Hewlett Packard Enterprise

HPE NonStop Security Hardening Guide

Part Number: 830852-004

Published: December 2018

Version: 1.5

Table of contents

Overview

This guide assumes basic familiarity with NonStop systems. If you do not have that background, begin by reading the HPE NonStop Security Overview technical white paper (see Documentation, below).

The intended audiences for this document are NonStop system security administrators looking for guidance on best practices and NonStop system auditors looking for both guidance on best practices and a more in-depth understanding of NonStop security at both the system and subsystem level.

The general structure is to provide relatively detailed explanatory information in sections labeled “background” and more concrete hardening recommendations in sections labeled “best practices”. The two sections are combined for topics having a relatively short description.

What's new in v1.5

New content on NonStop OmniPayment Switch has been added under the Remote Access Section. In addition, there are also a number of small updates to keep the document current.

HPE references

Documentation

A NonStop security technical overview can be found at www.hpe.com/info/nonstop-security, under "System Security". More detailed information is available in the NonStop manuals collection (www.hpe.com/info/nonstop-docs).The following documents are of particular interest. They can be found at www.hpe.com/info/nonstop-docs unless otherwise noted.

Abbreviation	Title
BACKBOXSSL	BackBox H4.02 SSL Setup
BACKBOXTAPE	BackBox H04.02 Tape Encryption Option
BACKBOXUSER	BackBox H04.02 User Manual
CIPMGMT	CLuster I/O Module (CLIM) Configuration and Management Manual (H06.16+, J06.05+)
DSMSCMTAPE	DSM/Tape Catalog Operator Interface (MEDIACOM) Manual
DSMSCMUSER	DSM/SCM User’s Guide
EMS	EMS Manual
FUPREF	File Utility Program (FUP) Reference Manual
FWMATRIX	NonStop Firmware Matrices
GUARDPROCREF	Guardian Procedure Calls Reference Manual
GUARDUSE	Guardian User’s Guide
ILOSEC	HPE Integrated Lights-Out Security; Technology brief, 7th edition
ILO4SEC	HPE Integrated Lights-Out Security Technology Brief
IPSEC	NonStop IPSec Overview
ITPADMIN	iTP Secure WebServer System Administrator’s Guide
JDBCMXT4REF	JDBC Type 4 Driver Programmer’s Reference for SQL/MX Release 3.2.1
MXDMUSER	MXDM User Guide for SQL/MX Release 3.2.1 (H06.26+, J06.15+)

MEDIACOM	DSM/Tape Catalog Operator Interface (MEDIACOM) Manual
NETBATCH	NETBATCH Manual
NSCINSTALL	NonStop System Console Installation and Configuration Guide
NSCPOLBP	NonStop System Console Security Policy and Best Practices
NSEINSTALL	NonStop Software Essentials Installation and Quick Start Guide
NSEUSER	NonStop Software Essentials User Guide
NSJSPADMIN	NonStop Servlets for JavaServer Pages 7.0 System Administrator’s Guide
NSJ6PROGREF	NonStop Server for Java 6.0 Programmer’s Reference
NSJ7PROGREF	NonStop Server for Java 7.0 Programmer’s Reference
NSSECOV	NonStop Security Overview
NSSOAPUSER	NonStop SOAP User’s Manual
NSSOAPUSER41	NonStop SOAP 4.1 User’s Manual
NSVLEGUIDE	NonStop Volume Level Encryption Guide
OPENSSLSHELL	NonStop cf SSL OpenSSL Shell Reference Manual
OSMCG	OSM Configuration Guide
OSMSCUG	OSM Service Connection User’s Guide
OSSMOG	Open System Services Management and Operations Guide (G06.29+, H06.08+, J06.03+)
OSSSHELL	Open System Services Shell and Utilities Manual (G06.29+, H06.08+, J06.03+)
OSSSYSCALL	Open System Services System Calls Reference Manual
PATHITSMGMT	Pathway/iTS System Management Manual
SAFEADMIN	Safeguard Administrator’s Manual (G06.29+, H06.08+, J06.03+)
SAFEAUDIT	Safeguard Audit Service Manual (G06.29+, H06.08+, J06.03+)
SAFEREF	Safeguard Reference Manual (G06.29+, H06.08+, J06.03+)
SAFEUSER	Safeguard User’s Guide (G06.29+, H06.08+, J06.03+)
SAMBA	Samba on NonStop User Manual
SCFKERNEL	SCF Reference Manual for the Kernel Subsystem
SECUREDATA	Pointer to HPE SecureData Enterprise manuals
SECMGMT	Security Management Guide (G06.29+, H06.08+, J06.03+)
SECURETAPE	NonStop cF Secure Tape Reference Manual
SQLMPMGMT	SQL/MP Installation and Management Guide
SQLMXCSREF	SQL/MX Connectivity Service Administrative Command Reference
SQLMXINSTALL	SQL/MX Installation and Upgrade Guide
SQLMXMGMT	SQL/MX Management Guide
SQLMXREF	SQL/MX Reference Manual

SSHLIB	NonStop cF SSH Library Reference Manual
SSHREF	SSH Reference Manual
SSLATLIB	NonStop cF SSL AT Library Reference Manual
SSLLIB	NonStop cF SSL Library Reference Manual
SSLREF	SSL Reference Manual
TACLREF	TACL Reference Manual
TCPIPMGMT	TCP/IP Configuration and Management Manual
TCPIPV6MGMT	TCP/IPv6 Configuration and Management Manual
TELSERV	Telserv Manual
TLSCONFIG	TLS Protocol Version Configuration for NonStop Products (coming soon)
TMFREF	TMF Reference Manual (H06.06+, J06.03+)
TSMPMGT	TS/MP System Management Manual (H06.05+, J06.03+)
TSMPPATH	TS/MP Pathsend and Server Programming Manual
TSMP26MGT	TS/MP 2.6 System Management Manual
TSMP25PATH	TS/MP 2.5 Pathsend and Server Programming Manual
VLEGUIDE	NonStop Volume Level Encryption Guide
VTSCONFIG	Virtual TapeServer Configuration Guide
XACREF	XYGATE Access Control Reference Manual
XMAREF	XYGATE Merged Audit Reference Manual
XSWREF	XYGATE Compliance PRO (XSW) Reference Manual
XUAREF	XYGATE User Authentication Reference Manual

Additional HPE documentation

Insight Remote Support:

http://www.hpe.com/info/insightremotesupport/docs white papers:

- (IRSSEC) HPE Insight Remote Support 7.x Security White Paper

HPE Security data-centric encryption and tokenization SecureData security solutions (MicroFocus products as of August 2017):

- (VOLTAGE) https://software.microfocus.com/en-us/products/voltage-data-encryption-security/overview
- (VOLTPAY) https://software.microfocus.com/en-us/solutions/data-security-encryption
- (VOLDPCI) https://software.microfocus.com/en-us/solutions/data-security-encryptionHPE ArcSight Security Information and Event Management (SIEM) solution (MicroFocus product as of August 2017):
- (ARCSIGHT) https://software.microfocus.com/en-us/products/siem-security-information-event-management/overview

- HPE Atalla Network Security Processor (NSP) - hardware security modules for high-performance cryptography and key management (MicroFocus product as of August 2017):
- (ATALLANSP) https://software.microfocus.com/en-us/products/hardware-security-module-hsm-atalla-nsp/overview

HPE NonStop security products

- Safeguard - included with the L, J and H series, optional for the G series
- XYGATE Merged Audit (XMA) – included with L-series, J-series and H-series commercial systems sold since September 2010, optional for other systems
- NonStop SSH – included with L-series, J-series and H-series commercial systems sold since September 2010, optional for other systems
- NonStop SSL – included with L-series, J-series and H-series commercial systems sold since September 2010, optional for other J-series and H-series systems
- XYGATE User Authentication (XUA) – included with L-series and NS56000c systems, optional for other systems
- NonStop Volume Level Encryption (VLE) – optional for CLIM-attached storage (L-Series, J-series systems, some H-series systems).
- NonStop Secure Tape - optional
- XYGATE Access Control (XAC) – optional
- XYGATE Command Monitor (XCM)
- XYGATE Compliance PRO (XSW) – optional
- SFTP API – optional
- XMA BASE24 plug-in – optional
- XMA BASE24 eps plug-in – optional
- XMA AJB RTS plug-in – optional
- XMA SECOM plug-in – optional

External sources for NonStop security products and packages

NonStop security partners

HPE has a robust community of security partners who develop products that extend the capabilities of NonStop servers. HPE encourages you to visit our partners' websites to learn about their offerings. The Resources section includes a list of security partner websites.

Community-maintained sources

The Connect user group's ITUGLIB team maintains a user-supported collection of open source packages that have been ported to NonStop servers. The packages are provided on an as-is basis, with usage subject to their individual licenses. Security-related ports include openssl, openssh, sudo, gnupg, libgcrypt, and prngd, among others. The link to ITUGLIB is http://ituglib.connect-community.org/apps/Ituglib/SrchOpenSrcLib.xhtml.

Security hardening implementation approaches

Any document on how to harden a specific platform can only give some general guidelines. To establish your specific NonStop security strategy, work with your internal security groups and auditors to determine your system security requirements and map them against your current configuration and practices to identify gaps. Also factor in recommendations on best practices and security settings from your software providers. When that investigation is reasonably complete, develop a phased remediation plan for tightening controls as identified. You should consult with both your auditors and your internal security group on acceptable phasing. You should regard security hardening as an ongoing project that will require continuing attention and updating.

While achieving compliance with applicable laws and regulations is very important, make sure that your focus is on developing a robust security strategy that is appropriate to the risk management needs of your business and customers. There are a number of resources available for you to consult, including but not limited to XYPRO1, XYPRO2, various Connection articles, Connect presentations such as KNIGHT1, and NonStop security partners' websites. You can find some external references at the end of this white paper.

Here is an example of a framework to use for phased implementation on the NonStop platform. Your choice of levels would depend on what products are going to be deployed and the amount of control/auditing/monitoring desired.

Control

Level 0 – Guardian security vectors, OSS file/directory permissions

Level 1 – Basic Safeguard configuration including user management, required security administration groups, coarse-grained Safeguard and OSS Access Control Lists (ACLs), and Safeguard granular OBJECTTYPE controls

Level 2 –More granular ACLs, enhanced user management (XUA or partner product, partner password quality SEEP)

Level 3 – Access management (XAC or partner product), two-factor authentication (XUA or partner product)

Level 4 – Partner authorization (Guardian and OSS) SEEP

Auditing

Level 1– All changes to Safeguard’s configuration, failed object access attempts, all logon operations

Level 2 – All Guardian access (via Safeguard global AUDIT-CLIENT-GUARDIAN), successful object access attempts

Level 3 – Most or all OSS access (via Safeguard global AUDIT-CLIENT-OSS and AUDITENABLED set for sensitive OSS filesets)

Monitoring

Level 1– XMA with local basic reporting

Level 2– XSW basic monitoring, XMA export to enterprise Security Incident Event Management (SIEM) and/or customization

Level 3 – XSW customization

Data in Motion

Level 1– SSH/SFTP, TLS/SSL, and/or IPSec

Data at Rest

Level 1 – Selective CLEARONPURGE, disk sanitization, tape encryption (VLE or partner)

Level 2 – Disk volume-level encryption (VLE)

Level 3 – Application field-level or column-level encryption or tokenization (Partner products)

When you are building your implementation approach, be sure that you also review the sections on NonStop System Console (NSC) and Cluster I/O Module (CLIM) security.

Initial system hardening overview

In order to align NonStop system security with your enterprise's security policies and practices, you will undoubtedly need to adjust some of its security configuration elements prior to putting a new system into use. Areas to review include:

Safeguard configuration and management

Safeguard configuration should be at the top of your list, so you will have appropriate system-level security in place and audit enabled from the beginning.

HPE expects that you will need to tailor various attributes to match your enterprise's policies. Use the SAFECOM ALTER SAFEGUARD command to make changes.

SAFECOM is Safeguard's command line interface. GUI interfaces to Safeguard are available from partners.

When SAFEGUARD is started, immediately create all of its OBJECTTYPE records and set their ownership and access rights to appropriate values for your environment.

Audit

Enable audit according to your enterprise's policies, including audit of Safeguard configuration changes (AUDIT-OBJECT-MANAGE-PASS, AUDIT-OBJECT-MANAGE-FAIL, AUDIT-SUBJECT-MANAGE-PASS and AUDIT-SUBJECT-MANAGE-FAIL) and authentication (AUDIT-AUTHENTICATE-PASS, AUDIT-AUTHENTICATE-FAIL).

Note: OSS audit must be enabled at the fileset level in addition to having AUDIT-CLIENT-OSS enabled via Safeguard.

You can configure audit at both the object level and globally. Test your initial configuration prior to putting it into production to ensure that it will not have negative performance impact and to help you size your audit pools.

Configure the Safeguard audit service, including audit pools and the audit failure recovery policy, to match your company's policies, including audit retention online and in archived form.

Caution: Ensure that audit pools are large enough to sustain the anticipated audit rate and are placed on disk volumes where they will not adversely affect production application performance. Move the Safeguard audit pools onto disk volumes other than $SYSTEM.

Caution: Audit pool volumes for Safeguard and any configured Safeguard Security Event Exit Processes (SEEPs) should be configured between CPUs 0 and 1 to avoid potential hangs during system load due to audit volume inaccessibility.

You have several configuration options for handling situations where no file is available for writing audit, including recycling (overwriting) the oldest online audit file, allowing activity to continue unaudited until you make a file available, or denying authorization and access requests that require auditing – which almost certainly will have a negative impact on application and system operations.

Unless your corporate policy is different, HPE recommends configuring audit recycling. If you do configure recycling, make sure that your audit pools are large enough to support your audit retention policy.

Ensure that Safeguard audit trails are backed up in accordance with your audit retention policy.

You can configure Safeguard to either write audit records to disk individually as they are generated (WRITE-THROUGH CACHE ON) or to cache the records in memory and write multiple records at a time (WRITE-THROUGH CACHE OFF). For performance reasons, HPE recommends setting WRITE-THROUGH CACHE to OFF (the default).

Note: You run a small risk of losing some audit records after a failure with WRITE-THROUGH CACHE set to OFF.

If possible, export Safeguard and other subsystem audit to your enterprise SIEM (HPE ArcSight or other). XYGATE Merged Audit (XMA) can be configured to export audit in syslog format, and also can export to HPE ArcSight in CEF format. See XMAREF for details.

Safeguard configuration files

Verify that the various user ID files (which might include some or all of USERID, USERAX, USERIDO, USERIDKO, OUSERAX, OUSERID, and OUSERA) in $SYSTEM.SYSTEM belong to SUPER.SUPER and are secured “----“.

Safeguard has a subvolume named SAFE on each disk volume. For disks other than $SYSTEM, SAFE contains files specific to that disk (e.g. GUARD, PATGUARD). There are additional files in $SYSTEM.SAFE. Verify that the files in the SAFE subvolumes belong to SUPER.SUPER and are secured “----“.

Use the SAFECOM command DISPLAY COMMANDS or a partner product to save your current Safeguard configuration on a regular basis.

Passwords

Set PASSWORD-ALGORITHM to HMAC256. (Note: PASSWORD-ENCRYPT defaults to ON.)

Set PASSWORD-COMPATIBILITY-MODE to OFF on new systems and systems that have been running with HMAC256 for some time. Its main purpose is as a temporary mechanism to help smooth migrations from DES to HMAC256 password protection; after that, it needs to be set to OFF to get the full benefit of passwords longer than eight characters.

If using a partner Password Quality (PQ) Security Event Exit Process (SEEP), configure Safeguard to consult it and configure both the SEEP and, insofar as is possible, Safeguard according to your enterprise’s password standards. If your enterprise requires checks for banned passwords such as dictionary words and frequently-used passwords such as “password”, “1234” or the user name, then you will need to acquire a PQ SEEP; Safeguard has basic checks such as minimum length and required character types. Typically, PQ SEEPs also are able to synchronize password changes across multiple systems.

Set the various password quality attributes, insofar as possible, to match your enterprise standards, including PASSWORD-HISTORY and PASSWORD-MAY-CHANGE as well as the various –REQUIRED, -REQ and –ALLOWED attributes. Set PASSWORD-EXPIRY-GRACE to a non-zero value if you want to allow a grace period for logging on after a password has expired. Make these changes even if you are configuring a PQ SEEP so you still have some degree of protection in place if you ever have to temporarily disable the SEEP.

Set PASSWORD-MINIMUM-LENGTH to match your corporate policy (HPE recommends a minimum of 12 characters), and set PASSWORD-MAXIMUM-LENGTH to at least that value. If your policy allows pass phrases, set PASSWORD-SPACES-ALLOWED to ON. You need to decide whether to configure PASSWORD-REQUIRED = ON or OFF. This attribute affects the ability of SUPER.SUPER to “log down” to any other user ID without supplying a password, and of group managers to do likewise to individual members of their groups. See the discussions under SUPER.SUPER and Group managers for additional information.

See SAFEADMIN Section 2, Controlling User Access, for details.

Users and aliases

The initial configuration established by HPE includes only SUPER.SUPER (user 255, 255) and NULL.NULL (0, 0). Change the password for SUPER.SUPER and delete or at least freeze

NULL.NULL unless needed as it is a well-known user and an easy target for attacks. If you do keep it, change its password. For individual accountability, consider either adding a SUPER.SUPER alias for each distinct user who will be using the ID or installing a product such as XYGATE User Authentication (XUA) or XYGATE Access Control (XAC) to enforce “logging up” to SUPER.SUPER from another, already-authenticated user. Also consider freezing SUPER.SUPER and any aliases that you create for it. On a properly-configured system, use of SUPER.SUPER for daily operation should not be required. For more information on options, see the detailed discussions of user and SUPER.SUPER management below.

Consider integrating NonStop user management with your enterprise user management system. You can configure XYGATE User Authentication (XUA) as an LDAP client; see XUAREF Section 12, Configuring the LDAP Interface, for details. LDAP client support also is available in partner products. Kerberos support is available in partner products. Multi-Factor Authentication (MFA) is a form of authentication that requires presentation of multiple pieces of information, typically from at least two of the following categories: something you know, something you have, and something you are. Requirements for MFA for command-level administrative access, particularly for non-console administrative access, are increasing. You can configure XUA to use the RSA SecurID interface to consult an RSA server; see XUAREF Chapter 16, Configuring the RSA SecurID Interface, for details. MFA support also is available in partner products.

Add other users and aliases as needed. For system management purposes, other members of the SUPER administrative group (referred to collectively in this document as SUPER.notsuper) have a number of inherent system configuration modification privileges, so their assignment and usage should be carefully controlled and monitored. You can use XAC or a partner access control product to grant particular non-SUPER users specific SUPER group privileges, which would allow you to avoid configuring any SUPER.notsuper users.

If you find it necessary to share a user ID among multiple people, one alternative is to define an alias for each person. Aliases assist in providing accountability, as Safeguard audit will reflect the alias name rather than the underlying user ID. However, note that all forms of system-level authorization checking (including Guardian, Open System Services (OSS), Safeguard ACLs, OSS ACLs, and SQL/MX GRANT/REVOKE) are configured and performed at the user ID level rather than the alias level.

If you are defining any Safeguard-protected terminals, you can set system-wide defaults for the defined command interpreter and associated parameters by altering the terminal definition.

See SAFEADMIN Section 7, Securing Terminals, for additional information.

Security groups

Populate the appropriate Safeguard security groups with the users who you want to have manage Safeguard and OSS audit configuration. Until you do so, all SUPER group members have full operational privileges. See SAFEREF, Section 13, for details.

Use SAFECOM ADD SECURITY-GROUP SECURITY-ADMINISTRATOR to add the authorization record for that group, and give execute authority on the ACL to the users who will belong to the group. This will allow designated users to perform certain Safeguard management functions without being authenticated as SUPER.SUPER. You also might choose to configure and populate the SYSTEM-OPERATOR group, whose members have day-to-day operational privileges but cannot fundamentally alter Safeguard’s configuration.

Consider populating all Safeguard security groups to avoid having the default memberships apply.

See SAFEADMIN, Section 6, Managing Security Groups, for details.

Note: Adding a non-SUPER-group user to the SUPER group as a file-sharing member does not grant that user all SUPER group privileges. See the discussions of the SUPER group and file-sharing membership in SAFEREF, Section 8 for details.

Administrative groups

Administrative groups are implicitly defined when the first user is added. For groups that will have managers (user IDs *, 255), one option is to define the manager first and allow him/her to add new members to that group directly. If you will have multiple applications running on the same system, you may wish to give each application its own group and group manager. Alternatively, you can use the OBJECTTYPE USER command to define a set of users who have the authority to add other users to the Safeguard database.

See SAFEADMIN Sections 2, 3, 5 and 6 for details.

Authentication

The NonStop system uses password-based authentication. Support for two-factor identification is available in XUA and partner products, along with more advanced controls (IP address, time of day etc.). See XUAREF for details.

Authentication attempts from the Guardian environment can trigger consultation with a $CMON process which might be XCM or a home grown or a partner product. If you are using one, you need to configure Safeguard accordingly. If you do use a $CMON process, make sure that you also have a process for keeping it in sync with Safeguard (e.g., not denying a logon after Safeguard has authenticated a user). If you do not use a $CMON process, use a Safeguard PROCESS record to protect that process name from misuse.

Configure BLINDLOGON = ON (default) to make it possible to prevent password display by not allowing the password to be entered on the same line as the user name or ID. If you want all users to log on by user name (recommended as it is harder to guess) instead of user ID, set NAMELOGON = ON (default).

Safeguard has several configuration attributes (AUTHENTICATE-MAXIMUM-ATTEMPTS, AUTHENTICATE-FAIL-TIMEOUT, AUTHENTICATE-FAIL-FREEZE that control how authentication failures are handled. These attributes apply uniformly to all users, including SUPER.SUPER. Adjust the defaults as needed to match your security policy. Safeguard also keeps per-user information with respect to authorization attempts, displayed by the SAFECOM INFO USER, DETAIL command. Consider setting AUTHENTICATE-FAIL-FREEZE to OFF to prevent a Denial-of-Service (DoS) attack that freezes all users.

If you need to control authentication failure handling on a per-user basis, use XUA or a partner access control product with those capabilities. One good reason for implementing per-user controls is to prevent SUPER.SUPER from being frozen due to authentication failures while still allowing other users to be frozen.

Authentication attempts can be audited through Safeguard, XUA, and/or partner authentication products.

See SAFEADMIN, Section 3, Controlling User Access, SAFEREF Section 5, User Security Commands, and XUAREF for details.

Access control

Follow the principle of least privilege required when securing system resources.

Guardian uses a read/write/execute/purge vector to control file access.

You can augment base Guardian access controls with Safeguard Access Control Lists (ACLs) for finer granularity.

Safeguard supports the use of an authorization SEEP for Guardian resources. If you are using a partner Guardian authorization Security Event Exit Process (SEEP), configure Safeguard to consult it. Otherwise, use Safeguard ACLs and, optionally, XYGATE Access Control (XAC).

Note: Use of an authorization SEEP does increase the authorization pathlength. Make sure that you understand when the SEEP might access a remote node for information, as that action potentially could increase authorization latency noticeably if communication over Expand is temporarily slow due to network issues.

Base OSS permissions are the standard POSIX/UNIX permissions:

You can augment base OSS (POSIX) permissions with OSS ACLs for finer granularity.

The OSS Name Server supports the use of an authorization SEEP for OSS resources. If you are using a partner OSS authorization SEEP, use SCF and Safeguard to configure the Name Server to consult it. Otherwise, use OSS ACLs.

If you are using authorization SEEPs, you still should also configure at least minimal Safeguard/OSS ACLs so you still have some level of protection should you need to temporarily disable one or both SEEPs.

Access control attempts can be audited through Safeguard, XAC, and/or partner authorization products. If you have a requirement for it, you can configure keystroke-level logging at a granular level for one or more users in XAC.

See SAFEADMIN, SAFEREF, XACREF (Guardian SEEP) and OSS MOG (OSS SEEP) for details.

Default passwords

Use SAFECOM to change the NonStop server passwords for:

- SUPER.SUPER
- NULL.NULL (if not freezing or deleting the user)

Consider changing passwords on underlying system components. Use the following interfaces and commands:

- CLIM root user (each CLIM): CLIMCMD <climname> passwd, logged on as a SUPER group member
- CLIM user (each CLIM): CLIMCMD <climname> passwd user
- NonStop System Console (NSC) Administrator and other users (each NSC): Log on as Administrator using the default password, press Ctl+Alt+Del, select Change a password and enter your new password. While still logged on as Administrator, click Start->Administrative Tools->Services and right-click on OpenSSH Server and then Properties. Under the Log On tab, check that this account shows .\Administrator as the user. Enter and reenter the new password for Administrator, click Apply, and then restart the OpenSSH service. To change the other default users' passwords, log on as each one and follow the same procedure as for Administrator. Users other than Administrator do not have passwords associated with the OpenSSH service.
- iLO Admin user (each CLIM, processor blade and NSC): use OSM Service Connection to select each CLIM , then select the "Launch iLO" action to log in and change the CLIM iLO Admin user password. See “Change the CLIM Passwords” in CIPMGMT for details. For blade iLOs, use the Launch iLO URL action to log in and change the password. For NSCs, you need to obtain the iLO IP address and connect with the iLO through the browser.
- Onboard Administrator (OA) Administrator user (each c7000 enclosure) and iLO Admin user (each blade): use OSM Service Connection to select each enclosure, then select the "Launch the Onboard Administrator" action to log in and change the OA password. While logged into each enclosure's OA, select each CPU blade, then select iLO, then select the Web Administration link and change that blade's iLO Admin user password.
- Internal InfiniBand switch Admin user: use OSM Service Connection to select each switch, run the "Launch Switch Management Console" action, then log in using the default password and change the switch's Admin user password.

- External InfiniBand switch Admin user: use OSM Service Connection to select each switch, run the "Launch Switch URL" action, then log in using the default password and change the switch's Admin user password.

