conf` file includes `agentuser` and `agentgroup` directives, however these are inoperative in Windows.



You can override the default search path by setting the environment variable `SNMPCONFPATH` to a colon-separated (":") list of directories for which to search.

34.2.3.1. 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.reloadConfig.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.

34.2.3.2. 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` configuration file is not required for the agent to operate and report MIB entries.

34.2.3.3. 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 |
|---------------------------------------|---|
| <code>statstimeout</code> | This directive specifies the length of time for which statistics commands are cached. The default is 5 seconds. |
| <code>admintimeout</code> | This directive specifies the length of time for which administrative commands are cached. The default is 60 seconds. |
| <code>keytable</code> | This directive sets the initial state of the key table to <code>none</code> , <code>all</code> , or <code>query</code> . See ListKeys in Administration sub-tree overview . |
| <code>enable_trap_zero_suffix</code> | This directive appends the '.0' suffix to object identifiers (OIDs) for backward compatibility. The default is <code>0</code> (disabled); the directive can be set to <code>1</code> to restore the suffix. Valid values are 0 and 1. |
| <code>memoryUsageOkThreshold</code> | 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. |
| <code>memoryUsageHighThreshold</code> | 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. |



There may be a tolerance gap between the `memoryUsageOkThreshold` and the `memoryUsageHighThreshold` values.



The timeouts should be set to values that achieve a balance between receiving up to date information whilst preventing excessive load.

34.2.4. 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.



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.

34.2.5. 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.

34.2.6. 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.

34.2.7. 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
```

34.2.8. 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 corresponding 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.

34.3. 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 |
| 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>.



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 [\[n5s-ug:snmp-monitor:::usm-user\]](#). The user can then be referenced in access control directives(s) after which it can be used.

34.4. 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 [noauth | 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 user1
```

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 **GETNEXT**) 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:

```
rocommunity public
```

- Setting up a read/write community named `private` that can only be accessed from the machine on which the agent is running:

```
rocommunity 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.

34.5. VACM configuration

The full flexibility of the VACM, for example allowing access to two or more distinct OID subtrees, or different access requirements for reading and writing, is available using four configuration 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 (`com2sec` 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 community 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

security 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 community-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 subtree, 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 an 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. 111111110100000) 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 | | | | | | | |

34.6. Trap Configuration

The distribution of SNMP supports SNMPv1, SNMPv2 and SNMPv3 traps. Control over these traps is defined with a number of directives:

34.6.1. 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 manager 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
```

34.6.2. 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
```

34.7. 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.

34.7.1. 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.



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).

34.7.2. 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**).

34.7.3. 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 |
|-------------------------------|-----|------------|---|
| <code>hardserverFailed</code> | 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. |
| <code>modulesFailed</code> | R | TruthValue | True if any modules have failed. |
| <code>load</code> | R | Unsigned32 | Percentage of total available capacity currently utilized. |

34.7.3.1. 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 |
|--------------------------------|---|
| <code>hardserverAlert</code> | This trap is sent when the hardserver fails or is shut down. |
| <code>hardserverUnAlert</code> | This trap is sent when the hardserver restarts. |
| <code>moduleAlert</code> | This trap is sent when a module fails. |
| <code>moduleUnAlert</code> | This trap is sent when a module is restarted after a failure. |

| Node name | Description |
|-----------------------------------|---|
| <code>psuAlert</code> | This trap is sent when a PSU fails. |
| <code>psuUnAlert</code> | This trap is sent when a previously-failed PSU is working again. |
| <code>fanfailureAlert</code> | This trap is sent when a fan fails. |
| <code>fanfailureUnAlert</code> | This trap is sent when a previously-failed fan is working again. |
| <code>memoryUsageHighAlert</code> | 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. |
| <code>memoryUsageOkAlert</code> | 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.



