

EAL3 Evaluated Configuration Guide for Red Hat Enterprise Linux

Klaus Weidner <klaus@atsec.com>

June 28, 2004; v1.1

atsec is a trademark of atsec GmbH

IBM, IBM logo, BladeCenter, eServer, iSeries, OS/400, PowerPC, POWER3, POWER4, POWER4+, pSeries, S390, xSeries, zSeries, zArchitecture, and z/VM are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based products are trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds.

UNIX is a registered trademark of The Open Group in the United States and other countries.

This document is provided AS IS with no express or implied warranties. Use the information in this document at your own risk.

This document may be reproduced or distributed in any form without prior permission provided the copyright notice is retained on all copies. Modified versions of this document may be freely distributed provided that they are clearly identified as such, and this copyright is included intact.

Copyright (c) 2003, 2004 by atsec GmbH, and IBM Corporation or its wholly owned subsidiaries.

Contents

1	Introduction	5
1.1	Purpose of this document	5
1.2	How to use this document	5
1.3	What is a CC compliant system?	5
1.3.1	Hardware requirements	6
1.3.2	Software requirements	6
1.3.3	Environmental requirements	6
1.3.4	Operational requirements	6
1.4	Requirements for the system's environment	6
1.5	Requirements for the system's users	7
1.6	Overview of the system's security functions	8
1.6.1	Identification and authentication	8
1.6.2	Audit	8
1.6.3	Discretionary access control	8
1.6.4	Object reuse	8
1.6.5	Security management and system protection	8
1.6.6	Secure communication	8
1.7	Overview of security relevant events	9
2	Installation	9
2.1	Supported hardware	9
2.2	Selection of install options and packages	10
3	Secure initial system configuration	12
3.1	Creating additional user accounts for administrators	13
3.2	Installing required updates	13
3.3	Automated configuration of the system	15
3.4	Configuring filesystem parameters	16
3.5	Add and remove packages	17
3.6	Disable services	20
3.7	Remove SUID/SGID root settings from binaries	20
3.8	Update permissions for 'su'	21
3.9	Disable root login over the network	21
3.10	Setting up SSH	22
3.11	Setting up xinetd	23
3.12	Setting up FTP	23
3.13	Setting up Postfix	24
3.14	Setting up the audit subsystem	24
3.14.1	Installing the packages needed for auditing	24
3.14.2	Setting up the audit configuration files	25
3.14.3	Starting auditd at boot as a system service	25
3.14.4	Starting auditd in fail-secure mode from init (OPTIONAL)	25
3.15	Introduction to Pluggable Authentication Module (PAM) configuration	26
3.16	Required Pluggable Authentication Module (PAM) configuration	27
3.16.1	/etc/pam.d/system-auth	27
3.16.2	/etc/pam.d/login	28
3.16.3	/etc/pam.d/other	29
3.16.4	/etc/pam.d/sshd	29
3.16.5	/etc/pam.d/su	30
3.16.6	/etc/pam.d/vsftpd	30
3.17	Configuring default account properties	31
3.18	Configuring the boot loader	32

3.18.1	GRUB boot loader configuration	32
3.18.2	Yaboot boot loader configuration	33
3.18.3	ZIPL boot loader configuration	33
3.18.4	iSeries kernel slots	34
3.19	Reboot and initial network connection	34
4	System operation	34
4.1	System startup, shutdown and crash recovery	35
4.2	Backup and restore	35
4.3	Gaining superuser access	36
4.4	Installation of additional software	36
4.5	Scheduling processes using cron and at	37
4.6	Mounting filesystems	38
4.7	Managing user accounts	39
4.8	Using serial terminals	41
4.9	SYSV shared memory and IPC objects	41
4.10	Configuring secure network connections with <i>stunnel</i>	41
4.10.1	Introduction	41
4.10.2	Creating an externally signed certificate	42
4.10.3	Creating a self-signed certificate	44
4.10.4	Activating the tunnel	45
4.10.5	Using the tunnel	46
4.10.6	Example 1: Secure SMTP delivery	46
4.10.7	Example 2: Simple web server	47
4.10.8	Example 1: system status view	47
4.11	The Abstract Machine Testing Utility (AMTU)	48
4.12	Setting the system time and date	49
5	Monitoring, Logging & Audit	49
5.1	Reviewing the system configuration	49
5.2	System logging and accounting	50
5.3	Configuring the audit subsystem	51
5.3.1	Intended usage of the audit subsystem	51
5.3.2	Selecting the events to be audited	51
5.3.3	Reading and searching the audit records	52
5.3.4	Starting and stopping the audit subsystem	52
5.3.5	Storage of audit records	53
5.3.6	Reliability of audit data	53
5.4	System configuration variables in <i>/etc/sysconfig</i>	53
6	Security guidelines for users	54
6.1	Online Documentation	54
6.2	Authentication	54
6.3	Password policy	55
6.4	Access control for files and directories	56
6.5	Data import / export	57
7	Appendix	57
7.1	Online Documentation	57
7.2	Literature	58
7.3	The file <i>/etc/audit/audit.conf</i>	58
7.4	The file <i>/etc/audit/filter.conf</i>	60
7.5	The file <i>/etc/audit/filesets.conf</i>	68