See SAFEREF, CIPMGMT Chapter 2, “Change the CLIM Passwords”, OSMCUG and OSMSCG for details.

Caution: If you do change the passwords on underlying components, ensure that you have a reliable process for communicating the new passwords to your HPE service providers whenever required for maintenance or repair.

Software installation using DSM/SCM and Software Essentials

One of the early tasks in installing a new system is configuring the software installation management subsystem, DSM/SCM. The default DSM/SCM user ID is SUPER.SUPER, but you may assign that role to a different user. The DSM/SCM user is the initial database user and is responsible for specifying which users are allowed to use DSM/SCM and what tasks they are permitted to perform. DSM/SCM has several user roles: planner, operator, and database administrator. Make sure that you have assigned appropriate users to each role.

Note: Certain DSM/SCM programs run as SUPER.SUPER; if run as SUPER.SUPER, INITENV sets their ownership accordingly and sets PROGID. Otherwise, this must be done manually during DSM/SCM installation.

There are additional DSM/SCM security parameters, including:

- If Safeguard is required to be running whenever DSM/SCM places files (should be checked/yes)
- Ownership for new files (the default is the DSM/SCM user)
- Security for new files (the default is NCNC; tighten as needed to match your policy)

Note: Files installed by DSM/SCM that replace existing files assume the existing files’ ownership and security.

The activation of a software revision requires that the user have adequate privileges perform the necessary functions, which may include:

- Stopping applications
- Stopping subsystems and/or devices
- Updating firmware
- Running ZPHIRNM (requires purge access to files being replaced)
- Performing a coldload (system load)

On the H and J series, it might be necessary to start an iTP WebServer instance. Its setup should be performed by a SUPER group user, but not SUPER.SUPER. See DSMSCMUSER Section 4, task 3 for descriptions of roles and the associated permissions that may be assigned, Section 5, Setting DSM/SCM Security for details on security-related parameters, Sections 9-12 on the actual installation process, Section 20, Installation of Web Interfaces, and Appendix E, ASSIGNs and PARAMs, for details on various security-related settings. For NonStop Software Essentials, see NSEUSER Section 3, task 3 for descriptions of roles and the associated permissions that may be assigned, Section 4, Setting User Security for details on security-related parameters, Section 5, Configuring OSM for enabling SSL if not enabled by default, and Appendix E, ASSIGNS and PARAMs, for details on various security-related settings.

Development and test systems

Development and test systems should be set up and managed with no less care than your production systems. In particular, developers and testers should not have routine access to SUPER.SUPER and privileged activity should be audited and monitored.There are a few additional considerations:

Developers and testers also normally should not have remote access to production systems. Do not add their user IDs and remote passwords to production systems unless there is a compelling need, and do not share passwords between development/test and production systems.

Other

Create and appropriately secure xxLOCL files for TACL, FUP, Safeguard etc.

NonStop user management: best practices

User management can be performed in both the Guardian and OSS environments. The OSS user management tools have a subset of the features available in Safeguard.

User creation

When you create a new user through Safeguard, be sure to:

- Set the user's primary group (PRIMARY-GROUP attribute) if desired. By default, the user's primary group is its administrative group. If the user is to belong to one or more file-sharing groups, you should consider changing its primary group to one of the file-sharing groups.
- Set adequately restrictive GUARDIAN DEFAULT SECURITY, preferably "OOOO" (local owner only), and an appropriate GUARDIAN DEFAULT VOLUME.
- If your protection model is file-oriented, consider using DEFAULT-PROTECTION to establish appropriate default protection attributes for files created by the user. You might wish to instead use a model that is oriented toward subvolume ACLs, to minimize the proliferation of file-level ACLs. A third option is to primarily rely on XAC or a partner product for tailored access control. If you choose the third model, consider adding basic Safeguard access control with volume, subvolume and/or disk-pattern ACLs as fallback protection.
- Force the new user to change passwords at first logon by setting the password as expired with a grace period (PASSWORD-EXPIRES, PASSWORD-EXPIRY-GRACE).
- Force password changes through PASSWORD-MUST-CHANGE in accordance with your corporate policy.
- If desired, restrict how often a user can change passwords through PASSWORD-MAY-CHANGE.
- If desired and using Safeguard-protected terminals, supply non-default values for the user's start-up command interpreter.
- Configure the user-level AUDIT attributes according to your enterprise's policy.
- Use the TEXT-DESCRIPTION string to capture information about the user (name, contact information, etc.)
- If you know that the user will need access to the system for only a limited time period, set USER-EXPIRES appropriately.
- If you are adding a functional user (e.g., SUPER.SPOOLMGR) and multiple individuals will need to act in that role, either create a separate new alias for each individual or (preferably) use controlled "logging up" through an access management tool such as XAC or a partner product for accountability. See below for more information on functional users.

If the user is to have OSS access:

- Assign an initial working directory (INITIAL-DIRECTORY). HPE recommends creating a separate fileset for initial working directories.
- Assign an initial program (INITIAL-PROGRAM) – typically, the OSS shell (/bin/sh).

You may assign one or more owners for an individual user through OWNER and OWNER-LIST. Owners have full control over the user's authentication record, including the ability to change passwords and delete the user. The default is that the user who created the new user is the primary owner.

By default, new users do not have access to other systems through Expand. If inter-system Expand access is required through either Guardian or OSS, use the REMOTEPASSWORD attribute to configure a matching remote password on each system where access is to be allowed.

Note: REMOTEPASSWORD is unrelated to the password the user provides for authentication; you should think of it as a persistent token enabling inter-system access by that user. Once a remote password is set, it stays in force until explicitly altered or removed.

See SECMGMT Section 2, SAFEADMIN Sections 2, 3, 5 and 6, and OSS MOG Section 8 for details.

A few individual subsystems, including Safeguard, newer versions of SQL/MX and SQL/MX connectivity, Software Essentials and DSM/SCM, have the ability to configure specific roles for individual users. If the new user is to have any of these roles, configure them appropriately. See the sections on individual subsystems for details.

Role-Based Access Control (RBAC)

Some forms of RBAC are inherent in the user capabilities described below; others can be implemented through various forms of access control including Guardian, OSS and OSM ACLs, XYGATE Access Control (XAC), Guardian and OSS Security Event Exit Processes (SEEPs), Safeguard security groups and SQL/MX security administrators.

SUPER.SUPER

There are multiple approaches for managing SUPER.SUPER in both the Guardian and OSS environments, so you will need to select the configuration options that best match your security policy. Do not use SUPER.SUPER except for specific operations that require those privileges; in particular, system administrators should not do their routine work while logged on under that ID. Maintain a list of procedures that require SUPER.SUPER access, such as licensing files, and train your administrators to not use it for anything else. Most day-to-day system administration activities can be performed as other SUPER group members, and most activities that do require SUPER.SUPER access, such as backing up the system, can be performed through a separate copy of the appropriate utility that belongs to SUPER.SUPER, has PROGID set, and has execute access restricted through Safeguard ACLs, XAC, or a partner authorization control product with appropriate capabilities.

Note: You will see references to “the super ID” in NonStop manuals, which usually includes both the SUPER.SUPER user ID and any aliases. In this version of this document, references to SUPER.SUPER generally include its aliases unless it is either otherwise noted or obvious from the context.

At an absolute minimum, physically secure access to the SUPER.SUPER password and log all requests for it (who, when, for what purpose). Consider having two separate people each enter part of the password when creating it and sealing their parts in separate envelopes, so nobody knows the entire password. Also consider changing it after each use. HPE much prefers that you also adopt the recommendations made below rather than just relying on physical password security.

User configuration

The SUPER.SUPER user should be owned by the Safeguard security manager. If it has an explicit owner, you have the option of keeping it in the FROZEN state by default and thawed only when needed; otherwise, do not freeze the ID.

Note: The owner of SUPER.SUPER is a sensitive ID in its own right and needs careful management.

There are multiple, conflicting schools of thought on SUPER.SUPER password management, with different advantages and disadvantages. Here are several:

- The SUPER.SUPER password should be configured to never expire, especially if SUPER.SUPER is not owned by a different user. The advantage is that you will never be confronted with an emergency only to discover that you can’t log on because the password is expired.

- The SUPER.SUPER password should be changed regularly.
- The SUPER.SUPER password should be changed after every use.

If you need to restrict which users are allowed to log on as SUPER.SUPER, and from which terminals and IP addresses, configure XUA accordingly. Partner user authentication products might have similar capabilities.

If you are not using XUA or a partner product that always requires logging up to SUPER.SUPER from another user ID rather than directly logging on as SUPER.SUPER, then consider creating a separate alias to SUPER.SUPER for each user who will be allowed to use it. Use of an alias increases accountability because it will show up in place of SUPER.SUPER in Safeguard audit. However, some activity in other subsystems such as SCF and FUP will be logged under SUPER.SUPER, and unless explicitly controlled all of the aliases for a user ID have the ability to modify that user's files such as the xxxxCSTM files for TACL, etc. If you do create aliases, consider freezing the underlying SUPER.SUPER user ID.

Note: Since the aliases have the power of SUPER.SUPER, they will need the same level of controls, including password management.

Note: While the UNIX super ID has a UNIX UID of 0, the OSS super user ID (UID) is 65535. See Relationship between Guardian and OSS user and group IDs, below, for the explanation. Do not delete the SUPER.SUPER user ID.

Some customers have renamed SUPER.SUPER, but that approach could have side effects and requires caution. One potential side effect is an application that attempts to connect as SUPER.SUPER would not be able to authenticate itself.

Restricting privileges

Logging down

By default, SUPER.SUPER can log down to any other user without supplying a password, and can successfully log on as another user even if that user ID is frozen. If your policy does not allow this behavior, use SAFECOM to set PASSWORD-REQUIRED = ON. Consider how you would handle emergency scenarios if you do not allow SUPER.SUPER to log down without a password.

Accessing Guardian files

On a newly-delivered system, SUPER.SUPER is configured as DENIABLE. This setting allows SUPER.SUPER file access to Guardian files to be prohibited through use of DENY in Safeguard ACLs and through Guardian authorization SEEP rulings.

If your policy allows SUPER.SUPER to have unfettered access to all Guardian files, you might wish to reconfigure SUPER.SUPER as UNDENIABLE so it will have full system access in an emergency. The change to UNDENIABLE access has to be made during system generation rather than online through SAFECOM. To do so, add the following line to the ALLPROCESSORS paragraph of the CONFTEXT file prior to DSM/SCM system generation:

```
SUPER_SUPER_IS_UNDENIABLE;
```

The change will take effect when the system is loaded with the resulting OS image, and applies to SUPER.SUPER and all of its aliases.

The SAFECOM INFO SAFEGUARD command will display the following header line if the change is in effect:

```
SAFEGUARD IS CONFIGURED WITH SUPER.SUPER UNDENIABLE
```

See SAFEADMIN Section 10, Installing the Safeguard Software and SAFEREF (INFO SAFEGUARD).

Accessing OSS files

As of H06.22 / J06.11, OSS supports functionality comparable to the Guardian DENIABLE setting. Behavior prior to those RVUs is comparable to UNDENIABLE in the Guardian environment. The

Guardian DENIABLE/UNDENIABLE setting has no effect on OSS file access rulings. To maintain compatibility with prior RVUs, the default is to not restrict SUPER.SUPER (65535) access. Restricted access can be enabled at the fileset level through SCF. If you do enable restricted access, there are special considerations for backups and other special forms of file access; you will need to populate the new Safeguard SECURITY-PRV-ADMINISTRATOR (SPA) group as well as the SECURITY-OSS-ADMINISTRATOR (SOA) group. To avoid subverting the new controls, SUPER.SUPER should not own either group, and you should configure the ACLs on the two groups to explicitly deny SUPER.SUPER any access.

See OSSMOG (chapter 5, Managing filesets, Chapter 9 Managing security and chapter 13, OSS Monitor SCF Command Reference Information) and SAFEADMIN Section 6 for details.

Subsystem commands and subcommands
If you need to provide command-level and subcommand-level controls on SUPER.SUPER, configure XAC accordingly or use a partner product with those capabilities.

Audit

For the most complete audit, enable keystroke logging for SUPER.SUPER through XAC or a partner access control product.

Enable auditing in both user authentication and access control products (XUA, XAC and/or partner products).

Enable authentication audit, both pass and fail, through SAFECOM.

Enable priv logon audit through SAFECOM.

See SAFEADMIN Section 9, Configuring Safeguard Auditing, XUAREF and XACREF for details.

Functional users

SUPER.SUPER is an example of a functional user. A functional user is associated with a specific role that may need to be performed by more than one individual user. As with SUPER.SUPER, you need to provide accountability and possibly command-level controls at the individual level, and most of the techniques described above apply here as well, including assigning each individual in that role an alias and freezing the underlying user ID. An example of a functional role might be a SUPER.SPOOL user that is allowed full control over the Spooler subsystem instances. If you are not using aliases for functional roles, consider giving them less guessable names – for example, don't use SUPER.OPER or SUPER.OPERATOR.

Other SUPER group users ("SUPER.notsuper")

The original system administration model, which still persists in many subsystems, allows users who do not belong to the SUPER group to retrieve most system-related information, but not perform any sensitive operations that might affect the system's configuration or operating environment, In these subsystems, members of the SUPER group other than SUPER.SUPER (collectively referred to in this document as SUPER.notsuper) generally have the ability to perform "day-to-day" sensitive operations but not the most sensitive ones such as licensing program files containing privileged code. The place where this model is most prevalent is in the subsystems managed by SCF. If it better fits your security policy, you can slightly alter this model for tape management (DSM/Tape Catalog) through configuration of the Safeguard SECURITY-MEDIA-ADMIN security group and, for persistent process management ($ZPM), through configuration of the Safeguard SECURITY-PERSISTENCE-ADMIN security group. Consider freezing SUPER.notsuper users who are absent on vacation or leave. You also can place temporary access controls on individual users through XUA or a partner product.

Some versions of UNIX support a wheel group, with similar privileges to SUPER.notsuper users. OSS does not support this model.

See SAFEADMIN, Section 6, Managing Security Groups, MEDIACOM, Section 2, Managing Tape File Entries, and SCFKERNEL, Section 6, Sensitive and Nonsensitive Commands, for details.

HPE and partner security software users

If you are using HPE security products such as XUA or SSH or partner security software products, apply the same care in managing users associated with installing, configuring and managing those products as you do for SUPER.SUPER.

Administrative group managers

By default, group managers (*, 255) are able to manage members of their group, including adding them, logging on as them without knowing their passwords, etc. Use SAFECOM to set PASSWORD-REQUIRED=YES to prevent group managers from logging down to an individual's ID without supplying a password.

For more sophisticated control, consider using the Safeguard OBJECTTYPE USER to specify what users can manage users, aliases, and file-sharing groups.

Note: Users with Safeguard OBJECTTYPE USER CREATE privileges can create SUPER group users and group managers. Also, if you use OBJECTTYPE USER, group managers lose their automatic privileges. You can restore their privileges by adding them to the OBJECTTYPE USER ACL with CREATE authority, but that gives them full user management privileges – not just privileges for managing their own groups.

Consider freezing group manager IDs if they are not required for day-to-day use.

HPE recommends that group manager IDs be considered as functional IDs, and that users that may perform functions in that role also have their own individual user IDs for day-to-day work.

OSS does not have a group manager concept.

Individual users

If you need or wish to tailor authentication restrictions beyond what is available through Safeguard, use XUA or a partner authentication product with similar capabilities. For access control, use XAC or a partner authorization product with similar capabilities.

User impersonation

People

XAC can be configured to allow a specific user to perform controlled actions as a different user, including SUPER.SUPER, without providing that user's password. Partner products offer similar capabilities. See XACREF for details.

Programs

There are legitimate cases where a server process needs to be able to act on behalf of users other than the one that started it, including having all of that user's capabilities. The Safeguard PRIV-LOGON disk file attribute allows a program file to be sanctioned to log on as a different user without providing that user's password. The AUDIT-PRIV-LOGON disk file attribute can be enabled to trigger audit when a process with PRIV-LOGON enabled invokes a priv logon.

Caution: There can be a performance impact if a process invokes PRIV-LOGON frequently, as the disk file attribute setting is checked on every invocation and, if AUDIT-PRIV-LOGON is set, audit is generated.

See SAFEREF Section 7, ADD DISKFILE or ALTER DISKFILE, for details.

Relationship between Guardian and OSS user and group IDs

OSS UIDs are the scalar equivalents of Guardian user IDs; for example, the super ID UID, 65535, is (255*256) + 255. OSS group IDs (GIDs) 0 through 255 are the same as Guardian administrative group IDs 0 through 255. OSS group IDs in the range of 256 – 65535 are file-sharing groups.

File-sharing groups

File-sharing groups have group IDs of 256 and higher (0-255 are used for administrative groups). When a new user is created, its administrative group is its primary group.

The maximum user-specifiable file-sharing group is 65535; Safeguard has reserved IDs above that limit for security-related file-sharing groups. File-sharing group numbers above 65535 cannot be used in Safeguard ACLs.

Group number 65536 is reserved for a file-sharing group called the SECURITY-ENCRYPTION-ADMIN group. Its members are allowed to manage Volume-Level Encryption (VLE) disk and tape encryption attributes. Only SUPER group members may be added to this group.

See VLEGUIDE for more details on the SECURITY-ENCRYPTION-ADMIN group.

An individual user also can be a member of one or more secondary groups (supplementary groups, in POSIX terminology). Extensive use of file-sharing groups can be difficult to manage, so consider whether you can meet at least most of your needs through some combination of administrative groups and XAC or partner access control products.

Security groups

Safeguard supports the definition of a number of security groups. You must add these groups explicitly to your Safeguard configuration and then assign users to them. The complete list of groups includes:

- SECURITY-ADMINISTRATOR
- SYSTEM-OPERATOR
- SECURITY-AUDITOR
- SECURITY-OSS-ADMINISTRATOR (“SOA”)
- SECURITY-PRV-ADMINISTRATOR (“SPA”)
- SECURITY-ENCRYPTION-ADMIN
- SECURITY-MEDIA-ADMIN
- SECURITY-PERSISTENCE-ADMIN

Members of the SECURITY-ADMINISTRATOR group are responsible for Safeguard administration. Members of SYSTEM-OPERATOR can manage day-to-day Safeguard operations such as rotating audit files but cannot change its configuration. Depending on the size and nature of your security administration and operations staff, you might or might not need to make use of the SYSTEM-OPERATOR group.

HPE recommends that you configure membership in all of the Safeguard security groups rather than relying on default settings.

The SECURITY-AUDITOR group allows your auditors to be more self-sufficient. SECURITY-AUDITOR group members have read-only access to Safeguard subject and group records.

See SECMGMT, Section 6, Managing Security Groups, for additional details on the first five groups, DSMSCMTAPE, Section 2, MEDIACOM Commands, for SECURITY-MEDIA-ADMIN, and SCFKERNEL for SECURITY-PERSISTENCE-ADMIN.

User deletion

You should identify and remove any explicit privileges associated with a user ID or alias prior to deletion. For example:

- Give ownership of Guardian files to other users as appropriate (use FUP INFO or DSAP USER, ANALYSIS, DETAIL across all volumes to identify them). Examples:

  ```
  TACL> DSAP *, USER 245,111, DETAIL, TERAFORM
  TACL> FUP INFO $*.*.* WHERE OWNER = 245,111
  ```

- Give ownership of OSS files to other users as appropriate (use find <directory> -user <username> across all directories to identify them). Examples:

  – Step 1: Find out what the OSS user-id is:

  ```
  TACL> SAFECOM INFO USER <name> or <number>, GENERAL
  ```

  The output shows:

  ```
  GROUP.USER         USER-ID    OWNER     LAST-MODIFIED    LAST-
  LOGON     STATUS
  GROUP45.BERT         245,111    255,255   27JAN09,  0:04 25NOV13,
  16:17  THAWED
    UID                        =      62831
  ...
  ```

  – Step 2: Use the OSS `find` command to find what is owned by that user:

  ```
  OSS> find / -W NOG -W NOE -user 65391
  ```

  If you forget to do this before deleting the user, you can check later for unowned files:

  ```
  OSS> find / -W NOG -W NOE \(-nouser –o –nogroup\)
  ```

  – Step 3: Give the files and directories owned by that user to an appropriate user:

  ```
  OSS> chown <newuser>[:NEWGROUP] command.
  ```

  Note: Changing ownership requires SUPER.SUPER privileges.

- Give ownership of SQL tables, schemas, catalogs etc. to other users, using the Safeguard user ID in OSS format (groupID*256 + userID)

  For SQL/MX Catalogs;

  To identify the catalogs owned by the user:

  ```
  OSS> mxci
  >> SET SCHEMA nonstop_sqlmx_<system name>.system_schema;
  >> SELECT catowner, rtrim(c.cat_name) AS cat_name
   FROM catsys c
   WHERE cat_owner = 62831
   FOR read uncommitted access;
  ```

  To give the catalog away:

  ```
  OSS> mxci
  >> GIVE CATALOG catname TO authid
  ```

  Example: `give catalog usercat1 to "sql.user5"`;

  For SQL/MX schemas:

To identify the schemas owned by the user:

```
  OSS> mxci
  >> SET SCHEMA nonstop_sqlmx_<system name>.system_schema;
>> SELECT schema owner, rtrim(c.cat_name) || '.' ||
   rtrim(schema_name) as schema_name
     FROM schemata s, catsys c
     WHERE schema_owner = 62831
     AND s.cat_uid = c.cat_uid
     FOR read uncommitted access;
```

To give the schemas away:

```
OSS> mxci
>> GIVE SCHEMA schema TO authid [CASCADE]
```

Example: `mxci give schema usercat.userschema1 to "sql.user5" cascade`;

With the CASCADE option all objects below the schema will be given to the new user. If CASCADE is not used, you'll need the following steps to identify objects within a schema belonging to the user and give them to another user.

For SQL/MX objects within a schema:

To identify the objects belonging to the user:

```
OSS> mxci
>> SET <user catalog>.DEFINITION_SCHEMA_VERSION_<schema version>;
>> SELECT object_owner, object_type, rtrim(object_name) AS
 object_name
   FROM objects
   WHERE object_owner = 62831
   FOR read uncommitted access;
```

Where `<schema version>` is:

`3400` for SQL/MX R3.4
`3200` for SQL/MX R3.3, R3.2 and R3.2.1
`3100` for SQL/MX R3.1
`3000` for SQL/MX R3.0
`1200` for prior releases

Give away ownership of the individual objects:

```
GIVE object-type object TO authid
```

Where *`object-type`* is one of `{TABLE | TRIGGER | VIEW | PROCEDURE SEQUENCE}`.

- Remove the user from the SQL/MX security administration group, if present. See SQLMXMGT, Section 2, REVOKE SECURITY_ADMIN, for details.
- Remove the user from the ODBC/MX operator group, if present. See MXDMUSER Section 7, Editing MXCS user permissions for details.
- Beginning with SQL/MX 3.3, you can avoid creating orphaned objects by specifying CHECK-SQLMX-OWNERSHIP when using SAFECOM to delete a user. This option causes the deletion to fail if the user still owns any objects.
- Remove the user from security groups and file-sharing groups.
- Remove the user from the SSH database.
- Remove the user from the XUA database.

- Remove the user from partner product and application-specific databases.
- Remove the user from Safeguard and OSS ACLs.
- Use Software Essentials, the DSM/SCM Maintenance Interface or the DSM/SCM Planner Interface to check whether the user has any DSM/SCM roles and delete the user from those roles if found.
- Reassign ownership of any scheduled NetBatch jobs.
- Reassign ownership of any persistent processes under $ZZKRN.
- Remove Spooler jobs.
- Make sure that there are no processes currently running under that user ID.

You also might wish to first freeze the ID temporarily to help detect any remaining use by applications.

If the user has one or more aliases, be sure you delete them all. Also ensure that the user/alias has been deleted from all systems where it existed.

Because there are only 256 group IDs and 256 user IDs within each group, it is possible to subsequently add a new user with the same group and user IDs and either the same user name as or a different user name than the user being deleted. Do not reuse a user ID unless you are sure that the level of clean-up described above has been performed; in most cases, user privileges are associated with the user ID (underlying user ID in the case of aliases), so just changing the user name is not adequate protection against reuse issues.

See SAFEADMIN Section 2, Deleting Users, SECMGMT Section 6, Removing a User from the System, DSMSCMUSER Section 4, task 3 for descriptions of roles and the associated permissions that may be assigned, Section 5, Setting DSM/SCM Security for details on security-related parameters, NSEUSER Section 3, task 3 for descriptions of roles and the associated permissions that may be assigned, Section 4, Setting User Security for details on security-related parameters, OSSMOG, SQLMPMGMT, SQLMXMGMT, and NETBATCH Section 6, Commands.

Partner tools are available to help identify orphaned objects.

Guardian file security

Guardian security vectors: background

Guardian uses an RWEP (read, write, execute, purge) vector to control file access, where each of the four positions can be set to one of:

A = all locally-authenticated users

N = local/remote authenticated users

G = all locally-authenticated admin group members

C = local/remote authenticated admin group members

O = locally-authenticated owner

U = local/remote authenticated owner

- = locally-authenticated SUPER.SUPER

Remote refers to a user who has been authenticated on another system within the Expand network of NonStop systems. Note that unlike owner, group, and all, there is no local/remote security vector value for SUPER.SUPER.

A user can access a remote file only if he/she has both the authority to access the file and matching user names, IDs and remote passwords exist on both systems. The user names and IDs are created by the two systems' security administrators; either the security administrator or the user may set the remote passwords. Users do not specify the remote passwords at the time that they attempt to access a remote resource; system software checks whether the configured remote passwords match.

Note: If you create a file on a remote system, its default security vector will be based on your local default security vector but with any A, G or O settings remapped to N, C, or U so you retain access to the file from your local system.

See GUARDUSE, Section 16, Disk-File Security and Network Security, for details.