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 |
|---|--|
| <code>SNMPv2-MIB::coldStart</code> | This trap is sent when the SNMP agent is started |
| <code>NET-SNMP-AGENT-MIB::nsNotifyShutdown</code> | This trap is sent when the SNMP agent is stopped |

34.7.4. 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 `memoryUsageHighThreshold` 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:



34.7.5. Administration sub-tree overview

The administration sub-tree (`enterprises.nCipher.nC-series.administration`) contains information about the permanent state of the hardware 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 |
|--------------------------------|-----|---|---|
| <code>hardserverRunning</code> | R | Enum 1: Running 2: NotRunning | This variable reflects the current state of the hardserver (<code>Running</code> or <code>NotRunning</code>). |
| <code>noOfModules</code> | R | Gauge32 | Number of nC-series modules. |
| <code>hsVersion</code> | R | DisplayString | Hardserver version string. |
| <code>globalSpeedIndex</code> | R | Gauge32 | Number of 1024-bit signatures each second. |
| <code>globalminQ</code> | R | Gauge32 | Minimum recommended queue. |
| <code>globalmaxQ</code> | R | Gauge32 | Maximum recommended queue. |
| <code>SecurityWorld</code> | R | TruthValue | <code>True</code> if a Security World is installed and operational. |
| <code>swState</code> | R | DisplayString | Security World display flags, as reported by <code>nfkminfo</code> . |
| <code>listKeys</code> | 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. |
| <code>serverFlags</code> | R | DisplayString | Supported hardserver facilities (the <code>NewEnquiry</code> level 4 flags). |
| <code>remoteServerPort</code> | R | Gauge32 | TCP port on which the hardserver is listening. |
| <code>swGenTime</code> | R | DisplayString | Security World's generation time. |
| <code>swGeneratingESN</code> | 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`).

34.7.5.1. 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. |
| 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. |

34.7.5.2. 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. |

| Node name | R/W | Type | Remarks |
|-----------------|-----|---------|---|
| 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 nonvolatile 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. |
| adminQuorumFT0 | R | Gauge32 | The quorum required to use a Foreign Token Open Delegate Key. |

34.7.5.3. Module administration table

The following table gives details of the nodes in the module administration table (`enterprises.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. |

| Node name | R/W | Type | Remarks |
|---------------------|-----|---------------|--|
| 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). |
| 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. |

| Node name | R/W | Type | Remarks |
|-------------------------------|-----|--|---|
| <code>moduleSWState</code> | 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. |
| <code>moduleSWFlags</code> | R | DisplayString | Security World flags for this module. |
| <code>hashKML</code> | R | MHash | Hash of the module's secret key. |
| <code>moduleFeatures</code> | R | DisplayString | Features enabled on this module. |
| <code>moduleFlags</code> | R | DisplayString | Like <code>serverFlags</code> , but for each module. |
| <code>versionSerial</code> | R | Gauge32 | Firmware Version Security Number (VSN); see Version Security Number (VSN) . |
| <code>hashKNETI</code> | R | MHash | K_{NETI} hash, if present. |
| <code>longQ</code> | R | Gauge32 | Max. rec. long queue. |
| <code>connectionStatus</code> | R | DisplayString | Connection status (for imported modules). |
| <code>connectionInfo</code> | R | DisplayString | Connection information (for imported modules). |
| <code>machineTypeSEE</code> | R | DisplayString | SEE machine type. |

34.7.5.4. Slot administration table

The following table gives details of the nodes in the slot administration table (`enterprises.nCipher.nC-series.administration.slotAdminTable`):

| Node name | R/W | Type | Remarks |
|-----------------------------------|-----|--|---|
| <code>slotAdminModuleIndex</code> | R | Integer32 | Module number of the module containing the slot. |
| <code>slotAdminSlotIndex</code> | R | Integer32 | Slot number (1-based, unlike <code>nCore</code> which is 0-based). |
| <code>slotType</code> | R | Integer 1: Datakey 2: Smart card 3: Emulated 4: Soft token 5: Unconnected 6: Out of range 7: Unknown | Slot type. |
| <code>slotFlags</code> | R | DisplayString | Flags referring to the contents of the slot (from <code>slotinfo</code>). |