1 Introduction

1.1 Purpose of this document

The Red Hat Enterprise Linux (RHEL) distribution is designed to provide a secure and reliable operating system for a variety of purposes. Because security requirements obviously depend on the applications and environment, it is not possible to simply certify that the system is "secure", a more precise definition is needed.

The Common Criteria (CC) provides a widely recognized methodology for security certifications. A CC evaluation is fundamentally a two-step process, consisting of defining the "security target" which describes the features that are to be evaluated, and then testing and verifying that the system actually implements these features with a sufficient level of assurance.

This document is a security guide that explains how to set up the evaluated configuration, and provides information to administrators and ordinary users to ensure secure operation of the system. It is intended to be self-contained in addressing the most important issues at a high level, and refers to other existing documentation where more details are needed.

The document primarily addresses administrators, but the section "Security guidelines for users" is intended for ordinary users of the system as well as administrators.

Knowledge of the Common Criteria is not required for readers of this document.

1.2 How to use this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 <<http://www.ietf.org/rfc/rfc2119.txt>>.

Note that the terms "SHOULD" and "SHOULD NOT" are avoided in this document. Requirements are either absolute (and marked with MUST and equivalent terms), or entirely optional (in the sense of not affecting required security functions) and marked with RECOMMENDED, MAY or OPTIONAL.

If you follow the requirements in this document when setting up and using the system, your configuration will match the evaluated configuration. Certain configuration options are marked as OPTIONAL and you MAY modify them as needed, but you MUST NOT make other changes, because they will make the system fail to match the evaluated configuration.

Of course, you MUST always use common sense. This document is not a formal specification, and legitimate reasons may exist to modify the system setup in ways not described here if that is necessary for the system to fulfill its intended purpose. Specifically, applying security patches released by the vendor is strongly RECOMMENDED even though that will cause a deviation from the evaluated configuration.

In cases where the requirements and recommendations in this document conflict with those in other sources (such as the online documentation), the information in this Configuration Guide has higher precedence. You MUST follow the steps described here to reach the evaluated configuration, even if other documentation describes different methods.

The usual convention is used in this guide when referring to manual pages that are included in the software distribution. For example, the notation `ls(1)` means that running the `man -S 1 ls` command will display the manual page for the `ls` command from section one of the installed documentation. In most cases, the `-S` flag and the section number may be omitted from the command, they are only needed if pages with the same name exist in different sections,

1.3 What is a CC compliant system?

A system can be considered to be "CC compliant" if it matches an evaluated and certified configuration. This implies various requirements concerning hardware and software, as well as requirements concerning the operating environment, users, and the ongoing operating procedures.

Strictly speaking, an evaluation according to the CC represents the results of investigation of the security properties of the target system according to defined guidelines. It should not be considered as a guarantee for fitness for any specific purpose, but should provide help in deciding the suitability of the system considering how well the intended use fits the described capabilities. It is intended to provide a level of assurance about the security functions that have been examined by a neutral third party.

1.3.1 Hardware requirements

The hardware **MUST** be the one of the following IBM systems:

IBM xSeries 335 - AS and WS

IBM pSeries 630 - AS only

IBM iSeries 825 type 9406 (OS/400 V5R2 LPAR) - AS only

IBM zSeries 900 (within VM 4.3 logical partition) - AS only

IBM eServer 325 (AMD64 Opteron) - AS only

Running the certified software on other similar hardware may result in an equivalent security level, but the certification does not apply if the hardware is different from that used for the testing processes during the evaluation.

1.3.2 Software requirements

The software **MUST** match the evaluated configuration. In the case of an operating system, this also requires that the installed kernel, system, and application software are the same. The documentation (including this guide) will specify permitted variations, such as modifying certain configuration files and settings, and installing software that does not have the capability to affect the security of the system (typically those that do not require 'root' privileges).

1.3.3 Environmental requirements

Stated requirements concerning the operating environment **MUST** be met. Typical requirements include a secure location for the hardware (protected from physical access by unauthorized persons), as well as restrictions concerning permitted network connections.

1.3.4 Operational requirements

The operation of the system **MUST** be in agreement with defined organizational security policies, to ensure that actions by administrators and users do not undermine the system's security.