You can give away ownership of your files either through the FUP GIVE command or programmatically through the SETMODE procedure, subject to certain restrictions for files with Safeguard disk file authorization records. SUPER.SUPER and its aliases can give away ownership of any user's files.

Note: It is possible to give ownership of a file to a nonexistent user.

See GUARDPROCREF Section 14, SETMODE, for details.

Guardian security vectors: best practices

Each user's default Guardian file security should be set up for least privilege (OOOO or UUUU) unless your organization's security policy is less restrictive. You can check user default security using either TACL USERS, SAFECOM INFO USER, or the USER_GETINFO_ API, and can check alias default security using SAFECOM INFO ALIAS.

Guardian security vectors have relatively coarse granularity. HPE recommends use of Safeguard ACLs for finer-grained control.

PROGID and LICENSE attributes: background

A Guardian program file's security can be configured through the PROGID attribute to cause the program to run as its owner even if it is run by another user.

PROGID sets the process' Process Access ID (PAID), effective user ID, saved-set-user-ID, effective group ID, and saved-set-group-ID to the program file owner's ID. Real user ID remains set to the Creator Access ID (CAID) and real group ID remains set as the creator's primary group ID. The group list remains set to the creator's group list. Access decisions are based on the PAID, effective user ID, and group lists.

The PROGID attribute allows the running program to access files on behalf of the user that runs it but based on the privileges of the program file owner, including files that the user running the program might not otherwise have access to. The program may choose to limit access to data, e.g. displaying only a subset of its contents depending on the user that invoked it. Use of PROGID also prevents the running program from accessing files belonging to the user running the program unless the program file's owner also has access to them.

Note: If you remotely launch a program with PROGID set, you cannot stop it remotely as the resulting process is considered to be locally authenticated. Users authenticated as either its CAID or PAID can stop it locally.

Programs that contain one or more procedures designated as CALLABLE or PRIV are presumed to be capable of performing system-privileged functions. By default, such programs are runnable only by SUPER.SUPER. If they must be runnable by other users as well, they must have the LICENSE attribute set ("be licensed").

PROGID and LICENSE can be set either online through FUP or SAFECOM or programmatically through the SETMODE API. PROGID can be set by either the program file's owner or SUPER.SUPER. LICENSE can be set only by SUPER.SUPER.

LICENSE can be revoked explicitly through FUP REVOKE (non-Safeguard protected file) or SAFECOM ALTER (Safeguard-protected file). It also is implicitly revoked whenever a licensed file is opened with write access.

See GUARDUSE, Section 16, Adopting the Owner ID of a Program File and Licensing Programs, FUPREF Section 2, LICENSE and REVOKE, SAFEREF, Section 7, ADD DISKFILE Command, and GUARDPROGREF, Section 14, SETMODE for details.

PROGID and LICENSE attributes: best practices

All usage of PROGID and LICENSE should be monitored. HPE publishes a list of the programs that it ships with either attribute enabled. See SECMGMT section 2, Sanitizing a NonStop System, and Appendix D, HPE-Supplied Files with LICENSE and PROGID.

It is a good practice to secure program files so only the owner has write access. It is especially important to restrict write access for files with either PROGID or LICENSE set, so other users who have execute access cannot attach a malicious runtime library.

You should either use Safeguard ACLs to tightly restrict access to licensed utilities that have high destructive potential such as DIVER, SNOOP, GOAWAY, and SQLCI2 or reset the LICENSE attribute entirely so only SUPER.SUPER can run them. Similarly, HPE recommends that you do not license TANDUMP, FILCHECK or DIRCHECK.

Place program files under Safeguard protection, and use SAFECOM to license them or set PROGID where required. Examples:

```
> SAFECOM ALTER DISKFILE PRIVPROG, LICENSE ON
```

```
> SAFECOM ALTER DISKFILE APPSRVR, PROGID ON
```

Note: The FUP LICENSE, SECURE and REVOKE commands also can be used to manage PROGID and LICENSE.

Note: A few subsystems either repurpose or overload the meaning of LICENSE. NetBatch uses the LICENSE flag as an indicator of SUPER.SUPER approval of the existence or contents of a different type of file, such as a data file. SQL/MP uses licensing SQLCI2 to indicate that it may perform privileged operations; if it is not licensed, it cannot perform those operations even if run by SUPER.SUPER.

Safeguard authorization records and ACLs: background

When you secure an object with the SAFECOM ADD command, Safeguard creates an authorization (protection) record for that object. The authorization record contains several security attributes, including the ACCESS attribute, which is used to define an access control list (ACL).

Safeguard ACLs

An access control list specifies access authorities associated with a particular object (such as a disk file). Access control lists allow you to specify access to a greater level of detail than Guardian security strings allow. For example, with an access control list, you can grant access to one or two members of a group without having to grant access to the entire group.

You cannot specify aliases in an ACL, so authorization decisions for a user logged in as an alias are made on the basis of the underlying user ID.

You can specify authorities for an individual user, a user group, or all users. Consequently, an individual user’s authorities can be determined by more than one entry. For example, one entry can grant EXECUTE authority to an entire group, while another entry can grant READ and WRITE authority to an individual member of the group.

Some actions require more than one authority in order to be allowed. When creating ACLs, you need to understand what combinations of authorities are required. You must specify all authorities required for a given action in one ACL entry. For example, a user needs both READ and WRITE authorities to edit a disk file. If only READ authority is granted to every member of a group, and only WRITE authority is granted to an individual user in the group, the user cannot edit the file because READ and WRITE authorities do not appear in the same entry.

If you add authorities for an individual user, a new entry is not created. The existing entry for that user is updated.

Use DENY to explicitly deny a user certain authorities. DENY is useful when you want to deny access to a few members of a group while granting access to the remainder of the group. Also, you can use DENY to deny access to the super ID. Normally, the super ID has all access authorities unless explicitly denied.

A denial always takes precedence over a grant. For example, if a user is granted WRITE authority in one entry and denied WRITE authority in another entry, the user is denied WRITE authority.

See SECMGMT, Section 3, Security Objects, and SAFEADMIN for additional details.

Safeguard has an OBJECTTYPE protection record for each type of object that it can secure, and a top-level OBJECTTYPE OBJECTTYPE protection record that controls the creation of new OBJECTTYPE protection records.

You can configure ACLs at multiple levels for certain object types;

DEVICE: device, subdevice

DISKFILE: volume, subvolume, file

PROCESS: process, subprocess

For these object types, the same user may appear, possibly with conflicting authorities, in ACLs at different levels for the same object. Safeguard provides controls on the order of evaluation for each of these objects; their attribute names begin with CHECK- (what to check), DIRECTION- (evaluation order) and COMBINATION- (conflict resolution).

Diskfile patterns and saved diskfile patterns (which can be wildcarded at the volume level and applied to new volumes as added) can be used to simplify ACL management by supplying a single pattern that matches many subvolumes or file names. If you use diskfile patterns, you'll need to use CHECK-DISKFILE-PATTERN to specify whether patterns are searched before doing the normal protection record search, after the normal search if it results in NORECORD, or are the only search performed.

Caution: If you use diskfile patterns or saved diskfile patterns, you need to use the ADD OBJECTTYPE DISKFILE-PATTERN or ADD OBJECTTYPE SAVED-DISKFILE-PATTERN to specify which users can control diskfile patterns and saved diskfile patterns. The default is that any user can add a diskfile pattern, which allows anyone to effectively control file access across an entire volume.

Caution: If you set CHECK-SUBVOLUME ON and set DIRECTION-DISKFILE to VOLUME-FIRST or FILENAME-FIRST, any user can gain access to someone else's files. All files that are in subvolumes that have not been added to the Safeguard database are vulnerable. This situation occurs because any user can add the subvolume to the database and thereby own it. If this configuration is needed, use the ADD OBJECTTYPE or ALTER OBJECTTYPE command to specify who can control subvolumes.

See SAFEREF Section 1, Object-Access Authorization, for an overview. See SAFEADMIN Section 5, OBJECTTYPE Control, and Section 9, Configuring Device Control, Configuring Process Control and Configuring Disk-File Control and SAFEREF for details.

Safeguard supports warning modes that let you test whether your diskfile and process ACLs are configured properly. It can be configured through WARNING-FALLBACK-SECURITY to either grant access or fall back to Guardian security if access would have been denied based on the ACL. You can enable warning mode on either a system-wide (SYSTEM-WARNING-MODE) or selective (OBJECT-WARNING-MODE) basis.

See SAFEADMIN, Section 8, Warning Mode, for details.

Safeguard authorization records and ACLs: best practices

Do not grant a user more access authorities than needed. In particular, be careful when using the asterisk (*) to grant all valid authorities to a user.

Do not grant specific authorities to more users than needed. In particular, if you are installing Safeguard ACLs for the first time, you might begin by mapping the file's existing Guardian security vector into an equivalent ACL. At this point you might consider whether everyone covered by G, C, A or N really needs that authority and adjust the ACL accordingly, including restricting the set of nodes from which a user can have remote access to a protected object. In particular, users on development nodes should not have routine remote access to objects on production nodes.

You can simplify Safeguard management and ensure that you can get a unified view of how objects are protected by having a single security management user own all protection records.

By default, diskfile ACLs are enabled at the individual filename level but not for volumes or subvolumes. If you either are just beginning to use ACLs or are using XAC or a partner access control product, consider enabling volume and subvolume ACLs and creating appropriate ones to provide at least coarse-grained protection. Volume and subvolume ACLs also can support separation of duties by, for example, effectively allocating different disk volumes to different applications and restricting file placement on specific volumes such as $SYSTEM and disk audit volumes.

ACL-REQUIRED-DISKFILE, if set to ON, causes denial of access to the file rather than applying Guardian security checks if there is no protection record that includes the file in its scope. HPE recommends that you consider leaving the attribute set to OFF (default). In Safeguard implementations where retaining Safeguard protection of all files is a high priority, you might utilize this option to prevent the Safeguard implementation from growing obsolete over time.

Caution: If you do set ACL-REQUIRED-DISKFILE ON, you must immediately configure an access control list for the SAFECOM program file that grants execute authority to you. Otherwise, when you complete your current session, you cannot control Safeguard through the SAFECOM command interpreter because you cannot run SAFECOM.

Use diskfile patterns to cut down on ACL proliferation.

You can get very different protection for individual objects, depending on how you configure the ordering of ACL evaluation. In order to retain the ability to set specific security on an individual file and not have it be overridden by a more general ACL, use the FILE FIRST configuration. If you use diskfile patterns, a relatively straightforward configuration through SAFECOM would be:

```
DIRECTION-DISKFILE FILENAME-FIRST -- Start the search with the individual file's ACL, if present

COMBINATION-DISKFILE FIRST ACL    -- Use the first ACL found, whether it contains the user's ID or not

CHECK-FILENAME ON          -- Check disk file ACLs (default)

CHECK-SUBVOLUME ON         -- Check subvolume ACLs

CHECK-VOLUME OFF           -- Do not check volume ACLs (default)

CHECK-DISKFILE-PATTERN LAST      -- Only if there is no ruling from the normal ACL search
```

HPE recommends that you enable ALLOW-NODE-ID-ACL. It defaults to OFF, which gives you only local/remote granularity rather than letting you specify users on specific Expand nodes.

Note: Putting a remote user in an ACL does not obviate the need for that remote user to have matching remote passwords on both systems.

HPE recommends that you take advantage of warning mode when introducing ACLs on a system, especially if you are enabling any of the ACL-REQUIRED attributes. Using warning mode helps you catch both too-loose settings that would allow unexpected access and too-tight settings that could severely hamper application or system operation. Do not leave warning mode enabled any longer than necessary.

See SAFEREF Sections 7-16 and Appendix B for additional information on ACLs.

Safeguard authorization SEEPs

Safeguard supports configuration of a Security Event Exit Process (SEEP) that participates in authorization decisions for Guardian objects. You should consider using a SEEP if you want to provide access controls beyond what is available through Safeguard. Use of an authorization SEEP also has the potential to reduce Safeguard ACL proliferation.

You can purchase a partner authorization SEEP or, if you wish, write your own. Make sure that you understand how SEEPs interact with the rest of the system before deciding to include one in your environment. If you write your own (not recommended due to complexity and the availability of several partner SEEPs), you also will need to provide a user interface to configure its rules. See SAFEREF Section 15, Event-Exit-Process Commands, for details on configuring authorization SEEPs. That section also includes design considerations and message format details that you will need if you do decide to write your own SEEP.

OSS file security

OSS file and directory permissions: background

Users and permissions

Access to files and directories is controlled in three ways:

r (read: allows users to view or print the file, for directories to read the names of files)

w (write: allows users to modify the file, for directories to create, delete and rename files)

x (execute: allows users to run the file or to search directories, for directories to cd to that directory and open files in that directory)

Users on the system are classified as one or more of the following:

u (user/owner)

g (group)

o (all others; also known as “world”)

The user/owner of a file or directory is generally the person who created it.

The group specifies the group to which the file or directory belongs.

“Others” are all other users on the system.

You can give each of the three types of users separate permission to read, write, or execute each of your files and directories. By assigning permissions to read, write, and execute a file or directory, you can regulate not only who is able to access your files and directories but also how they can access your files and directories.

The meanings of the three types of permissions differ slightly between ordinary files and directories:

PERMISSION	FOR A FILE	FOR A DIRECTORY
r (read)	Contents can be viewed or printed	Contents can be read (e.g. ls), but not searched (e.g., ls –l) Normally, r and x are used together
w (write)	Contents can be changed or deleted	Entries can be added, removed or renamed
x (execute)	File can be used as a program	Directory can be searched (e.g.ls -l) and changed into (cd)

Taken together, all the permissions for a file or directory are called its “permission code.” For example:

drwxrwxrwx

A permission code consists of four parts:

- The first character in the code shows the file type. A - (hyphen) indicates an ordinary file and the letter d indicates a directory. The letter l indicates a link, which follows the same rules as files. Any other character indicates an I/O device.
- The next three characters show user (owner) permissions in the order r (read), w (write), and x (execute).
- The next three characters show group permissions in the order r, w, and x.
- The last three characters show permissions for all others in the order r, w, and x.

An s in the owner “execute” position indicates that both the set user ID (setuid) bit and execute are set. An S indicates that only the setuid bit is set. An s in the group “execute” position indicates that both the set group ID (setgid) bit and execute are set. An S in that position indicates that only the setgid bit is set. The setuid and setgid bits are used to allow programs to access files and processes that would otherwise be inaccessible. The setuid bit is similar to PROGID for Guardian program files.

A t in the others “execute” position indicates that both the sticky bit and execute are set. A T indicates that only the sticky bit is set. When the sticky bit is set for a directory, only the directory owner can delete the directory or its files, even though other users have write permission. In some older UNIX implementations, when the sticky bit is set for a process, that process is retained in the swap area even when not being executed. OSS does not support this usage.

Note: A hyphen in the permissions for owner, group, or others indicates that the permission (read, write, or execute) normally holding that position in the code is not granted.

As with Guardian, OSS security vectors have relatively coarse granularity. HPE recommends use of OSS ACLs for finer-grained control.

Modifying file and directory permissions and ownership

If you are the owner of a file, you can change the file permissions with the chmod command. SUPER.SUPER, its aliases, and members of the OSS-SECURITY-ADMINISTRATOR group also can change its permissions.

You can use chgrp to change its group ID to a group that you belong to (either your effective group or one of your supplementary groups). If you do not own the object and/or do not belong to the new group, then you can’t make the change; only SUPER.SUPER or one of its aliases or a member of SECURITY-OSS-ADMINSTRATOR group can do it.

Use chown to change file ownership. The only users who can change file ownership are SUPER.SUPER or one of its aliases (effective user ID = 65535), or a member of the Safeguard SECURITY-OSS-ADMINISTRATOR group. Unlike Guardian files, other users cannot give away ownership of their own files.

chown is not allowed on a remote system object.

OSS file and directory permissions: best practices

Use SAFECOM ADD SECURITY-GROUP SECURITY-OSS-ADMINISTRATOR (SOA) to add the authorization record for that group and give execute authority on the ACL to the users who will belong to the group and be allowed to execute commands that are restricted to the group. This will allow designated users to perform these functions without being authenticated as SUPER.SUPER.

Restricted filesets

If you plan to configure any OSS filesets as restricted filesets to control access by SUPER.SUPER and its aliases, also create and populate the SECURITY-PRV-ADMINISTRATOR (SPA) group. HPE recommends that you not let SUPER.SUPER either belong to or own either the SOA or SPA group if you are planning to use restricted filesets, as that would make it possible to undermine their purpose.

See SAFEADMIN Section 6, Managing Security Groups, for details.

File and directory other/world permissions

Read access:

- Directories that are generally searchable.
- Files in /etc
- Scripts that are generally executable
- Executables that are generally debuggable
- /home/*, depending on individual requirements
- Configuration files, depending on individual requirements

Write access:

Whenever a directory has “rest of the world” write access it is advisable to have the sticky bit on (i.e: drwxrwxrwt)

- /tmp/
- /var/tmp/
- /usr/tandem/hpjmeter/var/fifos/
- /usr/tandem/hpjmeter/var/log/
- /usr/tandem/sqlmx/log/
- /usr/tandem/sqlmx/USERMODULES/ - optional, depending on how you are controlling SQL/MX usage
- /usr/tmp (optional)
- /var/preserver (optional)

Execute access:

- Executables that are generally runnable
- Scripts that are generally executable
- Directories that are generally searchable.

setuid and setgid

Caution: Use setuid and setgid only where necessary. If possible, avoid using setuid/setgid scripts at all. If you must use them, be aware of their security implications and carefully review their contents.

To locate setuid and setgid scripts, use:

```
OSS> find / -W NOG –W NOE -type f -user SUPER.SUPER -perm -04000 | xargs ls –l
OSS> find / -W NOG –W NOE –type f –group SUPER –perm -02000 | xargs ls -l
```

HPE currently ships the following files with the setuid bit on:

- /bin/at
- /bin/atq

- /bin/atrm
- /bin/crontab
- /bin/ipcs
- /bin/newgrp
- /bin/rsh
- /usr/lib/expreserve
- /usr/lib/exrecover

Modifying file and directory permissions and ownership
Use chown <user> <file> rather than chown <user number> <file>, and be very careful with the recursive form (i.e., chown –R …) along with an argument list such as ".*". ".*" . That form will also match "..", which will result in applying the command recursively in the directory above the current directory. See OSSUSER, section 9, Managing Access to Files and Directories, and the chmod, chgrp and chown man pages for additional details.

OSS ACLs: background

OSS ACLs allow the security administrator to arrange more granular permissions given to multiple users or groups that otherwise would need to be given through the permissions for "other", which would allow everybody on the system to access those files or directories.

OSS ACLs also allows the administrator to set "default" ACLs on directories that automatically will be inherited by files and directories created in that directory.

OSS Version 3 catalog filesets support access control lists (ACLs), in addition to basic file security, for directories, regular files, FIFO special files, and bound AF_UNIX sockets. OSS ACLs allow a process whose effective user ID matches the file owner, super ID, or a member of the Safeguard SECURITY-OSS-ADMINISTRATOR security group to permit or deny access to a list of specific users and groups.

All OSS system calls that include pathnames are subject to the ACLs on any directory or file in the path.

There are two commands to show and control OSS ACLs.

- getacl - Lists access control lists (ACLs) for files and directories
- setacl - Modifies access control lists (ACLs) for files and directories

Many of the OSS utilities have been extended to work with ACLs. As examples, ls shows the existence of OSS ACLs and chmod can change an OSS ACL.

For a detailed description of OSS ACLs, including examples, see the acl(5) reference page either online or in the Miscellaneous Topics section of OSSSYSCALL.

OSS authorization SEEPs: background

Beginning with the J06.15 and H06.26 RVUs, an OSS Security Event-Exit Process (SEEP) is supported and provides additional file-access authorization. Unlike other SEEPs, OSS authorization SEEPs are configured in SCF at the OSS fileset level and it is the fileset's OSS Name Server rather than Safeguard that communicates with the SEEP. Partner OSS SEEPs are available.

You should consider using a SEEP if you want to provide access controls beyond what is available through OSS. Use of an authorization SEEP also has the potential to reduce OSS ACL proliferation.

You can purchase a partner authorization SEEP or, if you wish, write your own. If you write your own, you also will need to provide a user interface to configure its rules.

See OSSMOG Section 9, Using an OSS Security Event-Exit Process (SEEP), for details.

sudo

OSS supports the sudo program, which allows a permitted user to run some (or all) OSS commands as the super ID or another user as specified by the sudoers security policy.

By default, sudo audits the command that it explicitly runs to both syslog (EMS) and the sudo log file (default location: /usr/coreutils/var/log/sudo.log). If you are using XMA or a partner product that generates alerts based on EMS events, you can trigger alerts on activity by sudoers. As an alternative, if you have installed a mailer program you can configure the sudo mailto options to generate alerts directly.

See the sudo man page for details.

SQL/MP object security: background

SQL/MP itself executes within the Guardian environment and uses Guardian security for its objects. Safeguard ACLs can be used to restrict placement of SQL/MP objects at the volume/subvolume level, but ACLs cannot be placed on individual objects. Authorization SEEPs can control access at the object level.

SQL/MP programs can be either Guardian or OSS programs. Authorization for OSS programs is done through OSS.

The local owner of a table, view, index, collation, or program, the local group manager, the local super ID, or the remote owner with purge authority generally has the authority to perform DDL statements on these objects. Authority to purge an object is required to drop a table, index, view, SQL program stored in a Guardian file, or collation from the database.

A group manager (user 255) can read or write to any local table owned by a group member and can execute an SQL program that runs in the Guardian environment that is owned by any group member. Remote tables, views, and programs must be secured for remote access. When a statement requiring access to an object is compiled, the catalog that describes the object must be accessible by the group manager. To alter attributes of a table, view, index, collation, or SQL program stored in the Guardian environment, or to run a DDL statement, a group manager requires purge authority.

SQL/MP security issues cover two areas:

- Security of a catalog that contains descriptions of SQL objects
- Security of SQL objects

Note: Allowing access to the catalog does not automatically allow access to the objects described in that catalog. Access to the catalog is required in addition to access to the objects for execution of:

- DDL statements
- DML statement compilations for SQLCI or dynamic SQL
- Most utility commands
- SQL program compilations

Network databases require remote passwords (at the network level) and network security strings for both catalogs and objects to allow remote access.

When an SQL object is created, the ownership defaults to the owner of the session or program. The security of the object defaults to either the security of the underlying table or the current default security, unless the statement creating the object provides another security string. SQLMPMGMT, Section 5, Creating a Database, contains additional object-specific information about security.

The security attributes of a table, view, index, or SQL program that runs in the Guardian environment can be changed by an ALTER statement.

The security string for an object must be set to allow users who have write authority to also have read authority.

A change in the ownership of an object affects the interpretation of the security string. SQL interprets the security string at run-time against the user ID of the new owner. The change does not apply to a running SQL program until program execution ends.

The owner and security of an underlying table determine those attribute values for indexes on the table. If you change the owner or security string for the underlying table, SQL automatically changes the owner or security string for any indexes on the table.

The CLEARONPURGE and NOPURGEUNTIL attributes for a table do not dictate these attribute values for dependent indexes. You can set these two attributes independently for indexes.

The owner of a base table determines the owner of a dependent protection view. If you change the owner of a table, SQL automatically changes the owner of any dependent protection view.

If you change the owner of a program, SQL automatically sets the PROGID attribute to NO, regardless of the original setting.

See SQLMPMGMT Section 4, Planning Database Security and Recovery, Table 4-1, for details on the necessary authority for various types of database operations on tables and programs.

See SQLMPMGMT Section 5, Creating a Catalog, for details and examples of securing catalogs base tables, protection views, shorthand views and collations.

See SQLMPMGMT Section 7, Adding, Altering, Removing, and Renaming Database Objects, for the security considerations associated with these actions.

SQL/MP object security: best practices

You should create the smallest number of catalogs logically possible, as dictated by your business operations. For authorization, you should create the simplest authority and security scheme possible. Dependent views, indexes, and programs should be owned by the same user ID, and only that user ID should have purge authority. With this authorization scheme, DDL operations and utility operations that can affect the entire set of dependent objects, such as DUP, are simplified. Because anyone who has authority to purge an object can drop that object, an authorization scheme should limit the authority for purging.

For an authorization scheme, you should establish catalog boundaries along the lines of application and user access requirements. Associate catalogs with sets of tables logically associated or used together. With this scheme, security follows the catalogs you choose.

Configure Safeguard volume and subvolume ACLs as needed to restrict the placement of SQL/MP objects. This is particularly useful if you have multiple SQL/MP applications running on the same system and wish to allocate each one its own set of database volumes or subvolumes. If you place a Safeguard ACL on the subvolume that contains a catalog, that protects the entire catalog.

Caution: Altering the security of a collation can restrict access to any table, index, view or program that uses it.

See SQLMPMGMT Section 4, Planning Database Security and Recovery, for additional guidelines and examples.

Note: SQLCI2 cannot perform privileged SQL/MP operations even when run by SUPER.SUPER unless it has been licensed.

Caution: Do not license SQLCI2. If you must use it to perform privileged operations, license a side copy and purge or unlicense it as soon as you have completed the operations.

A licensed SQLCI2 process (licensed program) can perform privileged operations, such as deleting or updating rows in catalog tables. Normally only the super ID can perform these operations because of the potential risk to the database. The super ID must explicitly license program files before beginning. These operations can be extremely dangerous to the consistency of the database and the data dictionary. Only the most extreme situations should require the use of a licensed SQLCI2. Only the most knowledgeable SQL/MP manager should attempt to correct problems with a licensed SQLCI2 process.

See SQLMPMGMT Appendix A, Licensed SQLCI2 Process, for details.