| <code>slotState</code> | 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. |
| <code>slotListFlags</code> | R | DisplayString | Flags referring to attributes of the slot (from <code>getslotlist</code>). |

| Node name | R/W | Type | Remarks |
|-------------------|-----|---------------|---|
| slotShareNumber | R | Gauge32 | Share number of card currently in slot, if present. |
| slotSharesPresent | R | DisplayString | Names of shares present in card currently in slot. |

34.7.5.5. 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. |
| cardsetTimeout | R | Gauge32 | Token time-out period, in seconds, or 0 if none. |
| cardsetGenTime | R | DisplayString | Generation time of card set. |

34.7.5.6. 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 (`enterprises.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. |

| Node name | R/W | Type | Remarks |
|--------------------|-----|---|--|
| 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. |
| 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. |
| keySEInteg | 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. |

34.7.5.7. 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 |
|------------------------------------|-----|---|--|
| <code>keyQueryAppname</code> | R/W | DisplayString | Application name. |
| <code>keyQueryIdent</code> | R/W | DisplayString | Name of key, as generated by the application. |
| <code>keyQueryHash</code> | R/W | DisplayString | |
| <code>keyQueryRecovery</code> | R/W | Integer 1: Enabled 2: Disabled 3: No key 4: Unknown 5: Invalid 6: Unset | The value <code>unset</code> is never returned by the key table. If you set the value <code>unset</code> , the keys are not filtered based on any of the attributes. |
| <code>keyQueryProtection</code> | R/W | Integer 1: Module 2: Cardset 3: No key 4: Unknown 5: Invalid 6: Unset | The value <code>unset</code> is never returned by the key table. If you set the value <code>unset</code> , the keys are not filtered based on any of the attributes. |
| <code>keyQueryCardsetHash</code> | R/W | DisplayString | Hash of the card set protecting the key, if applicable. |
| <code>keyQueryFlags</code> | R/W | DisplayString | Certificate and public key flags. |
| <code>keyQueryExtraEntries</code> | R/W | Gauge32 | Number of extra key attributes. |
| <code>keyQuerySEESInteg</code> | R/W | DisplayString | SEE integrity key, if present. |
| <code>keyQueryGeneratingESN</code> | R/W | DisplayString | ESN of the module that generated the key, if present. |

| Node name | R/W | Type | Remarks |
|--------------------------------------|-----|---------|--|
| <code>keyQueryTimeLimit</code> | R/W | Gauge32 | Time limit for the key, if set (0 for no limit). |
| <code>keyQueryPerAuthUseLimit</code> | R/W | Gauge32 | Per-authentication use limit for the key (0 for no limit). |

34.7.6. netHSMNetwork sub-tree

`netHSMNetwork` sub-tree (`enterprises.nCipher.nC-series.netHSMNetwork`) contains information about the network interfaces of the network-attached HSMs (nShield Connect, nShield 5c).

To output the network interface information in SNMP, you must enable both NIC exposure settings. For information how to enable them, see [NIC Reporting](#).

34.7.6.1. netHSMNICTable sub-tree

`netHSMNICTable` (`enterprises.nCipher.nC-series.netHSMNetwork.netHSMNICTable`) provides information about the name and link state of the network interface.

| Node name | R/W | Type | Remarks |
|-----------------------|-----|---------------|--|
| <code>ifModule</code> | R | Integer | The remote module number of the HSM |
| <code>ifIndex</code> | R | Integer | The number of the network interface on the HSM |
| <code>ifState</code> | R | Integer | The link state of the network interface. 0 = Down 1 = Up |
| <code>ifName</code> | R | DisplayString | The name of the interface. The interface names are <code>eth0</code> and <code>eth1</code> . When interface bonding is enabled, a third interface will show as <code>bond0</code> . |

34.7.6.2. netHSMNICAddressTable sub-tree

netHSMNICAddressTable (`enterprises.nCipher.nC-series.netHSMNetwork.netHSMNICAddressTable`) provides the network address of the interface.

The maximum number of addresses that can be reported for an interface is 10.

| Node name | R/W | Type | Remarks |
|-----------|-----|---------------|--|
| ifAddress | R | DisplayString | The IPv4 or IPv6 address of the interface. When interface bonding is enabled, <code>eth0</code> and <code>eth1</code> interface addresses are displayed as <code>0.0.0.0</code> . |

34.7.7. 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](#).