SQL/MX object security: background

SQL/MX executes within the OSS environment, but database files reside within the Guardian environment. For SQL/MX tables, data access uses the ANSI GRANT/REVOKE authorization model. SQL/MX is not tightly integrated with Safeguard; however, with SQL/MX version 3.3 and later, Safeguard volume protection can be used to control where SQL/MX objects are created. SQL/MX uses Guardian security for SQL/MP objects,

Installation and fallback: cautions

The initial installation of SQL/MX, done with the InstallSqlmx script, must be run by SUPER.SUPER (not an alias to it).Do not change the security setting on the anchor file created by SQL/MX. If the file is modified, SQL/MX ceases to operate.

If you do not use DSM/SCM to do a SQL/MX installation, you must explicitly license the ZCLIPDLL public DLL. If this step is not performed, program load errors might occur.

Make sure that the following files are licensed in $SYSTEM.SYSTEM:

- IMPORT
- MXUTP
- MXIMPDDL
- MXAUDSRV
- MXCMP
- MXESP
- MXRTDSRV
- MXTOOL
- ZMXBRDLL

Make sure that the mxci file in /usr/tandem/sqlmx/bin has both read and execute permissions.

See SQLMXINSTALL Section 2, Installing SQL/MX, for details.

If you are falling back to a version of SQL/MX prior to 3.1, you should remove all security administrator grants before falling back.

See SQLMXINSTALL Appendix A, Removing security administrator grants, for details.

SQL/MX object security: best practices

Create a separate SQL/MX security administrator group if running 3.1 or a subsequent version. Security administrators manage access to SQL/MX data but they do not have access to the underlying data itself unless explicitly GRANTED access by an object owner or designee or through PUBLIC access. Use this group to administer security. Consider having only one user create database objects (such as a user designated as the database administrator). In that way, no users other than the security administrators and database administrator may grant object access.

Periodically audit the set of security administrators.

To obtain detailed security administrator information:

```
OSS> mxci
```

```
>> -- In the following example, substitute your system's name for <system name>.

>> -- For example: SET SCHEMA nonstop_sqlmx_mysys.system_security_schema;

>> SET SCHEMA nonstop_sqlmx_<system name>.system_security_schema;

>> SELECT * FROM privileged_users FOR read uncommitted access;
```

Use `GRANT/REVOKE CREATE CATALOG` to restrict the set of users who can create catalogs.

Use GRANT UPDATE (<column list>) to limit modification to specific fields of a table to those users entitled to update their contents.

Use ANSI VIEWs to limit SELECT access to specific fields of a table to those users entitled to see their contents.

See Section 1, Security, Section 2, GET ALL SECURITY_ADMINS, GRANT SECURITY_ADMIN and REVOKE SECURITY_ADMIN, and Section 10, System Security Schema Tables, in the appropriate version of SQLMXREF, and Section 3, Planning Database Security and Recovery and Section 5, Creating a SQL/MX Database, in the appropriate version of SQLMXMGMT for details.

See Section 5, Privileges Required to Execute Utilities, in the appropriate version of SQLMXREF for utilities-related details.

Consider using the OSS security settings on the MXCI program file to control access and limit the ability of users to execute adhoc queries against the database.

Normally, SQL/MX uses Guardian user names (groupname.username) for GRANT and REVOKE, externally, and userids (groupnumber, usernumber) internally. If a user is deleted from the system after permissions have been granted to it, then SQL/MX can no longer look up the userid associated with that username. To revoke permissions in such cases, use a userid instead of a username as the parameter to the REVOKE statement.

To prevent unauthorized access to SQL data, network access and firewalls must be properly maintained and monitored for suspicious activity. In addition to network security, application-level access control and audit is necessary in order to prevent unauthorized exposure of data or a possible denial of service attack.

JDBC and ODBC and persistence frameworks must be designed to reflect the user needs of each database associated to those access methods.

Ported Binaries on HPE NonStop, Java SE (NSJ)

Authentication and authorization

The Java Authentication and Authorization Service (JAAS) is integrated into NSJ. JAAS implements a Java version of the standard Pluggable Authentication Module (PAM) framework, allowing applications to be insulated from underlying authentication technologies. Applications enable authentication by instantiating either a LoginContext object that references configuration information to determine the authentication technology to use or a LoginModule that performs authentication itself. NSJ7 includes Oracle-produced LoginModule interfaces including JndiLoginModule and KeyStoreLoginModule, but does not provide a LoginModule that interfaces to Safeguard. You also may choose to implement your own LoginModule.

Secure communications

NSJ includes support for Java Secure Socket Extension (JSSE), which supports protocols such as TLS and SSL.

Cryptography

NSJ includes support for Java Cryptography Architecture (JCA), including the Java Cryptographic Extension (JCE), which provides a wide range of cryptographic services and algorithms for building secure applications.

HPE recommends enabling java.security.nativeRNG:

```
-Djava.security.nativeRNG=true
```

This option reduces the processing time required to generate the first random number in a Java application.

Note: This option can be used in NSJ 8.0, NSJ 7.0 and NSJ 6.0.

See NSJ6PROGREF and NSJ7PROGREF, About this Manual, Feature Changes, and Section 4, Java Authentication and Authorization Services (JAAS). Also see NSJSEC for general Java security information.

Securing operating system software: best practices

Unless your security policy states otherwise, SUPER.SUPER should own system files, including libraries, user tools, system utilities, and microcode files. Guardian write and purge permissions should be set to owner only (local or network, according to your policy) and read permission should be limited to the owner unless there's a compelling reason to grant it to other users. Execute permission should be limited to those that need it. For Guardian files, the primary locations are the $SYSTEM.SYSTEM and $SYSTEM.SYSnn subvolumes, along with various subvolumes whose names begin with Z and are typically also on $SYSTEM. The pmsearchlist is a good place to check for other candidate subvolumes to protect.

File integrity monitoring products such as XYGATE Compliance PRO (XSW) or similar partner products can be of assistance in identifying files of interest and appropriate security settings.

One option for locating all HPE-supplied files installed through DSM/SCM is to request and review an independent DSM/SCM snapshot of the system.

See SECMGMT Section 2, Sanitizing a NonStop System, Section 3, Securing Critical Objects, and DSMSCMUSER Section 17, Creating an Independent Snapshot of the Target System, for more information.

Securing and monitoring program files

Tightly restrict write security for program files to prevent the attachment of arbitrary libraries. Use HPE Compliance PRO (XSW) or a partner file integrity monitoring product to detect and alert on changes to important program files.

You can use XSW or a partner file integrity monitoring product to monitor changes to PROGID and LICENSE.

Use Safeguard ACLs or an authorization SEEP to tightly control CREATE access to subvolumes containing program files to prevent the installation of malicious programs or macros with file names such as $SYSTEM.SYSTEM.EDOT that are similar to commonly-used utilities.

Persistence manager

The persistence manager, $ZZKRN, is used to create, monitor, and restart (if necessary) processes as configured through SCF. Ensure that its configuration files and the configuration files for the processes that it manages are appropriately secured.

See SCFKERNEL, Section 3, Configuring and Managing Generic Processes for details.

Securing sensitive system data: best practices

Do not alter the security of files created by individual subsystems (configuration files, log files, KMSF files, etc.) unless you are sure that you know what you are doing.

If you choose to tighten security on any of these kinds of files, test out the changes on a non-production system first to ensure that you have not inadvertently rendered part or all of the subsystem inoperable.

Make sure that you can restore your current configurations, either by backing up the configuration files directly or, where available, having the subsystem create an obey file that can be replayed. Example:

```
TACL> SAFECOM
= DISPLAY AS COMMANDS ON
= info <object>, detail
= DISPLAY AS COMMANDS OFF
```

Information on individual subsystems' configuration files appears in their descriptions elsewhere in this document.

OSS security configuration files

Most of the security-related directories and files that are used in UNIX systems are absent from the OSS environment. User/group administration and audit are based on Safeguard rather than the usual OSS mechanisms; OSS administration of device access does not use certain files that normally are present in the /dev directory, and a number of other files are not used.

See OSSMOG, Section 9, Administrative Files and Directories for details.

Other OSS files and directories

All OSS Monitor and server database files must always be owned by the super ID (255, 255) and secured as having read, write, execute, and purge privileges for only the super ID (----) in the Guardian environment.

The /etc/install_obsolete directory and the Pcleanup command should be secured such that only the super ID or a trusted user has access to them. A single file surreptitiously added to /etc/install_obsolete could be used by an intruder to cause severe damage to mounted filesets.

The Guardian subvolume from which OSS files are installed by COPYOSS or PINSTALL should be carefully secured and its contents periodically checked for validity. A pax archive surreptitiously added to that subvolume could be used by an intruder to install files such as:

- Viruses, worms, Trojan horses, or rabbit programs
- Substitutes for standard utilities, containing logic bombs or back doors

For bash, ksh and sh:

/etc/profile must be secured for root-only write access, with world read access. Execute use is not needed. /etc itself should not be world-writeable.

$HOME should be secured for USER write access; OTHER, and probably GROUP, should not have write permission.

Execute access is not recommended for $HOME/.profile.

To secure the above directories and files appropriately, along with /:

```
OSS> chmod o-w / /etc /etc/profile /home /home/.profile
```

Alternative shells may have different default behaviors and local configuration requirements.

See OSSSHELL, sections ksh(1) and sh(1), or the OSS man pages for more details.

Securing application files

Application files – Enscribe files, SQL objects, SQL user catalogs, application-level obey files, etc. – should belong to appropriate non-SUPER users. As noted earlier, the best practice is to use XAC or a partner product with similar capabilities to restrict access by type to those users that require it (least privilege), with appropriate ACL and Guardian or OSS security settings as fallbacks. Otherwise, configure appropriate ACLs and fallback Guardian/OSS security.

Your application providers should specify security configuration options and best practices.

The user mask (umask) in OSS is used to establish the default permissions for a new OSS file or directory. The default OSS setting of umask is 022, which will have the effect of denying write access to group members and others but allowing them read and execute access even if the program creating the file or directory by default assigns looser permissions. Your policy may call for tighter default restrictions for all users; if so, edit the umask value in /etc/profile to adjust the default accordingly. If certain users require different umask settings, they can adjust it accordingly in their individual .profile files.

Owners and groups for OSS files and directories can be changed with chown and chgrp, respectively. Use of chown requires SUPER.SUPER privileges; use of chgrp also requires that the user be a member of the group to which he/she is changing the file ownership.

See SECMGMT, OSSMOG, Section 9, Managing Security, and OSSUSER, Section 10, for details.

Securing utilities and commands

For complete command-level control over all utilities, use XYGATE Access Control or a partner product with similar capabilities. In the absence of this type of control:

Most subsystem manuals include tables describing which users can issue which commands or subcommands.

Older utilities such as SCF effectively have a three-role model: users/aliases who do not belong to the SUPER group have information-only access, SUPER group users other than SUPER.SUPER (collectively, SUPER.notsuper) can perform most modifications, and commands with the highest destructive potential are reserved for SUPER.SUPER.

Safeguard supports the configuration of a number of Safeguard and OSS security administration groups, and you should use them. It also supports newer groups for NonStop Volume Level Encryption administration, persistent process configuration, and certain tape catalog functions.

xxxxLOCL and xxxxCSTM files

A number of subsystems, including TACL, FUP, and (as of L17.02/J06.21) Safeguard's SAFECOM, look for system-level customizations in files named TACLLOCL, FUPLOCL, SAFELOCL etc. in $SYSTEM.SYSTEM and user-level customizations in files named TACLCSTM, FUPCSTM, SAFECSTM, etc. in the user’s default volume/subvolume. You should create the appropriate set of xxxxLOCL files during system installation, secured either OOOO or UUOO. If you have no specific site-level customizations, create them as empty files.

When a new user is created, you should create the appropriate set of xxxxCSTM files for that user as empty files, secured either OOOO or UUOO. If FUP does not find an existing FUPCSTM file for the user who started it, it will create a default file. SAFECOM has similar behavior. See TACLREF, FUPREF and SAFEREF for additional information.

Safeguard: additional considerations

Background

Safeguard supports a number of OBJECTTYPEs, which allow your security administrator to define the user or groups of users who can add new subjects or objects to the Safeguard database. Each kind of subject and object can be given its own unique OBJECTTYPE protection record.

The OBJECTTYPEs are ALIAS, DEVICE, DISKFILE, DISKFILE-PATTERN, GROUP, OBJECTTYPE, PROCESS, SUBDEVICE, SUBPROCESS, SUBVOLUME, USER, and VOLUME.

OBJECTTYPE OBJECTTYPE is the "meta" object type that controls the ability to create new OBJECTTYPE records.

Each OBJECTTYPE has a default setting for who can place an object under Safeguard control. The default applies until you create an OBJECTTYPE protection record for that subject/object.

Note: OBJECTTYPE DISKFILE has no effect on default protection for a user's disk files; it only controls use of the ADD DISKFILE command.

Note: OBJECTTYPE USER controls who can add users, aliases, and groups.

Best practices

As mentioned earlier, you should configure protection records for all Safeguard OBJECTTYPEs during initial system setup, as at least some of the defaults are unlikely to be acceptable for a tightly-controlled system.

See SAFEREF, Section 12, OBJECTTYPE Security Commands, for more details.

Additional security considerations for individual subsystems

DSM/Tape Catalog (DSM/TC)

Tape media

DSM/TC by default allows all SUPER group members to issue a number of sensitive commands. HPE recommends that you consider defining and populating the SECURITY-MEDIA-ADMIN group to further restrict access to the more sensitive commands. See DSMSCMTAPE Section 2, MEDIACOM Commands, for additional details.

EMS

The default Guardian security vector for your EMS primary and alternate collector logs is COOO. If necessary, alter it by using the SECURITY command to change the value of the PROTECTION attribute. Consider adding Safeguard ACLs for the logs’ subvolumes if you need finer-grained control (e.g., read access for only a subset of SUPER group members).

If you wish to change the user ID that the primary collector uses to access log files from the default of SUPER.SUPER to another user ID, do through EMSCCTRL or the SPI CONTROL command. If you have alternate collectors, start each one under the desired user ID; the only way to change their user IDs are to stop them and, running as the correct user ID, restart them.

Although the EMS log is not primarily a security log, it contains some events that you might consider to be security-relevant. Insofar as practical, configure it to not lose events. As an example, the default for ROTATEFILES (ON) causes the oldest log file in the subvolume to be purged when the log files are full rather than stopping logging as it would with ROTATEFILES OFF. The details are slightly different for primary and alternate collectors.

See EMS Section 12, Configuring EMS, for additional information.

FUP

General: background

You need both read and write access to a file in order to issue an ALTER command against it. To rename a file, you also need purge access if you are not the super ID.

The FUP INFO command can be used to identify all Guardian files that have LICENSED, PROGID, CLEARONPURGE and/or TRUST set. It also can display the underlying Guardian security of a file protected by Safeguard at the individual file level.

FUP is used to license SQL/MP object program files. They cannot be licensed through SQLCI.

For SQL/MP files, GIVE applies only to object files. You must use SQLCI to give away ownership of other SQL/MP files. Similarly, PURGE and PURGEDATA apply only to object files and SQLCI must be used for other objects.

FUP can display information about OSS files (including their security vectors and whether they are protected by OSS ACLs), SQL/MP object program files and SQL/MX files, but cannot manipulate them in any way.

Non-Safeguard-protected files: background

You can preserve the source file's owner ID and Guardian security vector in the copied file if you use SAVEID or SAVEALL with FUP DUP. The LICENSE attribute is preserved only if the PAID of the current FUP is SUPER.SUPER and the target file resides on the node where FUP is running. The same rules apply to PROGID, except that it also is preserved if the PAID of the current FUP is the file owner. CLEARONPURGE is transferred unconditionally.

You need to be either the owner or SUPER.SUPER to GIVE ownership of a file. If you are not SUPER.SUPER, you also need purge access to the file.

GIVE clears PROGID. After the GIVE you need to use SECURE to set it again; the usual rules apply.

You can use REVOKE to reset CLEARONPURGE and PROGID if you are either the file owner or SUPER.SUPER. You must be SUPER.SUPER to revoke a file's LICENSE attribute.

Safeguard-protected files: background
You need create access to the destination volume and subvolume as well as read access to the input file in order to duplicate a Safeguard protected file.

Caution: If you use DUP with the PURGE option but do not have create access to the target file, the original file at the target location is purged but the new version is not created.

Caution: A file's Safeguard protection is not automatically inherited by the target file. It will inherit any applicable volume-level and subvolume-level ACLs on the target volume, and will not be Safeguard-protected if none apply. You will need to use SAFECOM to restore or set Safeguard protection for the new file.

As with non-Safeguard-protected files, for FUP DUP both SAVEID and SAVEALL transfer the source file's owner and corresponding security to the target file.

Caution: the GIVE rules are more complicated, especially if the file has a persistent protection record (one that stays around after the file is purged and is applied to any subsequent file created with the same name). The REVOKE rules for LICENSE also are more complicated. As with non-Safeguard-protected files, only SUPER.SUPER can revoke a file's license.

If a file is under individual Safeguard protection, i.e. has its own ACL, FUP is aware and will display **** for its security vector. FUP is not aware of volume-level or subvolume-level diskfile ACLs and hence can't flag files that have only that level of Safeguard protection, so you cannot rely on **** as an indicator of Safeguard protection.

General: best practices

You can set the NOPURGEUNTIL attribute to prevent a file from being purged before a specified date and time.

You can use FUP INFO to identify all of the Guardian files owned by a specific user.

You can find out what processes and associated users have an Enscribe or SQL/MX file open by using the LISTOPENS command.

Caution: PURGEDATA does not physically purge the file contents; it simply resets the end of file to zero. This applies whether or not the file (or system) has CLEARONPURGE set. PURGE access allows deletion of both a file and its contents, but you can effectively purge the contents of a file without PURGE access through a combination of PURGEDATA and DEALLOCATE.

Caution: When a file with CLEARONPURGE set is purged, its disk process is going to rewrite the contents with zeros up to the end of the last allocated extent. The disk process has some built-in pacing for the writes, but this activity still has the potential to negatively affect application and system performance.

See FUPREF, Section 2, DUP, GIVE, INFO, LICENSE, LISTLOCKS, PURGE, PURGEDATA, REVOKE and SECURE for additional details.

iTP Secure WebServer

iTP Secure WebServer: background
Secure transport
iTP Secure WebServer supports the following standards for data in motion encryption:

- Transport Layer Security (TLS 1.0, TLS 1.1 and TLS 1.2)
- Secure Sockets Layer (SSL 3.0)
- Secure Hypertext Transfer Protocol (Secure HTTP)

For TLS and SSL, you can specify the ciphers to be used. TLS and SSL session keys can be cached globally across httpd serverclass instances.

iTP Secure WebServer also supports:

- X509 version 3.0 certificates
- Certificate chains
- Security certificates with non-English characters
- Client authentication in SSL 3.0 and TLS
- Digest access authentication

Note: While iTP Secure WebServer can generate key pairs and manage certificates through its keyadmin utility, it does not function as a Certificate Authority (CA). It can, however, generate certificate signing requests.

Client access control

You can control access to the iTP Secure WebServer on the basis of factors such as host name, time of day, user name, browser type and version, and authentication method.

Audit

iTP Secure WebServer supports an easily-parsed extended log format (ELF) that includes the access, error and security information of each request, along with other useful information. It also supports the widely-used Common Log Format (CLF).

Administration

The Administration Server lets you establish and modify configurations and control and monitor one or more iTP Secure WebServer environments from a Web client. The admin httpd process provides the interface between your web client and the Administration Server.

iTP Secure WebServer: best practices

HPE recommends using iTP Secure WebServer, rather than iTP WebServer. iTP WebServer does not include:

- TLS/SSL support
- Crypto key-exchange methods
- Certificate support
- Certificate authorization lists
- Digest access authorization
- Client authentication through advanced TLS/SSL methods

Secure transport

Installation generates a 90-day self-signed test certificate, protected by a password that you choose. You must not use the self-signed certificate in production; obtain a valid commercial-grade certificate from a reputable CA and use the keyadmin utility to install it in the key database. See "SSL encryption: background and best practices", below, for additional details about best practices for installing and checking SSL certificates.

If you are running version 7.4 or later, you should use ciphers that provide Perfect Forward Secrecy (PFS). PFS prevents an attacker who is capturing the encrypted traffic from decrypting it even with after-the-fact access to the TLS/SSL certificates. PFS is achieved by using an ephemeral Diffie Hellman parameter; the respective ciphers supporting PFS can be usually identified by names beginning with “DHE-…“. For the ephemeral Diffie Hellman exchange, iTP has to be set up with initial Diffie-Hellman parameters during installation. Note that the generation of the initial Diffie Hellman parameters is very resource intensive and can take a long time. To improve overall SSL performance, configure global session key caching across all instances of the httpd serverclass.

The default setup for iTP WebServer creates ephemeral Diffie-Hellman parameters with a parameter size of 1024. Many Internet browsers now have tightened their security beyond that level. You can regenerate the dh_params file with a higher value as follows:

```
keyadmin –dhparams –out dh_params –length 2048 –overwrite
```

The maximum length is 4096. The dh_params file is referenced in your httpd.stl.config file in the AcceptSecureTransport configuration. An example:

```
AcceptSecureTransport –transport /T/ZTC0 –port 8043 –address someserver.com –cert :CN=someserver.com,O=HPE} –dh_paramsFilepath dh_params.
```

You can specify the set of acceptable ciphers, in order of preference, through the AcceptSecureTransport –cipher option. HPE recommends not enabling cipher suites that use MD5 as their MAC algorithm. Again, see “SSL encryption: background and best practices” for additional guidance.

You should explicitly disable use of SSL by specifying the –nossl option and disable use of TLS 1.0 by specifying the -notls1.0 option. If you need to disable TLS 1.1, specify the –notls1.1 option.

Client access control

You should restrict client access to what is necessary. There are a number of approaches to consider:

- HPE recommends that you enable client authentication whenever practical. Use the AcceptSecureTransport configuration directive to either request or require that the client present a valid certificate.
- You can restrict client access to secure transport only on a per-region basis by specifying RequireSecureTransport. You can use CGI environment variables in Region commands to restrict access based on client key length, DN, etc.
- You can either grant or deny access at the region level based on the client host name or IP address.
- You can use either basic or digest access client authentication to require a username and password at the region level. Based on your policy, you may choose to either rely on the underlying operating system authentication (-safeguard option) or create and manage a separate user/password database for use within this region. If you create a separate database, you will need to manage it through the useradm utility. HPE recommends that you configure Safeguard users (-safeguard option) unless there is a compelling reason to maintain a separate user/password database.

Key database file

You protect the key database (stored in two files, starting with release 7.5), which contains private keys and public key certificates, though a password. If you store that password in a configuration file or another file, protect that file at least as carefully as the key database. You also should ensure that the key database file is owned by the user ID of the webmaster and secured with read/write access granted only to that user.

You can use the bin/keyadmin –list option to display a list of the stored keys and certificates and their attributes.

You can use the –initdefaults option to update the default root certificates in the database.

Audit

Enable CLF logging, and configure XMA to include those logs among its security activity sources.

Determine your log rollover/rotation policy and configure the server accordingly.

Administration

HPE recommends creating a SUPER group functional user to perform the installation and serve as webmaster. Do not use SUPER.SUPER to install the software. This is especially important if you are allowing users to write and execute their own Common Gateway Interface (CGI) programs and scripts, as they run under the ID of the user who established the environment.

Note: The standard port for TLS and SSL is 443. If installation is done by a user who is not a member of the SUPER group then you will need to reconfigure it to 1024 or higher.

The default configuration gives all users execute and read permission for the bin directory, which allows them to specify a configuration file to start their own server copies. To restrict this ability, tighten the security on either the bin directory or the bin/httpd file.

You can modify the httpd.config and httpd.stl.config files to restrict listening to only secure ports and, if desired, reconfigure which ports are used. You also can reconfigure the ports used by the Administration Server.

HPE recommends that you configure the Administration Server to accept only secure connections. Modify the httpd.adm.config file to add a RequireSecureTransport command to the Region directive for the /admin/* region. If practical, choose the –auth option of RequireSecureTransport to require client authentication.

If you enable the PUT request method to replace or create content, specify a script to perform validation before permitting the update.See ITPWSADMIN Sections 1-4, and 6 for more details.

NetBatch

NetBatch has its own defined set of roles, including:

- SUPER.SUPER
- NetBatch supervisor, defined as all users with EXECUTE access to the NETBATCH program file
- NetBatch job owner
- Other users

Installation and configuration Securing files

Once you have validated all of its job definitions, secure the JOB file “OOOO” to the local super ID by using the FUP SECURE command with the PROGID option.

Caution: If the JOB file is not secured “OOOO” to the local super ID by using the FUP SECURE command with the PROGID option, warm-start processing fails.

Licensing files

The NETBATCH program file is the only program file you must license before use (the program contains privileged code). To license the file, use the FUP LICENSE command. For example:

```
> FUP LICENSE $SYSTEM.SYSTEM.NETBATCH
```

Note: The NetBatch database can be installed only by SUPER.SUPER. To verify that this has been done, NetBatch checks that the data files also have been licensed as a sign that the local SUPER.SUPER has approved its contents.

BATCHLIB configuration

NetBatch has a couple of security options that are configured by creating a small TAL procedure, NBFLAGS that you bind into the BATCHLIB library file. NBFLAGS controls the settings of the following options:

- Allow jobs owned by frozen users to run (by default, they are not allowed). For information on frozen users, see the FREEZE USER command in SAFEREF.
- Disable the BATCHCOM commands CHANGEUSER and RUN (by default, they are enabled)

HPE recommends retaining the default behavior of not allowing jobs owned by frozen users to run. However, this depends on your policies; if, for example, you routinely freeze a person’s user ID(s) while on vacation, you might still need to let their scheduled jobs run.

Consider disabling CHANGEUSER for day-to-day use. If you do disable it, you might wish to build a separate version of BATCHCOM with it enabled for use under specific circumstances such as using OBEYFORM to transfer a large number of jobs from one scheduler to another.

Caution: Passwords associated with CHANGEUSER are in the clear.