34.7.7.1. Statistics sub-tree

The following table gives details of the nodes in the statistics sub-tree, and 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 <code>moduleStatsTable</code>). |
| hsuptime | R | Counter32 | Uptime of the hardserver. |
| cmdCountAll | R | Counter32 | Returned aggregated for all modules and all commands. |

| Node name | R/W | Type | Remarks |
|-----------------|-----|-----------|--|
| cmdBytesAll | R | Counter32 | |
| cmdErrorsAll | R | Counter32 | Returned as for <code>cmdCount</code> , 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 <code>cmdErrors</code> . |
| 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 | |

34.7.7.2. 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 <code>moduleStatsTable</code>). |
| uptime | R | Counter32 | Uptime of the module. |
| cmdCount | R | Counter32 | Returned aggregated for all commands. |
| cmdBytes | R | Counter32 | |
| cmdErrors | R | Counter32 | Returned as for <code>cmdCount</code> all the different error states aggregated into one counter. |
| replyCount | R | Counter32 | |
| replyBytes | R | Counter32 | |
| replyErrors | R | Counter32 | See notes above for <code>cmdErrors</code> . |

| Node name | R/W | Type | Remarks |
|------------------|-----|---------------|--|
| 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 | |
| CPUVoltage9 | R | DisplayString | |
| CPUVoltage10 | R | DisplayString | |
| CPUVoltage11 | R | DisplayString | |
| nvmlFreeSpace | R | Counter32 | Free space available on the HSM's NVRAM, in bytes |
| nvmlWearLevel | R | DisplayString | Wear level of the HSM's NVRAM |
| nvmlWornBlocks | R | DisplayString | Worn blocks in the HSM's NVRAM |

34.7.7.3. 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 |
|----------------------------------|-----|---------------|---|
| <code>connStatsIndex</code> | R | Integer32 | Index of this entry. |
| <code>connNumber</code> | R | Integer32 | Hardserver connection number. |
| <code>connUptime</code> | R | Counter32 | Uptime of the connection. |
| <code>connCmdCount</code> | R | Counter32 | Number of commands submitted through this connection. |
| <code>connCmdBytes</code> | R | Counter32 | Number of bytes submitted through this connection. |
| <code>connCmdErrors</code> | R | Counter32 | Number of marshalling errors on commands through this connection. |
| <code>connReplyCount</code> | R | Counter32 | Number of replies received by this connection. |
| <code>connReplyBytes</code> | R | Counter32 | Number of bytes received by this connection. |
| <code>connReplyErrors</code> | R | Counter32 | Number of marshalling errors on replies through this connection. |
| <code>connDevOutstanding</code> | R | Gauge32 | Number of commands outstanding on this connection. |
| <code>connQOutstanding</code> | R | Gauge32 | Number of commands outstanding in the hardserver queue. |
| <code>connLongOutstanding</code> | R | Gauge32 | Number of long jobs outstanding for this connection. |
| <code>connRemoteIPAddress</code> | R | IpAddress | IP Address of connection client. |
| <code>connProcessID</code> | R | Integer32 | Process identifier reported by connection client. |
| <code>connProcessName</code> | R | DisplayString | Process name reported by connection client. |
| <code>connObjectTotal</code> | R | Gauge32 | The total object count for a connection |

34.7.7.4. 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 |
|---|-----|---------|---|
| <code>connModuleStatsConnId</code> | R | Integer | Identity of this connection |
| <code>connModuleStatsModuleIndex</code> | R | Integer | Index of the module entry |
| <code>connModuleStatsObjectCount</code> | R | Gauge32 | The object count on this module for this connection |

34.7.7.5. 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 |
|--------------------------------|-----|-----------|-----------------|
| <code>netHSMModuleIndex</code> | R | Integer32 | Module number |
| <code>netHSMFanIndex</code> | R | Integer32 | Fan number |
| <code>netHSMFanSpeed</code> | R | Gauge32 | Fan speed (RPM) |

34.7.7.6. Software versions table

The following table gives details of the nodes in the software versions table

(`enterprises.nCipher.softwareVersions.softwareVersionsTable`):

| Node name | R/W | Type | Remarks |
|-------------------------------|-----|-----------------------|------------------|
| <code>compIndex</code> | R | Integer | Table index. |
| <code>compName</code> | R | DisplayString | Component name. |
| <code>compOutput</code> | R | Component output name | Component name. |
| <code>compMajorVersion</code> | R | Gauge | |
| <code>compMinorVersion</code> | R | Gauge | |
| <code>compPatchVersion</code> | R | Gauge | |
| <code>compRepository</code> | R | DisplayString | Repository name. |
| <code>compBuildNumber</code> | R | Gauge | |

34.8. SNMP agent command-line

34.8.1. 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. |

| Option | Description |
|-----------------|---|
| <code>-s</code> | This option logs warnings/messages to syslog. |
| <code>-r</code> | This option specifies not exiting if root-only accessible files cannot be opened. |
| <code>-I</code> | <code>[-]INITLIST</code> This option specifies a list of MIB modules to initialize (or not). Run <code>snmpd</code> with the <code>-Dmib_init</code> option for a list. |
| <code>-l</code> | <code>LOGFILE</code> This option prints warnings/messages to a file <code>LOGFILE</code> (by default, <code>LOGFILE=log/snmpd.log</code>). |

34.8.2. 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 |
|----------------------------|--|
| <code>snmpctest</code> | This utility monitors and manages SNMP information. |
| <code>snmpget</code> | This utility runs a single GET request to query for SNMP information on a network entity. |
| <code>snmpset</code> | This utility runs a single SET request to set SNMP information on a network entity. |
| <code>snmpgetnext</code> | This utility runs a single GET NEXT request to query for SNMP information on a network entity. |
| <code>snmpstable</code> | This utility obtains and prints an SNMP table. |
| <code>snmptranslate</code> | This utility translates SNMP object specifications into human-readable descriptions. |
| <code>snmpwalk</code> | This utility communicates with a network entity using repeated GET NEXT requests. |
| <code>snmpbulkwalk</code> | 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/>.

35. HSM status and error codes

The Entrust nShield 5s HSM is fitted with a tri-color LED on the back panel. This LED will typically indicate the operational state of the HSM, see [LED status](#). However, the LED can also indicate if the HSM is in an unrecoverable error state, see [LED error states](#).

Unrecoverable error state codes can also be retrieved remotely using the [enquiry](#) utility, see [Error codes accessed remotely](#).

35.1. LED status

The Entrust nShield 5s HSM is fitted with a tri-color LED on the back panel. This LED typically indicates the status of the HSM.

The following states indicate a normal operational state:

| Colour | Pattern | Meaning |
|-----------|------------|--|
| N/A | Blank | No power or processors not working. |
| Green | Solid | Power is good. Main processor has not started booting. |
| Cyan | Solid | Main processor is booting. |
| Cyan/Blue | Slow flash | Security processor firmware upgrade in progress. |
| Blue | Solid | System has booted, now idle. |
| Blue | Flickering | System is active - normal operation. |

The following states indicate an error within the HSM:

| Colour | Pattern | Meaning |
|----------|------------|---|
| Blue | Morse code | Error state for when the HSM is in an unrecoverable state, see LED error states for more information. |
| Red | Morse code | Error state for when the HSM is in an unrecoverable state see LED error states for more information. |
| Red | Fast flash | Security processor bootloader failure. |
| Blue/Red | Flash | Security processor detected a tamper condition. |
| Other | Any | Contact Entrust Support. |

35.2. LED error states

If the Entrust nShield 5s HSM encounters an unrecoverable error, it enters an error state. In an error state, the HSM does not respond to commands and does not write data to the bus. The LED displays a Morse code pattern to indicate a specific error state, see [Error codes shown on the LED](#).

In some cases you can reset an HSM in an error state by powering down the HSM and then reapplying power, or with `hsmadmin reset`. Not all errors can be reset in this way.

If any HSM goes into an error state, except as a result of you issuing the `nopclearfail --fail` command, contact Entrust Support, and give full details of your HSM set-up and the error code.

Entrust recommends that you contact Entrust Support even if you successfully recover from the error.

For troubleshooting information, see the relevant *Installation Guide* for your HSM.

35.2.1. Error codes shown on the LED

If an HSM enters an error state, the LED flashes with a Morse code pattern corresponding to an error code.



Error codes can also be retrieved remotely using the `enquiry` utility, see [Error codes accessed remotely](#).

All the LED error codes have three digits:

- The first digit is indicated by a number of dots.
- The second digit is then indicated by a number of dashes.
- The third digit is then indicated by a number of dots.

There is then a longer gap and the error code repeats.

The following guidelines are useful when reading LED code messages from the HSM:

- The duration of a dash (-) is three times the duration of a dot (.).
- The gap between components of a letter has the same duration as a dot.
- The gap between digits has the same duration as a dash.
- The duration of the gap between repeating codes is seven times the duration of a dot.

The numbers of dots/dashes and the Morse code equivalent is shown in the table below.