Infiles

Ensure that all infiles are secured to restrict modification to only authorized users.
See NETBATCH Section 2, Software Installation, for more details.

Security-related Scheduler attributes
AT-ALLOWED: Determines whether users without execute access to the NETBATCH program file can submit jobs with the AT attribute - OFF

LOCALNAMES: Makes the scheduler treat jobs submitted from licensed requesters on the specified nodes as local jobs, not as remote jobs.

(An example of such a requester is NetBatch-Plus.) Through the scheduler, the remotely submitted jobs gain the same access privileges on the scheduler's node as that user would have with locally submitted jobs.

A scheduler without the LOCALNAMES attribute treats remotely submitted jobs as remote jobs subject to normal NonStop system remote-access restrictions. The LOCALNAMES attribute, which can specify up to 30 remote nodes, can only be set when NETBATCH is run by the local super ID (255,255). You can set it either by altering the LOCALNAMES attribute within BATCHCOM or by specifying the remote nodes in the command to run NETBATCH.

HPE recommends not specifying LOCALNAMES unless the same people have SUPER.SUPER access to all of the listed systems, as in the case of SUPER.SUPER LOCALNAMES effectively gives SUPER.SUPER on the remote nodes as complete control of the local system as the local SUPER.SUPER has.

Note: Unlike most NonStop subsystems, ability to exercise a specific NetBatch command might be contingent on the user's access rights to a related file rather than whether the user belongs to the SUPER group.

Caution: NetBatch by policy considers ANY user with the access rights to alter a job's IN file as equivalent to the job owner, and permits that user to alter all of the job's attributes, effectively giving that user the same access rights on the system as the owner. Make sure that IN files are tightly secured.

Attachment-sets control the files on which your job operates. You should allow only users who are responsible for maintaining those file definitions to have write access to the attachment sets. Unless multiple users perform that maintenance, you should secure your own attachment-sets as "AOAO" or tighter.

See NETBATCH, Section 6, Command Security, for command-level security considerations.

Audit

HPE recommends enabling logging of EMS errors, at a minimum. Set the level to EMS ON only if you have monitoring software that is interested in seeing successful job completions.

NonStop Servlets for JavaServer Pages (NSJSP)

Because NSJSP is a web-based application, it is important to establish a secure link between it and its web browser clients, authenticate users, authorize their access to specific resources, and prevent malicious or erroneous code in a JSP from affecting the NSJSP container. These considerations apply to both the web applications themselves and their corresponding administrative interfaces.

Secure link

Since iTP WebServer is the front-end WebServer for NSJSP, it is the first line of defense. You should configure iTP WebServer to use TLS for its connections (see iTP WebServer, above).

Authentication: background

If an incoming request contains a user ID and password or X.509 certificate, iTP WebServer can pass these credentials on to NSJSP to use for authentication. NSJSP stores user credentials in a database called a Realm, along with information about roles (analogous to UNIX-style groups).

Authentication methods include:

- HTTP basic authentication (user name and password, sent unencrypted)
- HTTP digest authentication (user name and password, sent in encrypted (digest) form)
- Form-based authentication (application controls display of the login screen; successfully-authenticated user might be redirected to a different web resource based on assigned role)

- HTTPS client authentication (certificate-based authentication using HTTPS - HTTP over SSL)

NSJSP also defines a Java interface that supports authentication plug-ins. There are a number of standard plug-ins supporting various authentication methods, including a Lightweight Directory Access Protocol (LDAP)-accessible directory server, Java Naming and Directory Interface (JNDI) named JDBC DataSource, and Java Authentication and Authorization Service (JAAS) framework that includes a Safeguard interface.

NSJSP supports single sign-on through the SingleSignOn valve.

Authentication: best practices

You should not configure HTTP basic authentication. At a minimum, use digest-based authentication.

Decide whether your web applications need their own set of users or should support Safeguard users. Authenticating the application clients as Safeguard users removes the need to perform separate maintenance of a different set of user names and credentials, but you might not wish to allow your web clients to have OS-related IDs.

Use the NSJSPLockOutRealm to lock out users after multiple successive failed authentication attempts.

If you are hosting multiple web applications, consider enabling single sign-on.

Authorization: background

Security constraints, as defined through the security-constraint element in an application-specific deployment descriptor, control access to a set of web content. You configure web resource collections to define the HTTP methods (POST, GET…) and URL patterns to which a security constraint applies. You can apply authorization and/or user data constraints to web resource collections. Authorization constraints are based on roles, i.e. knowledge about an authenticated user. User data constraints apply requirements to the transport. Specifying either INTEGRAL or CONFIDENTIAL requires the request to be delivered over a secure transport such as HTTPS.

Sender validation

You can validate the origin of a client request prior to servicing the request through a valve element – either a Remote Host Filter valve, based on the sender's hostname, or a Remote Address Filter valve based on the sender's IP address.

HPE recommends that you configure client filtering in the iTP Secure WebServer rather than in NSJSP.

Java Security Manager

The Java Security Manager (JSM) can be used to restrict access to system resources such as Java Virtual Machine (JVM) properties, methods, and sockets. To use it, configure the iTP_catalina.policy file and include the appropriate JSM command-line arguments when starting NSJSP.

Caution: Make sure that you understand the impact of these security permissions before enabling them.

Caution: Ensure that you provide adequate package definition and access for NSJSP's internal packages through the use of the package.access and package.definition properties in the catalina.properties file.

Note: To assist in debugging JSM issues, you can use the java.security.debug property to write debug logs. Use the appropriate options to minimize unnecessary log generation.

Manager web application and NSJSP Manager security

Use the NSJSP Admin Web application to administer container objects and manage resources such as users and roles.

Use the NSJSP Manager Web application or NSJSP Manager to manage your deployed web applications. You can configure multiple NSJSP installations in one iTP WebServer installation, with each having multiple hosts and each host having multiple applications. You can constrain management scope by role through the definitions in the host-access.properties file.

The Realms repository is used to validate user credentials for these applications.

See NSJSPADMIN, Section 8, for details.

NonStop SOAP

HPE recommends that you deploy your server into an iTP WebServer instance running in a non-privileged or lightly-privileged account. Especially during development, using non-privileged iTP WebServer instances gives your developers the freedom to stop/start their own WebServers, which enhances their productivity.

NonStop SOAP can communicate with either a TS/MP server class or a NonStop process.

Digital signatures:

NonStop SOAP can be configured to use digital signatures to validate the sender’s identity and the message’s integrity. HPE recommends the use of digital signatures for NonStop SOAP requests and responses. Digital signature support exists in the NonStop SOAP with Digital Signatures product.

To use digital signatures, you must include the following directives in the configuration file:

- SOAP_KEYDB_LOC (file name of the server certificate and key location database)
- SOAP_DN (Distinguished Name value stored in the database)
- SOAP_WEB (file name of the iTP Secure WebServer configuration file)

See NSSOAPUSER, Section 7, for more details on digital signatures.

As of T0603H01^ACK (2011/09/22), the OSS installation directories for the product have permissions of 0755, and most of the individual OSS files have default permissions of 0644. If you are running an earlier version of SOAP 3, consider altering your existing directory and file permissions accordingly. You are responsible for appropriately securing the binaries and data files associated with anything you add to your SOAP deployment, for example User-Exits.

NonStop SOAP 4

NonStop SOAP 4 utilizes the Axis2c Rampart module to implement WS-Security Policy V1.1. The WS-SecurityPolicy V1.1 provides a set of standards for validating the security properties of a received message. The Rampart module provides an implementation of the following WS-Security standards:

- SOAP Message Security V1.0
- Username Token Profile V1.0
- X.509 Certificate Token Profile V1.0

See SOAPMSG, SOAPUSER and SOAPCERT for the WS-Security standards. See NSSOAPUSER41, Chapter 14, WS-Security in NonStop SOAP 4 for setup and usage details.

As of T0865H01^AAP (2013/07/08), the OSS installation directories for the product have permissions of 0755 and the individual OSS SOAP configuration files have 0644, the .h include files have 0444, the binaries and shell scripts have 0555 permissions. If you are running an earlier version of SOAP 4, consider altering your directory and file permissions accordingly. You are responsible for appropriately securing anything you add to your SOAP 4 deployment, such as user-written modules, handler, and Message Receiver User Functions (MRUFs).

NonStop Software Essentials (NSE):

NSE ships with a self-signed Jetty certificate, which will trigger browser warnings. Consider replacing the original Jetty certificate with one of your own. Instructions for either generating your own certificate or exporting an existing certificate from a keystore and installing the certificate in Jetty will be included in NSEINSTALL beginning with the 5.1 release. If you need instructions in the interim, contact the GNSC.

OSM: background

Access control

OSM versions T0682^ACJ and later provide ACL support, enabled by default, for more granular access configuration for OSM infrastructure commands. The ACLs are configured at a system level rather than for specific NonStop Consoles (NSCs), and configurations can be migrated to other systems. There are three default users (roles), and you can create additional, non-default users. All new users will need to be explicitly given OpenSSH permissions.

Prior to this version of OSM, there is one predefined set of actions that require SUPER group privileges and another set that can be performed by any user.

See OSMSCUG Section 8, Controlling Access to Actions, and OSMCG Section 5, Configuring Non-Default Users, for details.

Secure transport

Starting with T0682H02 AAM, OSM enables SSL by default.

Note: There are a number of considerations around certificate generation, contents and checking.

Note: If you have the HPE SIM client installed, which will be the case if you are monitoring the system with HPE Insight Remote Support Advanced and may also be installed for other uses, SSL must be enabled.

The terminal emulator shipped on NonStop Consoles (NSCs), MR-Win6530, has built-in SSH support. MR-Win6530 is used by the OSM Low-Level Link during system startup and by the OSM Service Connection for TACL and FTP sessions.

See OSMCG Section 3, OSM Server-Based Components, and Section 4, OSM and Other HPE Client-Based Components, for more details.

Audit

User-initiated OSM actions are audited in the $SYSTEM.ZSERVICE.ZTRC* files.

OSM persistent process security

Starting with T0682H02 ACE, OSM sets the default process file security for its persistent processes to NCNC and allows you to alter the level if desired.

Interactions with TACLCSTM, CMON and XAC

Caution: Some OSM processes start non-interactive TACL processes, which can cause problems with TACLCSTM scripts that expect human interaction. There also can be issues with tight $CMON or XYGATE Access Control (XAC) scripts.

See OSMCG Section 3, Adding OSM Process Files to TACLCSTM and CMON Exceptions, for details.

Idle Event Viewer sessions

Starting with T0682H02 ABP, the Event Viewer will time out and eventually delete sessions after specific periods of inactivity.

OSM: best practices

Secure transport

If you are running a predecessor to T0682H02 AAM, SSL is not enabled by default; you should enable it.

See *Generating Certificates* in the NonStop SSL section for certificate best practices.

If your NonStop Consoles (NSC) are running Windows Server 2008 rather than Windows XP or Windows Server 2003, consider explicitly configuring use of the stronger AES ciphers in place of the older and less secure MD5, SHA and 3DES defaults.

The administrator and non-administrator private keys for logon to the NSC are stored in $SYSTEM.ZSERVICE. By default they are owned by SUPER.SUPER, with the administrator private key (NSCSUPER) and administrator public key (NSCSKEY) secured as CCCC and the user private key (NSCUSER) and user public key (NSCUKEY) secured as NCNC. These settings allow any SUPER group member to issue commands with administrator privilege and any local or remote user to issue commands with user privilege.

Caution: Do not loosen the administrator private key's read security setting; otherwise, you will allow non-SUPER-group members to execute commands on the NSC with administrator privilege.

Idle Event Viewer sessions

If you are running T0682H02 ABP or a successor, you can configure idle session expiration and deletion in accordance with your corporate policy.

See OSMCG Section 3 for more details.

SNMP Read/Write Access for UPS and Maintenance Switch

Modify the default community strings and their associated access privileges.

See OSMCG Appendix C, Configuring SNMP Access for Monitored Service LAN Devices, for more details.

Samba

HPE recommends not running Samba on a production system unless you have analyzed whether it is capable of being configured to support all of that system's security policies. Samba has a separate user/password management scheme with fewer capabilities than Safeguard, insecure configuration options such as the SWAT demo mode, and writes to its own audit logs rather than a central log such as EMS.

If you do use Samba, configure its user and password management policies as similarly as possible to your Safeguard policies. Modify its file permissions and related settings as needed, restrict administrator access, configure SWAT to use SSL and to not be runnable in demo mode, and configure log management appropriately.

See SAMBA Section 4, Security Considerations, for details.

SCF

SCF is used to manage multiple subsystems, including storage, comm, and the kernel (OS). As mentioned earlier, SCF has a straightforward set of roles with non-SUPER-group users generally able to issue only nondestructive commands, SUPER.notsuper able to issue the set of commands required for day-to-day operations, and SUPER.SUPER as the only user who can issue a few sensitive commands.

Persistent processes

It is possible to allow one or more users other than SUPER.SUPER to have full control over persistent process management via $ZPM by defining and populating the SECURITY-PERSISTENCE-ADMIN group. HPE recommends not using this feature unless necessary. See SCFKERNEL Section 6, SCF Commands for the Kernel Subsystem, for details.

Spooler

Do not run Spooler instances as SUPER.SUPER.

Control command-level access through either XAC or a partner product. Restrict the ability to add new print processes, which run as the same user as the Spooler instance. At a minimum, use Safeguard ACLs to restrict access to the supervisor and collector processes.

TACL

Do not hard-code passwords in TACL macros. Use #INPUT/NOECHO/… for password entry.See TACLREF Section 9, #INPUT Built-In Function.

Consider using #SET #TACLSECURITY "OO" to tightly control who can open a TACL process and prevent arbitrary users from using the ENQUIRY feature to access recent data written by a TACL process to its OUT file. The default value of the #TACLSECURITY built-in variable is "NN".

See TACLREF Section 9, #TACLSECURITY.

TMF

TMF configuration information is stored in the ZTMFCONF subvolume of the configuration volume (default is $SYSTEM). If the subvolume is not $SYSTEM, a file in $SYSTEM.ZTMFCONF points to the configuration volume. Do not alter the security of these files.

TMF audit trails reflect changes to user data, and have the same security considerations as the underlying data – so protect them accordingly. Make use of Safeguard ACLs where appropriate. Provide similar protection for TMF disk dumps, and properly control TMF tape dumps. Consider encrypting tape dumps.

TMF follows the typical security model with informational commands available to all users, most sensitive commands available to all members of the SUPER group, and the most destructive commands (DELETE CATALOG, DELETE TMF, DELETE TRANSACTION, and ALTER PROCESS with the DEBUG attribute) restricted to SUPER.SUPER. DISPLAY OPERATIONS is available to any user with read access to EMS log files.

See TMFREF Section 1, Table 3-1 for details.

Many TMF programs are shipped as licensed, and you should follow the same pattern on your system with the exception of SNOOP and SNOOPDR. If SNOOP or SNOOPDR is delivered licensed, revoke the license. SNOOP should be used only with guidance from the GNSC or NonStop development, and all usage should be logged – preferably at the keystroke level – as it can be used to view audit trail contents and alter its attributes.

TS/MP

PATHWAY configuration (per PATHMON)
The OWNER attribute specifies the owner of the PATHMON environment. That user can stop the PATHMON process, add programs, delete objects, and make other modifications to the global configuration.

The SECURITY attribute, whose value is relative to the OWNER's user ID, specifies who else can modify the PATHMON environment after you issue the START PATHWAY command. The default is "N" for TS/MP 2.0 and 2.1, and should be changed to "O". The default is "O" for TS/MP 2.3 and later versions. For all TS/MP versions, prior to issuing the START PATHWAY command the owner is the process access ID (PAID) of the PATHMON process and the SECURITY attribute is O (owner only).

HPE strongly recommends that the OWNER of a given PATHMON environment be a user ID associated with management of that specific application, rather than SUPER.SUPER or a member of the SUPER group.

SERVER configuration (per server class)

OWNER specifies the user ID that controls access from a Pathsend process to a specific server class. (The TCPs ignore this server attribute.) If not specified, OWNER defaults to the user ID who started the PATHMON process. Specify an appropriate user.

SECURITY specifies the users, in relation to the OWNER attribute, who can access a server class from a Pathsend requestor. (TCPs ignore this attribute.) The default is "N", which should be changed to the most restrictive setting that does not interfere with correct application operation.

For Guardian servers, ensure that server class ASSIGNs and DEFINEs point to the appropriate files, and that the server class volume is explicitly set. Similar considerations apply to CWD and ARGLIST for OSS servers.

For OSS server classes, UMASK specifies the default permissions for the owner, group and others for OSS files created by the server process instance (see General OSS file security for more information on umask). The default for UMASK is -1; HPE recommends that you set it to a more restrictive value (022 or tighter).

Network security

The PATHMON process controlling the server class has to have corresponding user IDs and remote passwords with all of:

- The system where the requesting process is running
- The system where the PATHMON process is running
- The system where the server class is running

This level of security is required because the LINKMON process or the ACS subsystem processes must be able to open the PATHMON process (to make link requests); the LINKMON process or the ACS subsystem processes must be able to open the server processes (to send user requests); and the PATHMON process must be able to open the server processes (to send startup messages). All of these opens are performed with the PATHMON user ID.

Note: If the user starting the PATHMON process is an alias, then the alias must have matching remote passwords on all involved systems. It is not sufficient for the underlying user ID to have matching remote passwords.

Note: The user ID of the Pathsend process need not have remote passwords to the PATHMON system or to the server-class system to access the server class. Moreover, the Pathsend-process user ID need not be known on the PATHMON or server-class systems.

Server-class security

LINKMON processes or ACS subsystem processes perform authorization checks on each server-class send operation to make sure that the user ID of the Pathsend process at the time of the send conforms to the server class's OWNER and SECURITY attributes. You set these attributes for server classes at configuration time if those server classes are to be accessed by Pathsend processes.

See TSMPMGT and TSMP25MGT, section 3, Configuring Objects in a PATHMON Environment, section 6, Managing the Pathsend Environment, and section 12, SERVER Commands, and TSMPPATH and TSMP25PATH for more details.

Pathway/iTS considerations

Pathway/iTS OWNER and PROGRAM configuration (per TCP)
You can specify security for a PROGRAM object by setting the OWNER and SECURITY attributes of the SET PROGRAM command. The OWNER attribute specifies the owner of the current PROGRAM object. The owner can alter the security attribute for the PROGRAM object. The SECURITY attribute values are the same as those for Guardian operating environment security attributes. The SECURITY attribute specifies who can run the PROGRAM object. The OWNER attribute defaults to the user ID of the user who started PATHCOM, and SECURITY defaults to “N”. The recommendations for TS/MP OWNER and SECURITY attribute settings apply here as well.

The following example specifies that only the owner—user ID 8, 61—can run this PROGRAM object:

```
= SET PROGRAM OWNER 8,61
= SET PROGRAM SECURITY "O"
```

The next example specifies that any local or remote user can run the PROGRAM object:

```
= SET PROGRAM OWNER 8,61
= SET PROGRAM SECURITY "N"
```

File security

The POBJCOD and POBJDIR files contain the individual SCOBOL programs, and access should be limited to those users who are authorized to change them.

Communication between PATHMON environments

When one PATHMON environment communicates with another, the PATHMON process for each environment controls and reports on its own objects. When a TCP in one PATHMON environment executes a screen program request directed to a server class controlled by another PATHMON process, these operations occur:

1. The TCP requests a link from the external PATHMON process that controls the server class.
2. The external PATHMON process grants the link to the server process.
3. The TCP uses the link to open the remote server process and send the request to that server process.
4. After the remote server process completes its work, it returns its reply to the TCP.

The same general operations take place whether the PATHMON environments are running on the same or on different nodes.

For communication between PATHMON environments, the TCP requesting a server link must pass all process security checks. If the TCP and server class are on different nodes, the security requirements demand the following:

- The user who starts the local PATHMON process must possess a remote password for the external PATHMON process’s node.
- The user who starts the external PATHMON process must possess a remote password for the local PATHMON process’s node.

This level of security is required because the local TCP must be able to open the external PATHMON process and server classes, and the external PATHMON process must be able to open the local TCP. All TCPs and server processes started by a PATHMON process run under the user ID of the person who started that PATHMON process.

See PATHITSMGMT, Sections 2 and 11, for further details.

Your Pathway applications should authenticate their users, preferably using NonStop server user IDs and passwords. If they maintain their own authentication routines and set of users, they must take responsibility for protecting passwords (including encrypting them on disk), enforcing password quality, and auditing authentication attempts. They also should protect themselves against unexpected open requests.

If possible, your communicating Pathway environments should run as users in the same group; otherwise Pathway security would need to be set to “A” or “N” to allow communication.

There are partner products that can provide granular control of which users’ programs can issue PATHSENDs to a Pathway server. This is particularly useful in situations where multiple PATHMON environments must communicate, but they are running under different user groups.

Remote system access

TCP/IP

The primary concerns for TCP/IP are configuration security, network administration, and denial-of-service attack prevention and detection.

All potentially sensitive data being transported across TCP/IP should be encrypted. See SSH, SSL and IPSec, below, for more information.

Do not run unneeded services. ECHO and FINGER are examples of services that are unlikely to be needed. If a service is needed and there is a choice between less-secure versions and more-secure versions (e.g., iTP WebServer and iTP Secure WebServer, inetd and xinetd, dns and dnssec), use the more-secure version.

Use a third-party tool such as nmap to check which ports actually are open, and map those against your list of needed ports. If possible, close unneeded ports in a front-end firewall, IP CLIM (iptables or ip6tables), NonStop SSH, NonStop SSL, or a web application firewall. If a service uses a well-known port that is configurable, consider using a different port in its place. Disable unneeded ports in the inetd or xinetd configuration file.

Configuration file security

In the Guardian environment, secure the following TCP/IP configuration files so they can be written or purged only by users who have the authority to reconfigure the TCP/IP subsystem.

The host definition file, $SYSTEM.ZTCPIP.HOSTS (or configured location)

The IPv6 host definition file, $SYSTEM.ZTCPIP.IPNODES (or configured location)

The domain name resolver configuration file, $SYSTEM.ZTCPIP.RESCONF

The networks definition file, $SYSTEM.ZTCPIP.NETWORKS (or configured location)

The services definition file, $SYSTEM.ZTCPIP.SERVICES (or configured location)

The LISTNER port configuration (PORTCONF) files. The primary PORTCONF file is in $SYSTEM.ZTCPIP, but can be reconfigured. If you configure different LISTNERs with different functions, you should create and configure a separate PORTCONF file for each LISTNER. Document the pairings. You may wish to place your PORTCONF files under source control.

Note: $SYSTEM.ZTCPIP.SMPLNETW, the sample NETWORKS file, is overwritten whenever a new version of TCP/IP is installed. Copy it to $SYSTEM.ZTCPIP.NETWORKS when doing your initial configuration, and make all changes in NETWORKS.

In the OSS environment, /usr/ucb/inetd (or xinetd) provides equivalent functionality to the LISTNER process.

Secure /etc/inetd.conf, or your custom config file if using one, to restrict write access to the subsystem owner, and do the same for its other files:

/usr/ucb/inetd/

/etc/hosts

/etc/networks

/etc/services

/etc/protocols

Note: The hosts, networks, services, and protocols files in OSS are initially set up as links to the Guardian TCP/IP configuration files. You could create OSS versions, but HPE recommends using links instead so the configurations stay in sync in the two environments. With this approach, the configurations are controlled using only the Guardian environment, as the links cannot be changed in OSS.

HPE recommends using xinetd in place of inetd, as it has better security. Allow inetd (or xinetd) and LISTNER to listen only on ports needed for your system and approved by your security auditors.

Caution: If a LISTNER process and inetd (or xinetd) are configured to use the same TCP/IP process and port number, collisions may occur.

Network administration

Only enable ports in PORTCONF or inetd.conf that are actually needed.

Like all services that use ports 0-1023, LISTNER, inetd and xinetd processes must be started by a SUPER group member.

HPE recommends not having LISTNER processes be started by SUPER.SUPER. Programs started by the LISTNER inherit the LISTNER's CAID and PAID and associated privileges. Any program started by a LISTNER should always authenticate its users.

Session termination

Where possible, configure idle sessions to be timed out and disconnected.

Several programs provide session timeout options:

NonStop SSH: By default, NonStop SSH does not disconnect inactive terminals. To establish an idle timeout, set the INPUT_TIMEOUT and BANNER_TIMEOUT parameters to the number of minutes of inactivity to trigger disconnection. The minimum value is 3 minutes.

NonStop SFTP: The SFTPIDLETIMEOUT parameter controls how long SFTPSERV keeps running without any SFTP protocol traffic before terminating itself. The default is that the parameter is omitted and there is no idle timeout; configure SFTPIDLETIMEOUT to establish a timeout value.

Telserv: By default Telserv does not disconnect inactive terminals. To establish an idle timeout, use SCF ALTER PROCESS to alter TIMEOUTVALUE and BANNERTIMEOUT to the number of minutes of inactivity to trigger disconnection. The minimum value is 3 minutes.For details, see SSHREF, STN Reference, STN Commands and TELSERV Part II, Section 4, SCF Commands, ALTER PROCESS.

Denial-Of-Service (DOS) attack prevention and detection

If your system has IP CLIMs, you can establish iptables or ip6tables restrictions on inbound traffic.

See CIPMGMT, Section 3, CIP Configuration and Management, Configuring CIP iptables/ip6tables (IP CIP), for additional details.

TCP/IPv6 supports the ICMP-FILTER-PKTS Monitor attribute to control the flow of ICMP packets into the system. CIP supports ICMP packet filtering through iptables and ip6tables. Currently, disabling ICMP redirect messages is not a default setting on the CLIM, as there are routing configurations which legitimately need to advise network endpoints of better routes. This functionality is a standard part of Internet Protocol, as defined by IETF RFCs and must be supported for use by some HPE customers. Disabling this capability by default would break backward compatibility with earlier NonStop products and versions. At the same time, we support disabling ICMP and other iptables functionality for customers who want additional security constraints limiting access from the directly connected LAN.