| Colour | Digits | Dots and dashes | Morse code | Meaning |
|--------|--------|-----------------|------------|---|
| Red | 1-1-1 | .-. . | E T E | Battery voltage out of spec |
| Red | 1-2-1 | .-- . | E M E | Crypto SerDes core voltage out of spec |
| Red | 1-2-2 | .---- . | E M I | Main processor SerDes core voltage out of spec |
| Red | 1-2-3 | .----- . | E M S | Main processor core voltage out of spec |
| Red | 1-2-4 | .----- . | E M H | Main processor SerDes core IO voltage out of spec |
| Red | 1-2-5 | .----- . | E M 5 | Crypto SerDes IO voltage out of spec |
| Red | 1-3-1 | .---- . | E O E | Main processor IFC IO voltage out of spec |
| Red | 1-3-2 | .---- . | E O I | DDR access voltage out of spec |
| Red | 1-3-3 | .---- . | E O S | DDR IO voltage out of spec |
| Red | 1-3-4 | .---- . | E O H | V12 voltage out of spec |
| Red | 1-3-5 | .---- . | E O 5 | Security processor voltage out of spec |
| Red | 1-5-1 | .----- . | E O E | Security processor temperature out of spec |
| Red | 1-5-2 | .----- . | E O I | Main processor temperature out of spec |
| Red | 1-5-3 | .----- . | E O S | Crypto temperature out of spec |
| Red | 1-5-4 | .----- . | E O H | Security processor app blank |
| Red | 1-5-5 | .----- . | E O 5 | Security processor app invalid |
| Red | 2-1-1 | ..-. . | I T E | Security processor secure state corrupted |
| Red | 2-1-2 | ..-. . | I T I | No bootloader heartbeat |
| Red | 2-1-3 | ..-. . | I T S | Board-ID PROM failed |
| Blue | 2-1-5 | ..- . | I T 5 | Firmware signature auth failure |
| Red | 2-2-2 | ..-- . | I M I | Crypto known-answer tests failed |
| Red | 2-2-3 | ..-- . | I M S | RNG driver failed |
| Red | 2-2-4 | ..-- . | I M H | FIPS DRBG failed |
| Red | 2-2-5 | ..-- . | I M 5 | OpenSSL failed |
| Red | 2-3-1 | ..--- . | I O E | OpenSSH failed |
| Red | 2-3-2 | ..--- . | I O I | Library signature verification failed |
| Red | 2-3-3 | ..--- . | I O S | FPGA initialisation failed |
| Red | 2-3-4 | ..--- . | I O H | Init script failed |

35.3. Error codes accessed remotely

If an HSM enters an error state, you can retrieve error codes using the `enquiry` utility. These codes appear in the `hardware status field` of the `Module` and are included in the hardserver log.

There are three error categories:

- [Runtime library errors.](#)
- [Hardware driver errors.](#)
- [Operational mode errors.](#)



Error codes are also indicated by the LED on the back of the HSM, see [LED error states.](#)

35.3.1. Runtime library errors

The runtime library error codes described in the following table indicate one of the following:

- There is a bug in the firmware.
- There is a hardware fault.

If any of these errors occur, reset the HSM.

| Code | Meaning |
|-------|---|
| O L C | SIGABRT: assertion failure and/or <code>abort()</code> called. |
| O L D | Interrupt occurred when disabled. This is more likely to indicate a hardware problem than a firmware problem. |
| O L E | SIGSEGV: access violation. This is more likely to indicate a hardware problem than a firmware problem. |
| O L J | SIGFPE: unsupported arithmetic exception (such as division by 0). |
| O L K | SIGOSERROR: runtime library internal error. |
| O L L | SIGUNKNOWN: invalid signal raised. |

35.3.2. Hardware driver errors

The hardware driver error codes described in the following table indicate one of the following:

- Some form of automatic hardware detection has failed.
- There is a bug in the firmware.
- The wrong firmware has been loaded.

If any of these errors is indicated, contact Entrust Support.

| Code | Meaning |
|---------|--|
| H L | M48T37 NVRAM (or battery) failed. |
| H C V | CPLD wrong version for PCI policing firmware. |
| H C X | No crypto offload hardware detected. |
| H P P | PCI Interface Policing failure. |
| H V | Environment sensors failed. For example, the temperature sensor. |
| H D | Failure reading unique serial number. |
| H R | Random number generator failed. |
| H R F O | FIPS continuous RNG failed. |
| H R A O | Periodic RNG test failed. |
| H R S | RNG startup failed. |
| H R T | RNG selftest failed. |
| H R T P | Periodic (scheduled daily) RNG selftest failed. |
| H R M | RNG data matched. |
| H R Z | Impossible RNG Failure (match after PRNG). |
| H S S | Security processor internal semaphore error. |
| H O | Token interface initialization failed. |
| H E | EEPROM failed on initialization. |
| H C | Processing thread initialization failed. |
| H C P | Card poll thread initialization failed. |
| H F | Starting up crypto offload. |
| H C V | CPLD version number incorrect. |
| H J V | IPC-watcher failed. |
| H J U | IPC-EPD failed. |
| H J R | Module reset notification failed. |
| K R | RSA selftest failed. |

| Code | Meaning |
|-------|--|
| H H D | Unique serial number detection failed. |
| H H P | PCI bus hardware detection failed. |
| H H R | RTC hardware detection failed or random number generator detection failed. |
| H S C | Error writing correct SOS message. |

35.3.3. Operational mode errors

The runtime library error codes described in the following table indicate one of the following:

- There is a bug in the firmware.
- There is a hardware fault.