For more information on various kinds of ICMP attacks, see ICMPATTACK.

See TCPIPV6MGMT, Section 8, SCF Reference for NonStop TCP/IPv6, ALTER MON Command for TCP6MAN for details on TCP/IPv6 ICMP packet filtering.

Network traffic segregation

If your system has IP CLIMs, the Multiple Providers on a CLIM (MPC) feature allows you to segregate network traffic. Conventional TCP/IP and IPv6 both support an equivalent feature, Logical Network Partitioning (LNP). Parallel Library TCP/IP does not support traffic segregation and, if possible, should not be used if you are supporting multiple applications on a single system.

See CIPMGMT, Section 3, CIP Configuration and Management, Setting Up Multiple Providers per CLIM, for additional details.

Host server processes: Guardian

ECHO

To prevent denial-of-service attacks, ECHO should be removed from the PORTCONF file unless there is an absolute requirement for its use. If it must be used, reconfigure it to not use the default port 7.

FINGER

To prevent leaking information about users logged onto the system and IP addresses, FINGER should be removed from the PORTCONF file unless there is an absolute requirement for its use. If it must be used, reconfigure it to not use the default port 79.

FTP (File Transfer Protocol)

To prevent transmission of sensitive data, including user credentials, in the clear – and a variety of other attacks - FTP should be removed from the PORTCONF file unless there is an absolute requirement for its use. Otherwise, use SFTP instead (see below for further information.)

If FTP must be used, reconfigure it to not use the default port 21 and remove it from the PORTCONF files of any subnets that do not require that service. Take advantage of FTP's configuration options to limit access. See Disallowing Logons in TCPAPPUTILS Section 7, FTP - Communicating with the FTP Server, for details.

Do not allow anonymous users unless absolutely necessary; if you must allow them:

- Create an alias to NULL.FTP to be used for anonymous access, freeze the underlying NULL.FTP user ID, expire its password, and assign it a safe default volume and subvolume to NULL.FTP and all of its aliases.
- Use partner products to prevent users from opening an FTP session.

Use a Safeguard protection record to restrict which users can run FTP, and volume/subvolume ACLs and/or appropriate OSS security to restrict file and directory access. See TCPAPPUTILS Section 15, Anonymous FTP, for details.

PING

To prevent denial-of-service attacks and confirmation of connection, external firewalls should restrict which devices on the network are allowed to send Internet Control Message Protocol (ICMP) packets to the system. PING uses ICMP packets.

TELNET (Telserv)

To prevent transmission of sensitive data, including user credentials, in the clear – and a variety of other attacks – Telserv should not be used unless there is an absolute requirement for it. If it is needed, front-end it with an SSL proxy, See SSLREF, Introduction, Secure Proxy for Telnet Access and "Blocking access to the 'plain port' of encrypted protocol" under NonStop SSH and NonStop SSL: best practices, below, for details.

If use of TELNET/Telserv is not required, use SSH instead (see below for further information.)

When configuring services, specify OWNER to cause the client to be authenticated and access to the service enforced by the Login Server. For details, see SERVICE configuration in TELSERV Section 5.

Audit

Set PARAM ZTNT^LOG^CONNECTS to YES to log client connections and disconnects to $0. If your connect rate is heavy, consider logging to an alternate connector instead of $0. Also consider using ZTNT^LOG^SERVICENAME to log the service name in connect and disconnect EMS event messages in the case of login failure.

TFTP (Trivial File Transfer Protocol)

To limit public access to data, TFTP should be removed from the PORTCONF file unless there is an absolute requirement for its use; otherwise, use SFTP instead (see below).

TFTP is UDP-based, with no authentication. Users can GET files secured “Nxxx”, and PUT files at locations where they have remote WRITE access “xNxx”. Configure the subvolume where remote users will be allowed to PUT files, and use Safeguard ACLs to restrict network access.

Host server processes: OSS

Echo

See considerations for ECHO, above.

Note: echo conflicts with the GUARDIAN ECHO; if you must run one, use the Guardian version.

Finger

See considerations for FINGER, above.

ftpserv(7)–FTP server

HPE recommends use of SFTP (see below) in place of FTP. If you must use the OSS FTP server, consider specifying SITE CHMOD to set the default security for transferred files to something tighter than 666 (-rw-rw-rw-). The FTP server uses normal Guardian authentication (USER_AUTHENTICATE_).

OSS uses the FTPSERV running under the LISTNER. Your view of the world (OSS or Guardian) is determined by the setting of your initial-directory. You can use the quote GUARDIAN and quote OSS commands to change your view.

See the considerations under FTP, above.

inetd(8) – Internet superserver

HPE recommends use of xinetd (see below) in place of inetd.

If you do use inetd, it should be run as a user other than SUPER.SUPER (specified by UserName in the config file). Use the –R option when starting it to reduce the chance of denial-of-service attacks.

rexecd(8)

Unless there is a specific reason to enable it, HPE recommends disabling rexecd by commenting out or deleting its definition in /etc/inetd.conf.

rndc(8), dnssec_named(8), named(8)

If the use of DNS is required, use the secure version and secure named.conf to restrict write access to the subsystem owner.

rpcinfo(8)

RPCINFO can modify the status of RPC services available from the local node, so execute access should be restricted appropriately. Write access to its program definition file, $SYSTEM.ZRPC.RPC, also should be restricted appropriately.

rshd(8)

Unless there is a specific reason to enable it, HPE recommends disabling rshd by commenting out or deleting its definition in /etc/inetd.conf.

See OSSSHELL, section 12, for further details.

Use SSH instead of Telserv, SFTP instead of FTP

Do not run LISTNER as SUPER.SUPER.

Known used ports

See SECMGMT Appendix E, HPE NonStop Port Details, and Console Security (below) as well as individual product manuals for information on known used ports.

TLS/SSL and SSH configuration: overview

Many products used in the NonStop server environment include TLS/SSL support. Collectively, they use a number of different code bases, which means that their configuration methods also vary. The individual product manuals are the ultimate configuration authorities, but you can find a consolidated overview of configuration of TLS/SSL protocol versions in the TLSCONFIG white paper in the manuals collection.

Configuring SSH or TLS/SSL can be complex, and configuration mistakes can introduce security weaknesses. When in doubt, get expert help. To help customers configure NonStop SSH and NonStop SSL correctly and securely, HPE offers services that include education on data in motion protection options and configuration, TLS/SSL certificate creation and management, and security performance tuning. Your company's security management group also can be a useful resource.

As a general security best practice you should disable remote access to all unused ports in your firewall.

Especially for privileged users, additional security can be achieved if full keystroke logs are captured for fraud detection and forensic analysis through XAC or a partner product with similar capabilities.

Choosing which protocol(s) to use

Where you have a choice, HPE recommends SSH as the configuration of “trust” between client and server is more automatic in SSH than with TLS/SSL. Further points to consider:

- Does your company have a policy about or a preference on the use of TLS/SSL vs. SSH?
- What types of systems do you exchange data with? The peer systems might not have both protocols available, or their security administrators might prefer the use of one or the other.

- Which NonStop products are you using? Some products have built-in support for one or both protocols.

NonStop SSH configuration

The considerations below apply to both NonStop SSH and cF SSH-Lib.

Blocking access to the “plain port” of the encrypted protocol
NonStop SSH replaces both FTP and Telnet (Telserv). If you still are running versions of FTP or Telserv that are not front-ended by NonStop SSL, you should disable these insecure original services as the last step in the migration from using insecure to secure connections.

Auditing

NonStop SSH can write both log messages (informational from an operations viewpoint) and audit messages for security-relevant events.

Audit messages are, by default, not written to either a disk file or a device/console. You can send them to the EMS subsystem by configuring AUDITCONSOLE to write them to either $0 or a specific device, and AUDITFILE to write them to disk. You can configure disk file audit rollover behavior using LOGFILERETENTION. AUDITMAXFILELENGTH controls the maximum audit file size, and defaults to 20MB.

AUDITLEVEL controls what audit messages are generated. The default level includes user logons and network events. You can adjust the level to capture additional data, up to and including a full byte dump.

You should secure the audit log files appropriately to prevent tampering; if you are logging at the byte dump level, also configure them to prevent read access as they will contain passwords and potentially-sensitive user data.

You can configure LOGCONSOLE to write messages to either $0 or a specific device, LOGEMS to write them to a specific EMS collector, and LOGFILE to write them to disk. HPE recommends writing log events to either a file or EMS. You can configure disk file log rollover behavior using LOGFILERETENTION. LOGMAXFILELENGTH controls the maximum audit file size, and defaults to 20MB.

You can set a generic value for what gets logged using LOGLEVEL, or tailor it with LOGLEVELEMS, LOGLEVELCONSOLE and/or LOGLEVELFILE. For normal usage, up to level 50 is appropriate; you might need to enable higher levels for troubleshooting at HPE’s request.

The SSH log and audit files are unstructured files with file code 0. They can be viewed with any PC editor if and only if they are transferred to the PC in binary. The SHOWLOG tool allows you to view the logs and audit on the NonStop system. In addition, in the OSS environment you can use more and tail, possibly in conjunction with filter tools such as grep, awk or wc.

You should review the access logs regularly – ideally as part of an automated process where the logs are fed into a central SIEM system.

Environmental considerations

If you use the TCP/IP round robin capabilities, make sure the value of your PTCPIPFILTERKEY adheres to common best practices for secure passwords and that it is protected in the CONFIG/CONFIG2 file from unauthorized access. If the PTCPIPFILTERKEY is known to an attacker, he or she might start a process listening on the same port as the NonStop SSH processes, subverting the security layer. For TCP/IPv6, the risk of someone starting a rogue service listening on the same port can be further decreased if you start one instance of your NonStop SSH process on each CPU, since only one process can be started per CPU with that specific PTCPIPFILTERKEY. When using CIP and IP CLIMs, parallel round robin filtering allows you to start multiple listening processes in the same CPU that share the same port.

Staying up to date on protocol or implementation vulnerabilities

SSH is visible and accessible for anyone with network access to your NonStop system. Therefore, it is important that you stay up to date on installing fixes or mitigations for protocol and/or implementation vulnerabilities. The SSH protocols does evolve over time to make sure it can withstand attacks. Make sure you pay attention to HPE's security update notifications and that you follow their guidelines for addressing identified vulnerabilities.

NonStop SSH encryption: background and best practices

HPE NonStop SSH configuration is a highly security-critical task as HPE NonStop SSH tightly integrates into the system. An improper configuration may lead to severe security weaknesses.

In particular, it is important to understand that the management of HPE NonStop SSH users is as critical as the user management within Safeguard itself since a user that is allowed to configure SSH USER records can create access to the NonStop system without Safeguard authentication.

Lock down access to sensitive files

To prevent unauthorized access trying to read and/or modify files, the HPE NonStop SSH files must be secured properly, e.g. by using Safeguard ACLs. In particular, files to be protected are:

- The program files themselves

  This also includes restricting access to SSH/SFTP client program files to only those users actually needing them to perform their job functions.

- The HOSTKEY file

  Contains the private key of the SSH server. The default key type and size are now set to RSA 2048 bits.

- The configuration files

  In particular, the files referred in CONFIG or CONFIG2 and, for SFTPSERV, files SFTPCONFIG or SFTPCONFIG2.

  Note: All parameters specified in the CONFIG2 file will have precedence over the ones specified in CONFIG.

- The SSHCTL file

  This is the user database

- The log and audit files

  All log and audit files should be locked down to prevent access, in particular to prevent any modification. Keep in mind that log and audit files will be rolled over, thereby creating file names that do not yet exist on your system but still have to be thought of in advance when locking down the system.

  See the corresponding parameters AUDITFILERETENTION, AUDITMAXFILELENGTH and LOGFILERETENTION, LOGMAXFILELENGTH for details.

“Real time” alerting should be configured with XMA or an alternative partner product to alert on any violation of access in order to quickly notice potential intrusion attempts.

Lock down user access

Authentication methods

HPE NonStop SSH supports several user authentication methods. For stronger security, methods that do not rely only on a user password should be used. HPE recommends using the “publickey” method rather than password-based authentication methods (“password”, and, unless a SEEP is configured that checks other authentication factors than the password, “keyboard-interactive”). See the description for ALLOWED-AUTHENTICATION / CLIENTALLOWEDAUTHENTICATIONS in SSHREF for details.

By using publickey authentication, further restrictions can be added by configuring the restriction profile name in the PUBLICKEY attribute RESTRICTION-PROFILE to designate an existing restriction profile.

If you do use password-based authentication methods, then you should use only SSH clients that support the handling of password expiration. Password expiration handling in the SSH protocol is defined for both methods “password” (RFC 4252) and “keyboard-interactive” (RFC 4256). However, in practice support for resetting the password via method “password” is far less widespread in SSH clients than doing so via “keyboard-interactive”. If in doubt, HPE recommends using the “keyboard-interactive” method instead of “password".

Strong authentication is supported via multi-factor authentication, allowing you, for example, to enforce use of both a successful password-based method and a successful non-password based method before access is granted.

See the description for USER attributes ALLOWED-AUTHENTICATIONS and REQUIRED-AUTHENTICATIONS, SSH2 parameter CLIENTALLOWEDAUTHENTICATIONS and ssh/sftp client option -oAllowedAuthentications in SSHREF for details.

Access to SSH subsystems

Use the ALLOWED-SUBSYSTEMS property in the HPE NonStop SSH USER record to restrict each user to have access only to those subsystems (TACL, SFTP) he or she is allowed to access.

For users allowed to access the SFTP subsystem (file transfer), restrict capabilities on a need-to-have basis via the SFTPSECURITY parameter.

Use the CLIENTALLOWEDSUBSYSTEMS parameter for allowing outgoing connections. The subsystem(s) listed in this parameter can be supplied as a single value or a comma-separated list.

Use the ALLOW-SHELL parameter to indicate whether the SSH user is allowed to request a shell. Also use ALLOW-PTY to indicate whether the SSH user is allowed to request a pseudo terminal (PTY).

Use the special command “ci” in conjunction with ALLOW-CI, CI-PROGRAM, CI-COMMAND and CI-PROGRAM-OVERRIDE and ALLOW-CI-PROGRAM-OVERRIDE in cases when users should not get general TACL access (TACL not in ALLOWED-SUBSYSTEMS) but just need to execute a specific TACL command, access a specific Guardian subsystem, or execute a specific TACL macro.

Use of application service users

If you have an application accessible over SSH (e.g. a PATHWAY application) that includes authentication itself, create a dedicated SSH service user that serves as the entry point for all users that want to reach it. Lock this service user down completely along the lines of the following user properties:

- SYSTEM-USER *NONE*
- CI-PROGRAM *MENU* <THE ONLY APP> FORCE
- ALLOWED-SUBSYSTEMS ()
- ALLOW-SHELL NO
- ALLOW-PTY NO
- ALLOW-TCP-FORWARDING NO
- ALLOW-GATEWAY-PORTS NO
- ALLOW-MULTIPLE-REMOTE-HOSTS YES

- RESTRICTION-PROFILE <app user restriction profile>

Note: The first statement prevents any actual system user from being assigned to that specific application service user.

If the target application does authentication on the application level, then and only then ALLOWED-AUTHENTICATIONS can be set to “none” for that particular application service user. Except for those fully locked down SSH application service users, do not use authentication “none”, because this means that no authentication will be performed on the SSH level.

If multiple services are required, then use *MENU* and USER attribute PTY-SERVER together with multiple differently configured PTY servers (not just the default one configured via parameter PTYSERVER) offering a user/group specific service menu.

Further Considerations

Do not allow access if the corresponding system user is frozen by setting “ALLOWFROZENSYSTEMUSER FALSE”.

Control whether users are allowed to issue commands ADD KNOWNHOST and ALTER KNOWNHOST by setting ALLOWADDINGKNOWNHOST to one of {ALL | PARTIALSSHCOMACCESS | FULLSSHCOMACCESS }.

Control whether users are allowed to issue commands IMPORT KEY, GENERATE KEY and ALTER KEY by setting ALLOWADDINGPRIVATEKEY to one of {ALL | PARTIALSSHCOMACCESS | FULLSSHCOMACCESS}.

Define the set of users that are allowed to execute the SSHCOM command INFO SSH2 by setting ALLOWINFOSSH2 to one of {ALL | PARTIALSSHCOMACCESS | FULLSSHCOMACCESS}

Prevent the SSH2 daemon from listening on the ANY address by setting “ALLOWLISTENONANYADDRESS FALSE”

Disable TCP/IP port forwarding by default by setting “ALLOWTCPFORWARDING FALSE”.

Configure every user with a RESTRICTION-PROFILE and enforce denial of access if no record exists by setting the parameter RESTRICTIONCHECKFAILEDDEFAULT to TRUE.

Lock down management user access

Restrict privileges to invoke (sensitive) SSHCOM commands to only those users who must have the ability to invoke the commands. This can be achieved with the help of the parameters ALLOWINFOSSH, FULLSSHCOMACCESSUSER<i>, FULLSSHCOMACCESSGROUP<j>, PARTIALSSHCOMACCESSGROUP<n>, and PARTIALSSHCOMACCESSUSER<k>.

Please also refer to SSHREF section “Security within SSHCOM”.

Understand the security implications of the CLIENTMODEOWNERPOLICY as described in the “Security within SSHCOM” chapter of SSHREF. In particular, if alias names are used and different aliases for a Guardian user are assigned to different persons, then use the value LOGINNAME to prevent one person from accessing the client mode records of another person using an alias with the same underlying Guardian user.

The capability of automatic adding of system users via AUTOADDSYSTEMUSERS can be very helpful, especially in a migration.

Note: This capability is only possible for SSH sessions that use authentication method “keyboard-interactive” or “password”. Automatic addition will fail for new users who are not allowed to use at least one of these two authentication mechanisms.

Logging and Auditing

Use the logging and auditing capabilities (parameters LOG* and AUDIT*) to track activity as described in the general section for both SSH and SSL.

Especially for privileged users, additional security can be achieved if full keystroke logs are captured for fraud detection and forensic analysis through XAC or a partner product with similar capabilities.

Security Parameters

Use the CIPHERS and CLIENTCIPHERS parameters to restrict the ciphersuites to only the most secure ones both peers support. If cipher algorithm compatibility in peer systems allows, use the strongest cipher(s), which is currently “aes256-ctr”. The use of CTR ciphers is recommended and the use of CBC ciphers discouraged by the US National Institute of Standards and Technology (NIST). In particular, do not use CIPHER arcfour (=RC4) unless absolutely necessary.

Similarly restrict the allowed message authentication codes (parameter MACS/CLIENTMACS) to the ones with the highest available security, currently “hmac-sha2-512”. Do not use any md5 based HMACs unless absolutely necessary. If you currently use SHA1-based MACs, plan on migrating to the stronger SHA2-based MACs in the near future.

If you use public key authentication, force rollover of user private keys with the help of the parameters INTERVALLIVEPRIVATEUSERKEY and INTERVALPENDINGPRIVATEKEY.

The recommended LIFECYCLEPOLICYPRIVATEUSERKEY setting is FIXED. For incoming connections, LIFECYCLEPOLICYPUBLICUSERKEY is of more importance. See the description in SSHREF for details.

HPE recommends the use of the stronger diffie-hellman-group-exchange-sha256 as the value in the ALLOWEDKEYEXCHANGEALGORITHMS parameter for configuration of key exchange algorithms. Use values for the SSHAUTOKEXBYTES and SSHAUTOKEXTIME parameters, which determine frequency of automatic key re-exchange in SSH sessions that fit your security policies and performance profile.

In order to not allow SSH client program users to connect to a target host that has a changed host identity (which can be an indicator of a man-in-the-middle attack), enforce STRICTHOSTKEYCHECKING.

Global restriction profiles can be configured for all incoming and outgoing connections. The RESTRICTIONPROFILE parameter supports both white listing and black listing options. See the description for RESTRICTIONPROFILE in SSHREF for details.

Configure the MAXAUTHTRIES parameter that specifies the maximum number of authentication attempts permitted per connection to a low value.

Similarly, raise the security awareness of users connecting in to the NonStop via NonStop SSH as the server so they will take client emulator messages warning of a changed SSH server key seriously, since it could be a man-in-the-middle attack.

Note: Many of the above security considerations also apply to OpenSSH ports to NonStop servers.

NonStop SSL: best practices

The considerations below apply to the entire suite of NonStop TLS/SSL products, including NonStop SSL, cF SSL-Lib, and cF SSL-AT.

Blocking access to the “plain port” of the encrypted protocol

Once protocols such as Telnet, FTP, and ODBC all have been encrypted with TLS/SSL you should block remote access to the “plain port” of the underlying protocol. This can be either done via an external firewall or through proper configuration of the product. Please contact HPE support for details on the latter.

In cases where you are using the NonStop SSL proxy (i.e. NonStop SSL or SSL-AT in non-native mode), you need to block remote access to the underlying “plain” port as the program running behind the proxy has that port open for the proxy to communicate with it. This is not the case for programs that have a native TLS/SSL library built in (i.e. SSLAT native-mode or SSLLIB), as they will not open the plain port to begin with.

Auditing

NonStop SSL can write both log messages (informational from an operations viewpoint) and audit messages for security-relevant events.

You should secure the audit log files appropriately to prevent tampering; if you are logging at the byte dump level, also configure them to prevent read access as they will contain passwords and potentially-sensitive user data.

NonStop SSL log messages are, by default, not written to either a disk file or a device/console. You can configure LOGCONSOLE to write them to either $0 or a specific device, LOGEMS to write them to a specific EMS collector, and LOGFILE to write them to disk. HPE recommends writing log events to either a file or EMS. You can configure disk file log rollover behavior using LOGFILERETENTION. LOGMAXFILELENGTH controls the maximum audit file size, and defaults to 20MB.

You can set a generic value for what gets logged using LOGLEVEL, or tailor it with LOGLEVELEMS, LOGLEVELCONSOLE and/or LOGLEVELFILE. For normal usage, up to level 50 is appropriate; you might need to enable higher levels for troubleshooting at HPE’s request.

You can use DONOTWARNONERROR to log user-specified socket errors as errors rather than warnings.

The NonStop SSL log and audit files are unstructured files with file code 0. They can be viewed with any PC editor if and only if they are transferred to the PC in binary. The SHOWLOG tool allows you to view the logs and audits on the NonStop system. In addition, in the OSS environment you can use more and tail, possibly in conjunction with filter tools such as grep, awk or wc.

You should review the access logs regularly – ideally as part of an automated process where the logs are fed into a central SIEM system.

Environmental considerations

If you use the TCP/IP round robin capabilities, make sure the value of your PTCPIPFILTERKEY adheres to common best practices for secure passwords and that it is protected in the CONFIG/CONFIG2 file from unauthorized access. If the PTCPIPFILTERKEY is known to an attacker, he or she might start a process listening on the same port as the NonStop SSL processes, subverting the security layer. For TCP/IPv6, the risk of someone starting a rogue service listening on the same port can be further decreased if you start one instance of your NS SSL process on each CPU, since only one process can be started per CPU with that specific PTCPIPFILTERKEY. When using CIP and IP CLIMs, parallel round robin filtering allows you to start multiple listening processes in the same CPU that share the same port.

Staying up to date on protocol or implementation vulnerabilities

TLS/SSL is visible and accessible for anyone with network access to your NonStop system. Therefore, it is important that you stay up to date on installing fixes or mitigations for protocol and/or implementation vulnerabilities. The TLS protocol does evolve over time to make sure it can withstand attacks (the TLS 1.3 standard is recently approved by IETF and support for this version is in the NonStop SSL roadmap). Make sure you pay attention to HPE's security update notifications and that you follow their guidelines for addressing identified vulnerabilities.

Configuration file security

The NonStop SSL configuration file, specified by CONFIG, is an edit file whose name is passed to the proxy program during startup. Since the SSL proxy program can be run with a number of personas depending on how it is started, you will have at least one configuration file per type (e.g., Telnet server proxy). Secure each configuration file with both read and write access limited to the user who is responsible for its maintenance, as it contains a private key and a pass phrase. Consider using a second config file (CONFIG2) to hold pass phrases, which will allow you to secure the original CONFIG file for more general read access.

Private Key file security

If both a private key file and its associated pass phrase are in the wrong hands, your encrypted TLS/SSL session can be decrypted. For this reason, limit read access to the private key files (SERVKEY and potentially CLIENTKEY) to as few people as possible.

TLS/SSL certificates

There is an important difference between a TLS/SSL connection which is “working” and a connection which is “properly secured”. Unless TLS/SSL certificate checking is properly configured, the TLS/SSL protocol is open to the man-in-the-middle attacks. See SSLMIM for a detailed explanation.

By nature of the TLS/SSL specification, the generation, configuration and verification of certificates is a somewhat complex task. We will describe the task in the following, however you should consider consulting with HPE and/or your internal security team to get advice.

TLS/SSL Server and Client Authentication: background

For each TLS/SSL connection between the server and the TLS/SSL client, there are two directions to potentially authenticate:

- Server Authentication: By design of the TLS/SSL protocol, the server will always send a certificate chain with its response to the client. It is important to (1) use a proper and singular certificate chain on the server and (2) to actually verify the chain on the client. If the client fails to properly authenticate (verify) the server, it allows for the aforementioned man-in-the-middle attack. For this reason, server authentication is very important.
- Client Authentication: This is optional. If configured, the server asks for a client certificate and upon receipt of the client certificate verifies it. If the client sends no certificate or the certificate verification fails, the server will terminate the connection. Without Client Authentication, any client will look the same to the server and any client can connect. Client authentication thus further enhances the security of the TLS/SSL session by allowing the server to control who can connect – at the price of a more complex installation.