| Code | Meaning | Action |
|-------|--|---|
| T | Temperature of the HSM has exceeded the maximum allowable. | Restart your host computer, and improve HSM cooling. For the cooling requirements for your HSM, see the <i>Installation Guide</i> . |
| D | Fail command received. | Reset HSM by turning it off and then on again. |
| G G G | Failure when performing C learUnit or F ail command. | Contact Entrust Support. |
| I J A | Audit logging: failed to send audit log message. This can occur for any type of log message. That is, a log message, signature block or certifier block. | Contact Entrust Support. |
| I J B | Audit logging: no module memory (therefore failed to send audit log message). | Contact Entrust Support. |
| I J C | Audit logging: key problem or FIPS incompatibility (therefore failed to sign audit log message). | Contact Entrust Support. |
| I J D | Audit logging: NVRAM problem (therefore failed to configure or send audit log message). | Contact Entrust Support. |

35.4. Other errors

| Code | Meaning | Action |
|------|--|---|
| SFP | The Security Fuse Processor (SFP) has failed and is unable to handle further requests sent from the client's hardserver. | Restart the HSM. This resets the SFP. If this does not resolve the issue, or the SFP fails again, contact Entrust Support. |

For information on error codes not listed on this page, contact Entrust nShield Technical Support: nshield.support@entrust.com.

36. Application Performance Tuning

36.1. 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.

36.2. 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.

36.3. 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`.

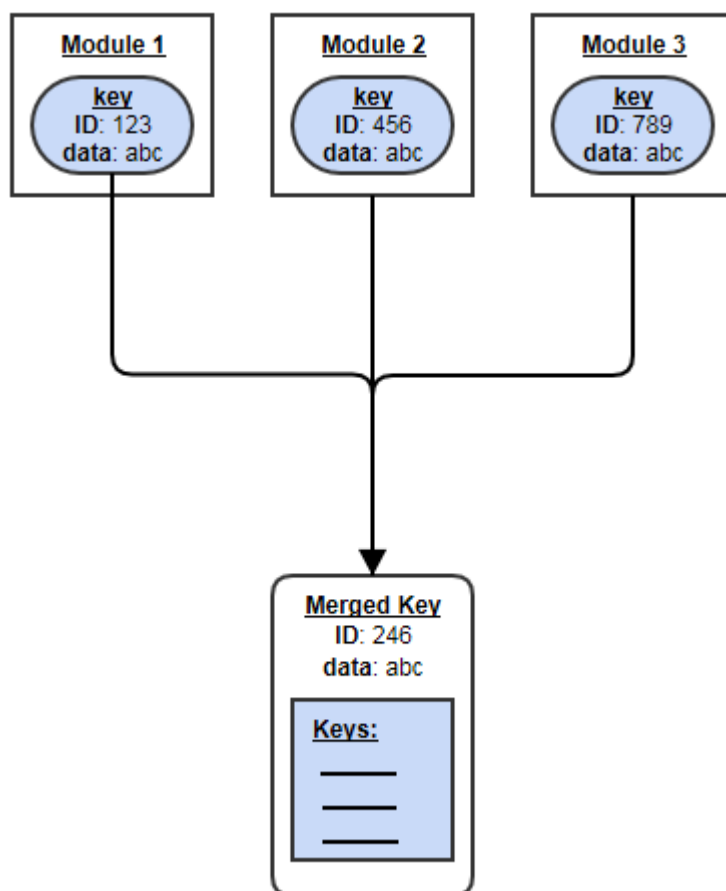
On some systems the default operating system scheduling algorithm is also not optimized for highly multi-threaded applications. A real-time scheduling algorithm such as the POSIX round-robin scheduler may yield noticeable performance improvements for your application.

36.4. 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.

37. 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

corresponding 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.

38. Product returns

If you need to return your nShield HSM, contact Entrust nShield Support for instructions:

<https://nshieldsupport.entrust.com>

Before returning your nShield HSM you should return it to factory state by following the procedure described in [nShield 5s modes of operation](#).

39. SSH Client Key Protection

39.1. SSH Services

This table explains what the different services are used for, to help inform what protection settings are appropriate for the client keys for those services in a deployment.

| Service | Service Description |
|-----------------------|--|
| <code>sshadmin</code> | Main authority for administration of the SSH services on the HSM. This key should have the strongest protection. |
| <code>ncoreapi</code> | nCore API service. Used by the hardserver for routine communication with the HSM. |
| <code>setup</code> | Setup service. Used for some administration options such as factory state. |
| <code>updater</code> | Updater service. Used for installing signed firmware upgrade packages. |
| <code>monitor</code> | Monitor service. Used for diagnostic operations such as retrieving logs with <code>hsmadmin logs</code> . |
| <code>launcher</code> | Launcher service. On versions with CodeSafe 5 support, this is used for starting CodeSafe 5 applications on the HSM. |

39.2. SSH Client Key Encryption

SSH client keys are protected by a passphrase derived from one or more inputs including machine IDs and user-supplied passphrases.