Configuring TLS/SSL certificates with the NonStop being the TLS/SSL server
When the NonStop is the TLS/SSL server, the following are the core parameters controlling the behavior:

- On the NonStop
  - CACERTS, SERVCERT, and SERVKEY: configure the (server) certificate chain and key file. All mandatory
  - CLIENTAUTH: activates client authentication. Optional
  - PEERCERTCOMMONNAME or PEERCERTFINGERPRINT: verifies content of client certificate. Optional, only relevant if CLIENTAUTH is used.
- On the remote system where the TLS/SSL Client runs:
  - Fingerprints or usage of Microsoft Certificate Store: controls which certificates (from the remote SSL system, which is the NonStop) to trust; therefore needs to logically match

CACERTS, SERVCERT, SERVKEY. This is important to avoid the aforementioned man-in-the-middle attack. Specifics will depend on the software used on the remote system.

  - If Client authentication is desired, the client certificate and private key file need to be configured. Need to logically match CLIENTAUTH.

Make sure to generate your own certificates for production instead of using the sample certificates and key files shipped with the product. Configure all your TLS/SSL clients to verify the certificates used by the TLS/SSL server. Use the CACERTS parameter to specify the file(s) containing your server certificate signing chain, and SERVCERT to specify your server certificates.

SERVCERT specifies the file containing the PEM or DER encoded X.509 server certificate. SERVKEY specifies the file containing the private key associated with the public key contained in the servercertificate configured by SERVCERT. Use SERVKEYPASS to specify the password needed to decrypt the private key.

Client authentication means that the remote TLS/SSL client will present a certificate to the NonStop server. The TLS/SSL handshake will succeed only if that certificate is trusted by the server. While client certificate authorization (controlled by the CLIENTAUTH parameter) is not enabled by default, you should consider enabling it if your environment is suitable. To enforce verification of the content of the remote leaf certificates presented to TLS/SSL, you can provide either the expected common name (PEERCERTCOMMONNAME) and/or a fingerprint (PEERCERTFINGERPRINT). In both cases, you can configure a comma-separated list to specify multiple entries.

All certificates to be used in HPE NonStop SSL have to be in standard PEM or DER format. Additionally, private key files have to conform to the PKCS#8 standard. If certificates are received in any other format they have to be converted, e.g. with the OpenSSL tools.

Configuring TLS/SSL certificates with the NonStop being the TLS/SSL client
When the NonStop is the TLS/SSL client, the following are the core parameters:

- On the NonStop TLS/SSL client:
  - CACERTS, CLIENTCERT, and CLIENTKEY: configure the client certificate chain and key file. All optional.
  - TRUST: controls which certificates (from the remote TLS/SSL system) to trust. This is important to prevent the man-in-the-middle attack.
- On the remote system:
  - The configuration of the certificates and key file on the remote system is mandatory by design of the SSL protocol. The configuration needs to logically match the TRUST parameter on the NonStop.
  - If TLS/SSL Client Authentication is desired, it has to be configured – details will vary by platform and software. The configuration parameters have to match the configuration of CACERTS and CLIENTCERT on the NonStop.

It is important to properly configure the TRUST parameter to specify a list of trusted CAs when running as a TLS/SSL client. The list can contain either one or more hash algorithm: fingerprint pairs or one or more trusted CA certificates in PEM or DER encoded format (preferred). The default is to not check the TLS/SSL partner's certificate chain, but HPE strongly recommends that you do check it.

CLIENTKEY specifies the file containing the private key associated with the public key contained in the client certificate configured by CLIENTCERT. Use CLIENTKEYPASS to specify the password needed to decrypt the private key. CACERT specifies the signing chain of the certificate configured in CLIENTKEY.

Testing your certificate setup

After you have configured both the TLS/SSL server(s) and TLS/SSL client(s) in your environment, HPE recommends that you test the setup for:

- Correctness of the SSL server certificate setup
- Validation that the SSL clients do indeed verify the SSL server certificates. You can check this by deliberately using another certificate chain on the server

If client authentication is configured, apply the same test principles.

If you do not feel comfortable verifying your configuration, consider getting outside help from your internal security organization and/or from HPE.

Generating certificates

You should generate your own X509 certificate for SSL rather than using the built-in shared certificate. Sign it with a SHA2 digest algorithm rather than MD5 or SHA1 if your infrastructure supports SHA2. HPE recommends using an RSA key size of at least 2048 bits for the SSL certificates; the key size is determined during the generation process.

There are several options for generating your own certificates:

- If your company has an in-house Certificate Authority (CA), that is probably the best option
- Many companies use service providers such as Verisign or Thawte to generate SSL certificates for them. This is a bit costly but reduces complexity a lot
- Finally, you can act as your own CA or use self-signed certificates. This is not recommended unless you have expertise in this area

Converting the file format of certificates and key files

Certificate files and private key files come in a rich variety of formats. The most flexible tool for format conversion is the OpenSSL shell, which is available for the Windows platform for free. If in doubt, consider using expert help from your security department and/or an HPE service.

Other considerations for TLS/SSL certificates

ALLOWCERTERRORS lets you specify the list of certificate error numbers (defined by OpenSSL) that can be overridden. Do not specify this parameter unless necessary, and change it online as soon as you have remediated the condition(s) that led to its use.

You can change the server certificate chain online using the SSLCOM RELOAD CERTIFICATES command. You might need to do this if you have reason to believe that your server private key file has been compromised. If your clients are authenticating the server, you should consider basing trust on the Root CA, which allows you to replace the server certificate without having to reconfigure the clients.

Using NonStop SSL’s firewall capabilities

NonStop SSL can control the remote IP addresses which connect into the HPE NonStop platform. Limiting the remote IP addresses to the minimal range is a recommended best practice. You can “white list” (ALLOWIP) and/or “black list” (DENYIP) either individual addresses or ranges. The default is that all IP addresses are allowed, so you should change the configuration to restrict the range if possible.

Note: Because NonStop SSL runs as a proxy process, the NonStop process that it front-ends will see the requester’s IP address as the loopback address (127.0.0.1 for TCP/IPv4 or::1 for TCP/IPv6) rather than the original requester’s IP address. You can obtain connection information by using the SSLCOM CONNECTIONS command or, for VI requesters, the CFWSADDR macro.

CONTENTFILTER can be used in certain proxy modes to filter all incoming messages and terminate the connection if a message does not match the rule set. This parameter can be helpful to filter the incoming data stream before it reaches the socket application on the NonStop, however it only works for very simple application protocols.

Other configuration parameters

HPE recommends that you set SSLCOMSECURITY to TRUE to restrict execution of sensitive commands to members of the SUPER group and the user under which the SSLOBJ process is running. The default is FALSE.
CIPHERSUITES controls the cipher suites allowed, in order of preference, and its default values may be adjusted by HPE at any time to add new, more secure cipher suites and/or exclude less secure cipher suites. The default value for CIPHERSUITES allows only secure cipher suites – only change the parameter if you know what you are doing. HPE does not recommend use of ADH ciphers (no authentication) or cipher suites 0.1 and 0.2 (no encryption). If you want to restrict your TLS/SSL server or client to a specific algorithm, we recommend using TLS 1.2 with ciphersuites that use AES-256 in the most secure encryption mode (GCM) and use SHA2 algorithms for MACing (e.g. SHA256 or SHA384). In addition, ciphers providing Perfect Forward Secrecy (PFS) should be enabled if possible to prevent an attacker from decrypting captured traffic even if he or she gets access to the TLS/SSL certificates after the fact. Cipher suites that support PFS can be identified by having "_DHE_"or "_ECDHE_" in the full cipher suite description provided in the documentation of the CIPHERSUITES parameter.

Ciphers meeting the criteria mentioned above include TLS_DHE_RSA_WITH_AES_256_GCM_SHA256 (0.107), TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (0.159), or their DSA or ECC counterparts, in case of DSA or ECC certificates.

In order to use DHE PFS ciphers when running HPE NonStop SSL in a server run mode, a DH parameter has to be generated initially (e.g. by using the openssl shell) and has to be referred to via the DHPARAM parameter. To use elliptic-curve-based DH (ECDHE) ciphers, the parameter ELLIPTICCURVENAME must be set. Not all peers might support TLS 1.2 and highly secure ciphers recommended above yet; if not, HPE recommends use of at least TLS 1.1 with TLS_RSA_WITH_AES_256_CBC_SHA, called AES256-SHA in OpenSSL and configured via the value of 0.53 for CIPHERSUITES.

You can configure the minimum and maximum admissible TLS/SSL protocol versions using MINVERSION and MAXVERSION. The mappings are as follows:

Protocol Version	**MAXVERSION/MINVERSION Long Parameter Value**	**MAXVERSION/MINVERSION Short Parameter Value**
TLS v1.2	TLSv1.2	3.3
TLS v1.1	TLSv1.1	3.2
TLS v1.0	TLSv1.0	3.1
SSL v3.1	SSLv3.1	3.1
SSL v3.0	SSLv3.0	3.0
SSL v2.0	SSLv2.0	2.0

MAXVERSION and MINVERSION defaults vary by SPR; see MINVERSION in SSLREF for details.

You should set MAXVERSION to TLSv1.2. If you are subject to PCI DSS, see the version 3.2 specification (PCIDSSV32) for detailed information on use of earlier protocol versions. This specification is a useful resource even if you are not subject to PCI DSS. If possible, set MINVERSION no lower than TLSv1.1. SSLv2.0, SSLv3.0 and early versions of TLSv1.0/SSLv3.1 have known protocol flaws and should not be enabled unless absolutely necessary.

HASHALGORITHMS defines which hash algorithms are used when verifying the TLS/SSL server side based on its fingerprint (relevant for the TRUST parameter). By default it does not enable either

MD5 or SHA1. You should not enable either of these algorithms unless absolutely necessary, as they are cryptographically weak.

To help prevent DoS attacks by a rogue TLS/SSL client, session renegotiation (ALLOWRENEGOTIATION) is off by default. Enable it only if necessary.

If you are allowing use of FTP and your environment requires that the FTP server make the data socket connection (active mode), then you must set PASSIVE to 0.

By default, NonStop SSL will not allow unencrypted sessions when running in FTPS mode. If you must allow them, set FTPALLOWPLAIN to TRUE.

The default ports for incoming connections range from 11011 to 11014, depending on the SSL run mode. If you need to use different ports in order to avoid conflicts, use the PORT parameter.

For FTP data connections, the default minimum port is 40000 and the maximum port is 41000. Use FTPMINPORT and FTPMAXPORT to adjust these values as needed to avoid conflicts with other services.

By default, NonStop SSL binds to any IP address for local binding on incoming connections. Use INTERFACE to set a specific IP address. Similarly, it binds to any IP address for local binding on outgoing connections. Use TARGETINTERFACE to set a specific IP address.

Make sure your PTCPIPFILTERKEY adheres to common best practices for secure passwords and that it is protected in the config file from unauthorized access. If the PTCPIPFILTERKEY is known to an attacker, he or she might start a process listening on the same port as the NonStop SSL processes, subverting the security layer.

Properly secure access to TCP/IP tracing facilities (i.e. PTRACE). Since SSLOBJ is a TCP/IP proxy, the local traffic between NonStop SSL and the target application will be transferred in plain text over the local loopback.

In FTPC mode, limit the access to outside FTP servers on a need-to-access basis. In addition to the ALLOWIP/DENYIP params applicable in all NonStop SSL runmodes, the TARGETHOSTFORCE and TARGETPORTFORCE parameters can help with this objective.

If running in TLS/SSL server mode, set parameter CIPHERSERVERPREFERENCE to TRUE to enforce that your preferred ciphersuite order is used rather than the client's preferred order.

See SSLREF, Installation and Configuring, for additional details.

Note: Many of the above security considerations also apply to openssl ports to NonStop servers.

IPSec: best practices

Internet protocol security (IPSec) provides application-transparent encryption services for IP network traffic. You can set up IPSec on an IP-address-to-IP-address basis, and optionally on a UDP or TCP port, but you cannot establish IPSec on a per interface basis. See IPSEC for a NonStop IPSec overview.

You will need to obtain X.509 certificates from a certificate authority, install them on the NonStop Console, move them (along with the private key files and certificate revocation lists) to /etc/racoon/certs on the CLIM along with any peer certificates, and establish your configuration including the security policy.

See CIPMON Section 3, Configuring IPSec (IP CIP) for details.

iptables/ip6tables (IP CIP)

You can use iptables and ip6tables to filter incoming TCP/IP inside of your CLIMs.

Note: The CIP iptables and ip6tables configurations are not failed over. You must pre-set the failover CLIM’s iptables and ip6tables configuration in anticipation of a failover.

Caution: If you plan to make extensive use of iptables and ip6tables, be sure to check for potential performance impact.

See CIPMON Section 3, Configuring CIP iptables/ip6tables (IP CIP) for details and examples.

Expand

Encryption options for Expand include NonStop SSL and, if CLIMs are present at both ends of the Expand line, IPSec,   Third-party hardware products also have been used to encrypt ServerNet long-haul traffic.

Caution: Be sure to measure the performance impact of enabling TLS/SSL or IPSec for Expand.

See SSLREF Installation, Installing an SSL Tunnel for EXPAND-over-IP Lines, and SSLREF, Advanced Configuration Topics, Load Balancing and Fault-Tolerance of EXPAND over SSL

JDBC/MX

The JDBC/MX type 4 driver includes native SSL support for communicating with the ODBC/MX server, starting with the J06.16/H06.27 RVUs. It uses the standard Java package (javax.net.ssl), which you configure using the runtime sslEncryption JVM property. You can either add the server certificate to the default Java trust store or create a new trust store. HPE recommends using SSL if any sensitive data is being transmitted between the JDBC/MX type 4 driver and ODBC/MX.

See JDBCMXT4REF Section 4, Type 4 Driver Properties, sslEncryption, for details.

ODBC/MX

The NonStop SQL/MX connectivity service, MXCS, includes support for remote connections to the SQL/MX database, such as ODBC and JDBC. The ODBC features can be configured at a system level through MXCI. By default, SUPER.SUPER is the only user ID with operator privileges, allowing the execution of sensitive commands against the MXCS configuration database. SUPER.SUPER can grant operator permissions to other users.

Note: MXCS configuration is done with user IDs; it does not support Safeguard aliases.

Caution: Remember to check for and, if present, remove operator privileges in MXCS when deleting a user.

See SQLMXCSREF, Sections 2 and 7 for details.

NonStop Software Essentials (NSE)

NSE ships with a default Jetty client certificate, whose localhost address does not match the system on which the NSE client is hosted. This is likely to trigger a browser warning. To avoid the waning, replace the certificate with one that matches the localhost. Instructions on replacing the default Jetty certificate are available from the GNSC, and will be included in the upcoming NSE 5.1

HPE NonStop OmniPayments Switch

OmniPayments Switch is a high availability Enterprise Service Bus (ESB) running on HPE NonStop servers. It is SOA compliant and orchestrates new services across hybrid platforms as well as legacy Pathway services. OmniSwitch (as the product name is known in short) supports flexible document interchange between applications. A rich range of adapters for data, message and transaction parsing, as well as rules engines and data transformation tools are readily available.

The switch has various security features. Some of them are described below.

TLS

OmniSwitch has TLS tunnel and supports TLS1.2 for secure communication over Internet.

Data Encryption:

For channels which need end to end encryption, OmniSwitch supports accepting (and sending) encrypted transactions. The encryption with 3DES is always with 192 bit key and OmniSwitch also supports and insist channels to use AES as well. OmniSwitch integrates with industry leading HSMs for encryption, thus ensuring that encryption takes place inside a tamper proof box.

Transaction Security:

OmniSwitch offers various security measures to ensure security of transactions. Application of these security measures is based on transaction profile.

- PIN Operations – Supports verification, translation and change of PINs
- Key Exchanges - In case of PIN operations, keys involved in PIN encryption are changed over the link periodically as per PCI mandates. We offer various different methods of key exchanges e.g.
    - Master Slave: Where the master generates the key, encrypts with key exchange key and sends to peer switch. OmniSwitch can act as master or slave depending on predefined configuration
    - DUKPT: OmniSwitch supports Derived Unique Key Per Transactions. With this method each transaction uses a new key that minimizes risk of PIN encryption from being compromised.
    - Manual Key Load: Loading of Key encryption key can be done manually using multiple key custodian method. OmniSwitch supports minimum 2 key custodians each having 1 key component and ensures each custodian does not have knowledge of another key component possessed on the other key custodian. OmniSwitch uses HSM's secure facility to achieve manual key loading with multiple key custodians
    - Remote Key Loading: For devices like ATM which support remote key loading for Master Key, OmniPayments supports the secure RKL protocols as laid down by ATM vendors like Diebold, NCR, etc.
- User-Passwords - The passwords in OmniSwitch are securely stored in the form of strong one-way cryptographic hash, databases holding secure data are protected with MAC validation so unauthorized tampering is detected. The channels/users are allowed to select a strong password only e.g. minimum length of password being 8, should include at least one numeric and special character.
- Message Authentication - To ensure messages are not altered during transit, MAC is used
- WS-Security - We provide WS-Security for the channels that support it. It ensures the transaction message is received by only intended parties and sensitive fields are never in clear.

OmniUI

Configuration, monitoring and operations within OmniSwitch are controlled and managed using a web based interface OmniUI.
OmniUI uses user-password method for authentication. OmniUI uses OmniSwitch's ACL functionality to allow or disallow which features of OmniSwitch a given user can touch for configuration and operations
ACLs - OmniSwitch has inbuilt access controlled list that governs which service a given channel can access.
All activity performed via OmniUI is audited for security purposes.
All sensitive information is masked.

Tokenization

Card numbers are never stored in the clear in OmniSwitch database. OmniSwitch provides tokenization service to intended customers like merchants who want to store cards details of customer so that next time when the customer returns his card details are not needed to be entered. For critical security reasons cards can’t be stored at merchants so we allow them to convert cards to tokens using OmniPayments tokenization service.

Data at rest protection

User data should be secured according to least privilege, with tokenization or encryption in force as required (for example, to comply with the Payment Card Industry Data Security Standard, PCI DSS – see PCIDSSV32 for details).

Production data that includes sensitive information should not be present on development/test systems.

Securing data

You can configure security for both Guardian and OSS files to deny access to SUPER.SUPER and its aliases. You can find more details above and in both SECMGMT, Section 6, Managing Security Groups, and OSSMOG.

For Guardian files, use SAFECOM to establish the appropriate ACLs.

You can protect OSS files from access by the super ID by placing them in a restricted-access fileset.

See OSSMOG Section 9, Using Restricted-Access Filesets and File Privileges, for details.

Note: In addition to application files and tables, user data may appear in other locations including TMF audit trails, RDF image trails, application log files, EMS events, swap files, CPU dumps and application saveabend files (ZZSA files).

Field/column-level tokenization and encryption

HPE does not offer this capability in its own product set. Partner products are available from MicroFocus (as of August 2017) and comForte that offer both application-transparent implementations and the ability to explicitly tokenize/detokenize or (MicroFocus) encrypt/decrypt data. Both products support Enscribe files, SQL/MP tables and SQL/MX tables.

Disk and tape volume encryption

NonStop Volume Level Encryption (VLE), in conjunction with the HPE Enterprise Secure Key Manager (ESKM), is available for protection of CLIM-connected disk storage (both SAS and external) and tape storage. VLE performs encryption and decryption within the storage CLIM; the data is in the clear and accessible subject to normal security protection within the NonStop server itself. Encryption management for VLE is configured through the SCF storage subsystem and requires creation of the Safeguard SECURITY-ENCRYPTION-ADMIN group.

Note: Additional protection mechanisms are needed if security standards require that logical access must be managed separately and independently of native operating system authentication and access control mechanisms (e.g., for PCI DSS V3-2, requirement 3.4).

Note: Encryption is configured at the physical drive level rather than the volume level, and each drive has its own encryption key. See VLEGUIDE for additional details. A number of best practices, including ESKM High Security Configuration, appropriate security for files in $SYS.ZENCRYPT, and post-installation cleanup are described in NSVLEGUIDE. An ESKM Best Practices document also is available if you are using NonStop VLE; contact your HPE account team to receive a copy.

NonStop cF Secure Tape provides a software-based encryption solution for securing files written to physical or virtual NonStop backup tapes. See the product data sheet and SECURETAPE for additional details.

Data sanitization

Memory

The operating system's memory manager always zeroes out memory pages prior to assigning them to processes. When the memory manager reads a partial page from disk, the rest of the page is zero filled.

Disk drives

The most reliable way to sanitize an entire disk drive is to physically destroy it. If the drive is still functional and you wish to reuse it, you can use OSM data sanitization to overwrite the entire drive's contents a configurable number of times with a configurable set of patterns.

See the OSM Service Connection User's Guide, Section 11, Perform Data Sanitization and the online help for the data sanitization guided procedure for additional details.

Disk files

Individual disk files can be configured through FUP for CLEARONPURGE, which will cause zeroes to be written over the file contents up to end of the highest-allocated extent. You also can configure CLEARONPURGE globally through Safeguard's CLEARONPURGE-DISKFILE attribute, which will cause the contents of all files, including temporary files, to be similarly overwritten.

Note: If you set CLEARONPURGE-DISKFILE ON, it remains in effect even when Safeguard is down.

Note: Performing a PURGEDATA on a file followed by a DEALLOCATE does not trigger CLEARONPURGE.

Caution: There can be a negative performance impact due to CLEARONPURGE, especially if multiple files or large files are being overwritten, even though the operating system does limit the amount of CLEARONPURGE I/O per second for an individual file by allowing CLEARONPURGE I/O to be preempted by other work for the volume and, when preempted, rescheduling the work with a delay.

Partner products also are available that can clean free disk space.

Disk volumes

Partner products are available that can overwrite most of the contents of a volume.

Note: Due to their nature they cannot overwrite either sectors that have been spared or the volume label itself.

VTS drives

There is no explicit command or process to sanitize the drives; however, creating a new filesystem will reformat the partition and wipe its data. See the section on reconfiguring vaults in the appropriate version of VTSGUIDE.

NonStop CLIM security

In newer NonStop servers, the I/O subsystem (cluster I/O protocols, or CIP) uses IP Cluster I/O Modules (CLIMs) for network I/O, Storage CLIMs for peripheral storage I/O, and Telco CLIMs for the Telco platform. These CLIMs are HPE ProLiant servers running the Debian GNU/Linux with HPE

CLIM Extensions, which is based on the Debian Linux distribution. Unless otherwise noted, in this section “CLIM” is a generic reference to all three types.

The CLIM is not a general-purpose Linux server. Because of its limited role, it needs to have only those software packages installed, and those services running, that provide IP, storage, or Telco functionality. Only HPE-required software is installed on the CLIM.

No development packages are installed. As examples, the following packages are not installed on a CLIM:

- gcc (GNU Compiler Collection)
- gdb (GNU Debugger)
- flex (Fast Lexical Analyzer)
- cpp (C pre-processor)
- g++ (C++ compiler)
- bison (GNU parser generator)
- make (Build tool)

CLIM connectivity

A CLIM provides two levels of network connectivity: maintenance connectivity and data connectivity.

Maintenance connectivity is required so that CLIMs can be monitored and managed as an integral part of the NonStop BladeSystem. The maintenance connectivity is provided physically via the Integrated Lights-Out (iLO) and eth0 (maintenance) ports. Both of these ports are connected to the NonStop maintenance LAN.

Data connectivity is required for IP and Telco CLIMS, to connect the corresponding NonStop system to the data LANs via the eth1, eth2, eth3, eth4, and eth5 ports. This is similar to the connectivity of the earlier-generation E4SA, FESA, GESA, and G4SA adapters.

Storage CLIMs with one or more drives configured to use NonStop Volume Level Encryption (VLE) require data connectivity for communication with HPE Enterprise Secure Key Managers (ESKMs). Communication between the storage CLIMs and ESKMs uses SSL and certificate-based authentication.

Only required services running on the CLIM

One of the most important ways to secure any system is to ensure that it does not have unnecessary services running. Thus, CLIMs have only required services running on them. As mentioned earlier, CLIMs have no unnecessary software installed on them. This helps to ensure that these unnecessary services do not run on the CLIMs.

An additional important security measure is to ensure that the services installed are accessible only at the access points (interfaces) where they need to be accessible. Keeping this in mind, no service runs on any non-maintenance interface.

On the maintenance interface (eth0), the only default services configured to run are those required for monitoring and controlling CLIMs. These services are:

- SSH (Secure Shell) / SFTP (Secure File Transfer)
- SNMP (Simple Network Management Protocol)

If a CLIM has been configured to run services to allow NonStop processors to boot from Halted State Services (HSS) images hosted on the CLIM rather than the NonStop Console (default), those CLIMs also run the following services:

- BOOTP (Bootstrap Protocol) / DHCP (Dynamic Host Configuration Protocol)

- DNS (Domain Name System)

All other services are disabled on the maintenance interface (in addition to all services being disabled on non-maintenance interfaces). Specifically, the following services are completely disabled on CLIMs:

- Telnet
- FTP (File Transfer Protocol)
- Finger (User information lookup)
- rlogin/rsh (Remote Login/Remote Shell)
- HTTP/Web server (for example, Apache)
- NFS (Network File System)
- X11 (X-window)
- Ident (Pattern search in a file)
- POP3 (Post Office Protocol)
- RPC (Remote Procedure Call)
- SMTP (Simple Mail Transfer Protocol)
- IPP (Internet Printing Protocol)
- Samba (File and print services)

Only required ports open on the CLIM

One consequence of unnecessary services not running on CLIMs is that no unnecessary TCP/UDP ports are open on the CLIM. On data (non-maintenance) interfaces, the following UDP port is left open:

- IKE (Internet Key Exchange)

On the maintenance interface (eth0), only the ports required by the following standard services are open:

- sshd/sftp
- snmpd

In addition, non-standard ports are open for HPE CLIM management services:

- confsync
- climagt

For Telco CLIMs, there are additional ports open for INS services.

All other ports are not open on the maintenance interface (in addition to all ports not being open on non-maintenance interfaces).

See SECMGMT Appendix E, CLIM TCP/IP Ports, for a complete list of ports used by the CLIM. You can use this list as input for your firewall configuration and block all traffic for other ports.

SLNP support on the CLIM (J06.14 and earlier CLIM software versions)

By default, processes that use sockets to bind an IP port (such as a TCP port or UDP port) can specify INADDR_ANY, to listen for activity on that port for multiple interfaces. In doing so, the process binds all available interfaces (eth0, eth1, eth2, eth3, eth4, eth5, and lo [the loopback interface]). However, eth0 is dedicated to the maintenance LAN, and some processes that must listen to the maintenance LAN have been coded to use INADDR_ANY.

SLNP (Simple Logical Network Partitioning) is HPE value-add software that is not present in other Linux distributions. It provides a mechanism to logically partition the network of a Linux system, through changes in the way bind () and connect () socket calls operate. With SLNP rules in place, when an application tries to bind or connect a socket, it is limited to only use those IP (v4 or v6) addresses or interfaces that are meant for its use. All this is possible without changes to the applications themselves.

The CLIM comes preconfigured with SLNP with rules that constrain the sockets (data) usage to only eth1, eth2, eth3, eth4, and eth5 and the maintenance LAN usage to only eth0. This enables the ports required to be open on the maintenance LAN (22 for sshd/sftp and 161 for snmpd) to not open on the data LAN. This means that there is no access to the CLIM from the data LAN.

All of the CLIM SLNP rules are hard-coded and the user is provided no external interfaces to modify them.

Subnet routing rules

The CLIM's multiple network interfaces typically are used to reach different IP subnets. Routing rules added for each interface instruct the CLIM as to what subnets are reachable through that interface.

When a packet is received by a CLIM, it will drop the packet unless it comes from one of the subnets configured in its routing rules. It will, however, permit the packet to be received if the packet arrives on a different interface than the routing rule is associated with.

The packet filter can be further narrowed to restrict packets to be received only if the specific interface has a matching route by altering the "rp_filter" sysctl value:

- rp_filter = 0: Packets are accepted from any source
- rp_filter = 1: Packets are accepted from IP addresses that are in the CLIM's routing tables for the interface that the packet arrived on.
- rp_filter = 2 (CLIM default): Packets are accepted if they come from an IP address that is in any of the CLIM's routing tables for that provider.

Consider setting rp_filter to 1 unless you have configured routes that can send over one interface and receive over another. Use the following command to alter the rp_filter setting for interface <interfaceName>

CLIMCMD <climname> climconfig sysctl –update net.ip4.conf.<interfaceName>.rp_filter <value>

For example:

CLIMCMD *NCLIM001* climconfig sysctl –update net.ip4.conf.bond0.rp_filter 1

Source routed packets

By default, the CLIM ignores source routed packets.

Redirects

By default, the CLIM is configured to allow redirects. Consider disabling redirects if practical in your environment. Use the following command to alter the accept_redirects setting for interface <interfaceName>

CLIMCMD *<climname>* climconfig sysctl –update
net.ip4.conf.*<interfaceName>*.all.accept_redirects <value>

For example:

CLIMCMD NCLIM001 climconfig sysctl –update net.ip4.conf.bond0.all.accept_redirects 0.

Similar considerations and syntax apply for the following settings:

- net.ipv4.conf.all.secure_redirects
- net.ipv4.conf.default.accept_redirects
- net.ipv4.conf.default.secure_redirects

Namespace support on the CLIM (J06.15 and later CLIM software versions)
Starting in J06.15, with the Multiple Providers on CLIM feature, the CLIM stopped using SLNP as the technology to isolate maintenance LAN traffic from the NonStop data LAN.

To replace it, a standard Linux facility, Linux Network Namespaces, is now used to provide a similar role. Network namespaces provide entire logically independent “network stacks”, ensuring that traffic for one network namespace is kept isolated from other namespaces.   On the CLIM, one network namespace is dedicated to all maintenance LAN services, while additional, separate namespaces, one per provider, are dedicated to data traffic using the eth1 through eth5 interfaces. The creation and assignment of traffic to these network namespaces is automatic, and there are no external interfaces to control the assignment of network resources to them.

iptables and ip6tables support on the CLIM (J06.13 and later CLIM software versions)

Starting in J06.13, the CLIM supports two facilities, iptables and ip6tables that can be used to provide additional hardening for data LAN traffic. These packet filtering facilities allows rules to be added to examine network traffic on a per-packet basis, allowing or dropping packets that don’t meet the user-configurable matching criteria. For example, incoming traffic may be dropped if it is not of an expected protocol, port, or from a trusted address.

CLIM access

As mentioned earlier, all IP-based access to the CLIM is restricted to the maintenance LAN over the eth0 interface. The maintenance LAN access to the CLIM is only via SSH/SFTP or SNMP. SNMP access is required by OSM (Open System Management) software running on the NonStop system so OSM can monitor CLIMs as an integral part of the system.

Access to a system via FTP, Telnet, or rlogin (rsh) is vulnerable to eavesdropping, so CLIMs do not have these insecure services running on them. All access to CLIMs for command/control is via SSH version 2 and SFTP.

SSH/SFTP access is available from the NSC or from the NonStop system. SSH/SFTP access from the NSC requires a username and password. This means that whenever users need to access a CLIM directly from the NSC (for example, to perform a CLIM software upgrade), they must supply two usernames/passwords: the username/password of the NSC and the username/password of the CLIM.

Similarly, access to the CLIM iLO interface is over SSH and HTTPS.

Under rare circumstances, the GNSC might recommend performing a specific action through PuTTY rather than CLIMCMD. If this is necessary, you should alter the PuTTY configuration (Session -> Logging) to capture appropriate audit.

SSH/SFTP access from the NonStop system uses certificate-based authentication, whereby a NonStop SUPER group user can run the SSOCLIM tool to establish the SSO (Single Sign On) for all NonStop SUPER and non-SUPER group users. SSOCLIM configures the SSO to provide SUPER group members with root privilege to execute commands on the CLIM and non-SUPER-group members with non-root (user) privilege. All NonStop programs that perform configuration and control on CLIMs (such as CLIMCMD, CLIMBKUP, and CLIMRSTR) use this design, thus providing clear access control.

The root and non-root CLIM private keys are stored in $SYSTEM.ZSERVICE. By default they are owned by SUPER.SUPER, with the root private key (SUPERKEY) and root public key (CLIMSKEY) secured as CCCC and the user private key (USERKEY) and user public key (CLIMUKEY) secured as NCNC. These settings allow any local or remote SUPER group member to issue commands with root privilege and any local or remote user to issue commands with non-root privilege.

Caution: Do not loosen the root private key's read security setting; otherwise, you will allow non-SUPER-group members to execute commands on the CLIM with root privilege.

You also can access a CLIM via the maintenance LAN over the iLO interface. This interface also requires username/password-based authentication.

As mentioned before, CLIM access from the NSC is based on username/password-based authentication. HPE software does not store the CLIM username/password anywhere on the NSC.

CLIM security updates

An important aspect of CLIM security is to apply security updates periodically. The Debian Linux community tracks security issues and vulnerabilities, and releases updates on a timely basis. HPE includes these security updates in the latest build of Debian GNU/Linux with HPE CLIM Extensions, on a periodic basis. All of the newly shipped CLIMs are shipped with the latest security updates included. All of the CLIM software updates also include the latest security updates.

Sensitive data on the CLIM

None of the IP CLIMs or Telco CLIMs store any sensitive data. All the data stored on these CLIMs is either dynamically created or configuration data. Thus, even if a malicious user gets access to the data on one of these CLIMs, this data cannot be used to get to any sensitive information.

All of the data buffers on a Storage CLIM are in an area of kernel memory that is not copied on a crash dump. The crash dump will not have any in-flight data. Also, the memory region used by the Storage CLIM for data buffers does not get swapped out to disk.

Persistent CLIM security configuration

CLIM configuration can be backed up on the NonStop BladeSystem using the CLIMBKUP tool. This configuration can be restored if the CLIM was re-imaged or replaced, or its hard disk was replaced. The backed-up configuration includes all security configurations other than VLE encryption keys, and thus any security configuration (such as changes in default passwords) does not need to be redone after CLIM re-image or replacement. The storage CLIM will retrieve the encryption key from the ESKM when an encrypted drive is started.

Authentication events

You should configure the CLIM Authentication Transport feature to transport authentication events from the CLIM to the NonStop system, where they are logged as EMS events. Their EMS collector, $ZCLA, should be preconfigured in the SCF Kernel subsystem persistence manager; make sure that it is configured and running. If $ZCLA is running, XMA will include it as a security activity source.

See CIPMGMT Section 3, Configuring Transport of Authentication Events from CLIM, for additional details.

Protection against broadcast storms and Denial-Of-Service (DOS) attacks

All devices that can be accessed via IP are potentially subject to network broadcast storms and DOS or Distributed-Denial-of-Service attacks. General rules for protecting devices from malicious traffic apply to CLIMs, including limiting physical access and keeping subnets as small as possible to reduce the potential impact of broadcast storms (which usually are accidental rather than malicious). Firewalls can help throttle traffic, but the throttling is likely to affect normal traffic as well as unwanted traffic. Logical Network Partitioning (LNP) can help limit the affected listeners. As discussed earlier, under TCP/IP, ICMP packet filtering can provide protection against some specific attack classes.

Customer responsibilities for CLIM security

HPE has taken the measures described above to deliver CLIMs with out-of-the-box security. You should take the following measures to make sure that the CLIM security is not compromised:

Do not install unapproved software on the CLIM

HPE helps ensure that a CLIM is shipped with only the software required to provide IP, storage, or Telco functionality for the NonStop environment. You should not install any non-HPE-approved software on the CLIM. Doing so will result in a violation of your support contract with HPE.

Change CLIM default passwords as part of initial installation. For ease of installation, all CLIMs are shipped with default passwords for both the iLO and maintenance (eth0) interfaces. You should change these default passwords immediately after setting up the CLIMs, before they are made accessible to the network.

Use difficult-to-guess passwords

Use the same guidelines in setting up CLIM passwords as you do when creating NonStop user passwords or any other password in your data center.

The following references can be used to create a difficult-to-guess password:

- UNIX/Linux man passwd command provides hints on how to create good passwords.
- Eric Wolfram's How to Pick a Safe Password (see SAFEPASS)

Do not configure or connect unnecessary Ethernet ports

In addition to the iLO port, all CLIMs are shipped with only one Ethernet port configured (eth0). This is required so that CLIM can be accessed over the NonStop dedicated maintenance LAN for configuration and software updates. You should make sure that unnecessary CLIM Ethernet ports are neither configured nor connected. For an IP CLIM, only configure and connect the Ethernet ports that are to be used. For a Storage CLIM, do not configure or connect any port other than iLO and eth0.

Do not use native Linux commands

Although a CLIM runs the Linux operating system, do not use it as a general-purpose Linux server. It is and should be treated like a NonStop system networking, storage, or Telco adapter. You should NOT use native Linux commands to perform any networking, storage, or Telco configuration or control. Only HPE-documented CLIM-specific commands should be performed on a CLIM. There are multiple reasons why Linux native commands should not be used:

- Linux native commands do not result in a configuration change that is persistent across CLIM reboots. On the other hand, HPE-documented CLIM-specific commands help ensure that any configuration change performed is persistent across CLIM reboots.
- Linux native commands do not understand the failover configuration. On the other hand, HPE-documented CLIM-specific commands make sure that any configuration change that has an impact on the failover configuration is passed on to the other CLIMs in the system as needed.
- Commands issued from NonStop TACL prompt (or from the HPE NonStop Cluster Essentials and the HPE NonStop I/O Essentials products) use certificate-based authentication between NonStop BladeSystem and CLIMs, after the user has been authenticated on the NonStop system, thus providing two layers of authentication.

Physically secure CLIMs in the data center

CLIMs are an inherent part of NonStop system and should be physically secured in the same way as the rest of the NonStop BladeSystem. CLIMs should be kept in locked data centers where access is restricted to people with the appropriate security clearance.

Secure NonStop maintenance LAN

CLIMs are connected to the NonStop maintenance LAN. It is important that the NonStop maintenance LAN itself be secured to ensure that CLIMs are not accessible by any malicious user. The main component on the NonStop maintenance LAN that needs to be secured is the NSC. See NSC security, below, for details.

Upgrade CLIM software periodically

As new security fixes are released by the Linux Debian community, HPE expects to package these security fixes in new versions of CLIM software. HPE recommends that customers upgrade the CLIM software on a regular basis to not only take advantage of the new functionality and functionality-related defect fixes, but also to gain protections that might be provided with any released security fixes.

NonStop System Console (NSC) security

The maintenance LAN has always been assumed to be a carefully guarded asset, and you need to protect it. A major element in securing the NonStop maintenance LAN is securing the NSCs. The NSCs exist to perform a specific set of NonStop server management functions performed over the maintenance LAN, such as resetting and loading processors, updating CLIMs, and performing service actions or software installation. HPE recommends that you not use the NSCs for day-to-day operations, and that you in no way consider them to be general-purpose PCs. Do not install additional applications on them other than antivirus software or software firewalls, or enable services not needed for NonStop system management. To the degree possible, physically secure your NSCs.

The NSC Security Policy and Best Practices Guide, NSCPOLBP, describes HPE and customer responsibilities and goes into detail on topics such as patch management, user management and audit configuration.

Beginning with NSC Installer DVD Update 26, HPE provides a package that allows installation of the Microsoft Security Compliance Manager recommended Windows Server 2012 R2 settings (with NonStop-required changes) and the Microsoft Enhanced Mitigation Experience Toolkit (EMET). NSCs do not ship with this package installed. See Support Note S16034 for details.

HPE strongly recommends that you install antivirus software and a software firewall on your NSCs.

You should take the same steps for the NSC as you would to secure any sensitive system, such as requiring use of strong passwords and individual user accounts rather than shared accounts. NSCs are shipped with only three configured users, all with administrative rights. You might need to configure additional administrative users in order to provide individual accountability, but you should not configure any non-administrative users unless required.

In addition to establishing and maintaining a Windows security configuration that aligns with your company's security policies, HPE encourages you to be proactive about keeping the NSC's shipped versions of Windows, Oracle Java, Adobe Reader and open source software such as OpenOffice and PuTTY up to date with respect to security patches. OpenOffice is present primarily for use by HPE Support; if you prefer, you can remove it.

You can tighten Remote Desktop security by configuring it to require encryption, and also should select the 'Allow connections only from computers running Remote Desktop with Network Level Authentication' setting if it is available.

For additional details about Windows Server 2012 Remote Desktop encryption configuration, see NSCRDWS1012. For additional details about Windows Server 2008 Remote Desktop encryption configuration, see NSCRDWS2008.

See the Insight Remote Support (IRSSEC) security documentation for details on securing that connection.

For related security considerations, especially with respect to protecting data in transit, see the OSM section, above.

iLO Security

The NonStop system, CLIM and NSC processors all contain iLO management processors. Each Intel processor architecture is supported by one of the iLO types:

x86	iLO 5 (x86-based)
x86	iLO 4 (x86-based Gen9 and Gen8 processors)
x86/Itanium	iLO 3 (different variants for x86, Poulson, and Tukwila processors)
x86/Itanium	iLO 2 (variants for earlier x86-based processors and Itanium predecessors to Tukwila)

The security capabilities of the various iLO types vary both by type and by version within type. ILO4SEC and ILOSEC offer excellent discussions of the iLO 4 and iLO 3 and earlier security capabilities, respectively. ILO4SEC is based on the latest HPE version (currently 2.50); ILOSEC is based on relatively old iLO 3 and iLO 2 versions, so you also should consult the iLO 3 or iLO 2 User Guide as appropriate. For all types, you may have different configuration options available, depending on the versions that you currently have installed and what newer versions currently are supported by HPE NonStop.

To locate the appropriate iLO manuals, search the Hewlett Packard Enterprise Support Center (HPESC) for "iLO 4", "iLO 3" or "iLO 2".

Note: In some cases, older protocols, cipher suites and individual algorithms can be disabled only by configuring either Enforce AES/3DES or FIPS mode. If you are considering use of these options, pay attention to the cautions about them in the iLO documentation - especially for FIPS mode.Securing the NonStop VTC / BackBox

The Virtual Tape Controller (VTC) component of BackBox uses a Windows Server OS, and the recommendations with respect to security are similar to those for the NSC: be proactive about installing patches for its software, and apply your normal Windows policy for installation of antivirus software and firewall configuration.

Note: HPE recommends that you install Microsoft patches manually rather than configuring Windows for automatic updates. Automatic updates may impact VTC usage if the update causes services to be restarted or the VTC to be rebooted. If you have two VTCs, the best practice is to choose a time when none of the tape devices connected to an individual VTC are in use, stop those devices, perform the update and then restart the devices. Repeat this procedure for the other VTC at a time when its devices are not in use.

As with the rest of the NonStop server, you also should consider physical access restrictions to the hardware.

To protect sensitive data, HPE recommends the use of BackBox virtual volume encryption.

See BACKBOXUSER for general security configuration information, BACKBOXSSL for information about configuring SSL communication among BackBox components and BACKBOXTAPE for information about configuring tape encryption.

Securing the NonStop VTS

Primary documentation is the appropriate Virtual TapeSystem (VTS) Configuration Guide (there is one per release).

The VTS uses a Linux-based OS. It is configured out of the box to disable unused services, use SSH for copying data between locations, use access controls and role based privileges, and appropriately configure TCP/IP security. HPE recommends that you implement appropriate network boundary and network security measures including firewall rules, use of only local network access, etc. A list of ports and protocols used by VTS may be found in Appendix D. The VTS is considered to be a closed appliance and does not require installation of antivirus software.

As with the rest of the NonStop server, you also should consider physical access restrictions to the hardware.

To protect sensitive data, HPE recommends the use of Data Encryption (formerly SecureVTS).

A list of ports and protocols used by VTS may be found in Appendix D.

Users

Information on configuring user accounts is in Section 15. The vtsa account has routine maintenance privileges. For accountability it is preferable to add a new user account for each user who will need maintenance privileges, rather than having them share vtsa. Added users have the same level of privilege as vtsa.

Access to the bill account should be restricted as described, as it has the same access rights as root.

The web interface has three default user accounts, with different privilege levels. Default passwords should be changed, and users added to the Operations, Supervisor and Administration groups as needed. Enable a closed system to require authentication. You can adjust group rights if appropriate. You also can create a group that does not have access to virtual tapes; the process is described.

Data encryption

See section 12. Both virtual and physical tapes can be encrypted.

Note: Encrypted data is decrypted as it is read by the host server.

Other configuration

See Section 17.

VTS uses self-signed certificates by default. You can import certificates from trusted CAs.

Configure session timeouts (Configuration > System > Edit System Settings > Timeout.

Configure the SMTP port used for email notifications (also under Configuration > System > Edit System Settings).

Tune the log rotation settings as needed (configuration files are in /etc/logrotate.d).

Consolidated monitoring and reporting

Due to the complexity of modern cyber attacks, the central collection of "security events" across multiple source systems has become a best practice and is mandated in multiple security standards. You can integrate your NonStop systems with enterprise Security Information and Event Management (SIEM) devices such as HPE ArcSight, allowing your security group to both have visibility of local events through the well-established, existing audit reports on NonStop systems and feed NonStop events into the consolidated view of activity across the entire data center.

For more details on HPE's SIEM, see ARCSIGHT.

XYGATE Merged Audit (XMA) provides a consolidated view of security-related activity on your system. Configure XMA to collect all relevant audit, filter as needed, send security events to your enterprise SIEM device, and report on activities of interest and alert you to potentially-serious problems such as repeated authentication failures.

Partner products also are available that provide similar functionality.

File integrity checking

Ensuring the integrity of critical operating system files and settings is an important part of your security strategy.

You can use XYGATE Compliance PRO (XSW) or partner products with similar functionality to monitor, analyze and report on system security settings and configuration, including monitoring the integrity of important files by collecting a baseline measurement, comparing it to a current measurement and reporting on any discrepancies found. XSW can perform that comparison either periodically or on demand.

Centralized monitoring and reporting and file integrity checking both assist you in demonstrating compliance to security regulations such as PCI DSS, SOX and HIPAA.

Hardening against known vulnerabilities

HPE pays serious attention to reports of security vulnerabilities in its products. For NonStop servers, the normal procedure is to issue product or documentation fixes as needed and to inform customers of their availability through both an HPE Security Bulletin (HPESB) and a NonStop Hotstuff. For major vulnerabilities in open source used in NonStop products, HPE usually publishes information about known impact while investigations are still in progress. Increasingly, HPE also is making early information on the impact of other open source vulnerabilities available to the GNSC prior to publication of Hotstuffs and HPE Security Bulletins. For vulnerabilities identified in NonStop-originated code, HPE typically does not publish any information until fixes are available, and avoids publishing a level of detail that might make it obvious how to exploit the vulnerability.

HPE Security Bulletins include:

- Characterization of the effects of a successful exploit (for example: Denial of Service, Unauthorized Access, and/or Gain Privileged Access)
- The CVE identifier (see CVESITE for more information on CVE identifiers, also called "CVE-IDs," "CVE names," "CVE numbers," and "CVEs", which are unique, common identifiers for publicly known security vulnerability information). A CVE identifier will be assigned even for internally discovered vulnerabilities.
- The CVSS score (See CVSS for more information on CVSS scores.)
- Affected products and versions
- Resolution (TCF SPR, documentation update, or, in rare cases, a workaround)
- Information on how to identify affected systems and how to determine that the vulnerability has been resolved

The link to HPE product security vulnerability alerts is https://www.hpe.com/us/en/services/security-vulnerability.html. From this site you can subscribe to alerts for the products that you use and also view HPE-wide security advisories for certain critical vulnerabilities and access the HPE Security Bulletin archive.

NonStop Hotstuffs for security vulnerabilities contain similar information, along with the identifier for the associated HPESB. They are sent out via ExpressNotice and are available in Scout for NonStop customers.

HPE encourages you to notify us of any potential security defects that you may encounter in HPE products. To report a security defect to HPE, send email to: security-alert@HPE.com.

HPE strongly recommends that you keep all of your security products and packages, from whatever source, as up to date as possible in order to minimize exposure to vulnerabilities.

Viruses and Virus protection

The NonStop server has some inherent architectural advantages that greatly reduce both the potential for malevolent software to be imported and the effects if such software were introduced to the system. As described in the HPE NonStop Security Overview (HPESECOV), nonprivileged (user) processes cannot modify either their own object code (binaries) or those of other processes. As long as you have properly secured your NonStop system, its object files also cannot be modified by a non-privileged process. As a consequence, even if a virus written for a UNIX, LINUX, or POSIX environment on any other platform were introduced to your NonStop server, it could not have its intended effect. For example, common UNIX attacks that exploit buffer overflows to gain root access capabilities will not work on a NonStop Server because a non-privileged process cannot escalate its set of privileges.

Additional references

Books:

(XYPRO1) Securing HPE NonStop Servers in an Open System World ISBN-13: 978-1-55558-344-6
(XYPRO2) HPE NonStop Server Security – A Practical Handbook ISBD -13: 978-1-55558-314-9

Note: While the books listed above are useful references, both were published some time ago, so they do not reflect more recent security enhancements in Safeguard and other subsystems or current thinking about certain best practices such as minimum/maximum password length.

Web:

Links to third-party websites are provided solely as a convenience and HPE does not endorse or make any representations about such websites.

Middleware:

An overview of Java security technology:

(NSJSEC) http://www.oracle.com/technetwork/java/javase/tech/index-jsp-136007.html

SOAP WS-Security standards:

(SOAPMSG) SOAP Message Security V1.0 (http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf)

(SOAPUSER) Username Token Profile V1.0 (http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-username-token-profile-1.0.pdf)

(SOAPCERT) X.509 Certificate Token Profile V1.0 (http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0.pdf

Payment Card Industry Data Security Standard:
PCI DSS version 3.2:

(PCIDSSV32) https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2.pdf

A consultant's view of approaches to meeting PCI DSS V3 on NonStop servers

(KNIGHT2) http://www.knightcraft.com/pci-dss-3

An interesting discussion of the differences between being secure and being compliant:

(KNIGHT1) http://www.knightcraft.com/2014-HPE-nonstop-advanced-technical-boot-camp:

An e-book and two downloads on SSL and TLS:

(SSLTLS) Bulletproof SSL – https://www.feistyduck.com/books/bulletproof-ssl-and-tls

(SSLCOOKBOOK) https://www.feistyduck.com/books/openssl-cookbook

(SSLDEP) https://www.ssllabs.com/projects/best-practices/index.html

NSC security:

(NSCRDWS2008) https://technet.microsoft.com/en-us/library/cc770833(v=ws.11).aspx

(NSCRDWS2012) https://blogs.technet.microsoft.com/askperf/2012/10/30/windows-8-windows-server-2012-remote-desktop-management-server/

Common vulnerability identification and scoring:

All things CVE:

(CVESITE) http://cve.mitre.org

An explanation of how HPE does CVSS vulnerability scoring:

(CVSS) HPE Customer Notice: HPSN-2008-002 at
http://h20000.www2.HPE.com/bizsupport/TechSupport/Document.jsp?lang=en&cc=us&objectID=c01345499

Other topics:

(SAFEPASS) http://wolfram.org/writing/howto/password.html

NonStop Security Partners:

NonStop security partners include:

3rd Data: www.thirddata.com

4tech Software: www.4techsoftware.com

ACI: www.aci.com

Bowden: www.bsi2.com

Brightstrand: www.brightstrand.com

CAIL: www.cail.com

comForte 21: www.comforte.com

Crystal Point: www.crystalpoint.com

CSP Security: www.cspsecurity.com

Crossroads: www.crossroads.com

ETI-NET: www.etinet.com

GreenHouse: www.greenhouse.de

IdentityForge: www.identityforge.com

Knightcraft Technology: www.knightcraft.com

Opsol Integrators: www.opsol.com

TSI: www.tributary.com

XYPRO Technologies: www.xypro.com

Visit their websites for product information, white papers, services, etc.