These passphrases are derived automatically by applications which use the SSH keys, and with the exception of the option of user-supplied passphrases do not prompt the user.

39.2.1. Available SSH Key Protection Options

The following are the supported protection options for SSH keys. Multiple options can be combined in any order.

On Linux, although the hardserver runs as `nfast` user, it starts as root and then drops privileges. The hardserver derives any SSH key passphrases before dropping privileges, and so can use protections that are available only to root.

| Option | Description |
|----------|--|
| K | Per-key nonce. Ensures that the derived passphrase is unique to the key. |
| F | Fixed nonce present in the nShield client software. |
| S | System UUID. Ties the key to the local machine. On Linux, this is only readable by root. This is available on most systems. |
| B | Baseboard UUID. Ties the key to the baseboard (motherboard) of the local machine. On Linux, this is only readable by root. This may not be available on all systems. |
| G | Global nonce. Ties the key to the local OS install. The global nonce is readable only by root or Administrators and is generated the first time the option is used to protect a key. |
| P | User passphrase. The key will require a user-supplied passphrase (in addition to any other protection options specified). |

39.2.1.1. User-supplied SSH Key passphrases

If a key is generated with protection by a user-supplied SSH key passphrase, there will be an interactive prompt on the console to enter and confirm the passphrase. When a key is loaded with passphrase protection, there will be an interactive prompt on the console to enter the passphrase.

To avoid interactive prompts for automation purposes, the user passphrase can be supplied using the environment variable `NFAST_KEYPROT_PASS`. If a user passphrase is specified for the `ncoreapi` service SSH key, then the environment variable is the only way to supply the passphrase as interactive prompts are disabled for the hardserver service.

39.3. Setting Protections on SSH keys

Protections are set on SSH keys when they are generated, either during `hsmadmin enroll` (if the keys are not already present) or during `hsmadmin keys roll` (if switching to a freshly generated set of SSH client keys).

Protections can be overridden using environment variables that are set in the environment of the above commands when the keys are generated. The protections for all SSH service keys can be overridden using the `NFAST_KEYPROT` environment variable. Individual SSH service keys can have their protections set directly using per-service environment variables as specified in the table below. If both `NFAST_KEYPROT` and the per-service environment variable are set, the per-service environment variable takes precedence.

| Service | Default Protection | Environment Variable |
|----------|--------------------|------------------------|
| sshadmin | KFSG | NFAST_SSHADMIN_KEYPROT |
| ncoreapi | KFSG | NFAST_NCOREAPI_KEYPROT |
| setup | KFSG | NFAST_SETUP_KEYPROT |
| updater | KF | NFAST_UPDATER_KEYPROT |
| monitor | KF | NFAST_MONITOR_KEYPROT |
| launcher | KF | NFAST_LAUNCHER_KEYPROT |

39.4. Permissions on SSH keys

Access to SSH keys is controlled by permissions on the directories they are created in. The key directories are created with permissions during installation of the nShield software.

Linux

All keys are owned by user `root`, except for `ncoreapi` which is owned by user `nfast`. The table below shows what group can read each of the service keys by default. The group that can read each key can also be overridden with the environment variables listed below if set in the environment when running the install script `/opt/nfast/sbin/install`. The install script will always set the owner and group on the key, so if custom groups are used, they must be specified every time the install script is run. If the group specified in the environment variable does not exist, it will be created automatically by the install script.

| Service | Default Group | Environment Variable |
|----------|---------------|----------------------|
| sshadmin | root | NFAST_SSHADMIN_GROUP |
| ncoreapi | root | NFAST_NCOREAPI_GROUP |
| setup | root | NFAST_SETUP_GROUP |
| updater | nfastadmin | NFAST_UPDATER_GROUP |
| monitor | nfastadmin | NFAST_MONITOR_GROUP |
| launcher | nfastadmin | NFAST_LAUNCHER_GROUP |

Windows

SSH keys are created under the directory `%NFAST_SERVICES_HOME%\client` (`C:\ProgramData\nCipher\services\client` by default). The installer sets permissions on this directory to be read/write by the built-in local Administrators group only. If different permissions are wanted on particular keys to enable particular users or groups

to perform certain operations, then these must be set manually on the keys and parent directories to permit the required access.