1.4 Requirements for the system's environment

The security target covers one or more systems running RHEL, networked in a non-hostile network, with a well-managed and non-hostile user community. It is not intended to address the needs of an Internet-connected server, or the case where services are to be provided to potentially hostile users.

You **MUST** set up the server (or servers) in a physically secure environment, where they are protected from theft and manipulation by unauthorized persons.

You **MUST** ensure that all connections to peripheral devices and all network connections are protected against tampering, tapping and other modifications. Using the secured protocols SSHv2 or SSLv3 is considered sufficient protection for network connections. All other connections must remain completely within the physically secure server environment.

All components in the network such as routers, switches, and hubs that are used for communication are assumed to pass the user data reliably and without modification. Translations on protocols elements (such as NAT) are allowed as long as those modifications do not lead to a situation where information is routed to somebody other than the intended recipient system.

Be aware that information passed to another system leaves the control of the sending system, and the protection of this information against unauthorized access needs to be enforced by the receiving system. If an organization wants to implement a consistent security policy covering multiple systems on a network, organizational procedures **MUST** ensure that all those systems can be trusted and are configured with compatible security configurations enforcing an organization wide security policy. How to do this is beyond the scope of this Configuration Guide. If you set up a communication link to a system outside your control, please keep in mind that you will not be able to enforce any security policy for any information you pass to such a system over the communication link or in other ways (for example, by using removable storage media).

Every person that has the ability to perform administrative actions by switching to root has full control over the system and could, either by accident or deliberately, undermine the security of the system and bring it into an insecure state. This Configuration Guide provides the basic guidance how to set up and operate the system securely, but is not intended to be the sole information required for a system administrator to learn how to operate Linux securely.

It is assumed, within this Configuration Guide, that administrators who use this guide have a good knowledge and understanding of operating security principles in general and of Linux administrative commands and configuration options in particular. We strongly advise that an organization that wants to operate the system in the evaluated configuration nevertheless have their administrators trained in operating system security principles and RHEL security functions, properties, and configuration.

Every organization needs to trust their system administrators not to deliberately undermine the security of the system. Although the evaluated configuration includes audit functions that can be used to make users accountable for their actions, an administrator is able to stop the audit subsystem and reconfigure it such that his actions no longer get audited. Well trained and trustworthy administrators are a key element for the secure operation of the system. This Configuration Guide provides the additional information a system administrator should obey when installing, configuring and operating the system in compliance with the requirements defined in the Security Target for the Common Criteria evaluation.

1.5 Requirements for the system's users

The security target addresses the security needs of cooperating users in a benign environment, who will use the system responsibly to fulfill their tasks.

Note that system availability is *not* addressed in this evaluation, and a malicious user could disable a server through resource exhaustion or similar methods.

The requirements for users specifically include:

- User accounts **MUST** be assigned only to those users with a need to access the data protected by the system, and who **MUST** be sufficiently trustworthy not to abuse those privileges. For example, the system cannot prevent data from being intentionally redistributed to unauthorized third parties by an authorized user.
- All users of the system **MUST** be sufficiently skilled to understand the security implications of their actions, and **MUST** understand and follow the requirements listed in section §6 "Security guidelines for users" of this guide. Appropriate training **MUST** be available to ensure this.

It is part of your responsibility as a system administrator to verify that these requirements are met, and to be available to users if they need your help in maintaining the security of their data.

1.6 Overview of the system's security functions

This section summarizes the security functions that were covered by the evaluation. Please refer to the appropriate sections for information on configuring, using and managing these functions.

1.6.1 Identification and authentication

Pluggable Authentication Module (PAM)

Sections §3.15 "Introduction to Pluggable Authentication Module (PAM) configuration", §3.16 "Required Pluggable Authentication Module (PAM) configuration" of this guide; and the documentation in */usr/share/doc/pam*/* and the *pam(8)* man page.

OpenSSH

Section §3.10 "Setting up SSH" of this guide; and the *sshd(8)*, *ssh(1)*, *sshd_config(5)* man pages.

vsftpd

Section §3.12 "Setting up FTP" of this guide; and the *vsftpd(8)*, *vsftpd.conf(5)* man pages.

su

Sections §3.8 "Update permissions for 'su'", §4.3 "Gaining superuser access" of this guide; and the *su(8)* man page.

1.6.2 Audit

Sections §3.14 "Setting up the audit subsystem" and §5.3 "Configuring the audit subsystem" of this guide; and the *laus(7)* man page, whose "SEE ALSO" section points to the remaining LAuS man pages.

1.6.3 Discretionary access control

Sections §6.4 "Access control for files and directories" and §4.9 "SYSV shared memory and IPC objects" of this guide.

1.6.4 Object reuse

See the RHEL High Level Design document, the kernel automatically ensures that new objects (disk files, memory, IPC) do not contain any traces of previous contents.

1.6.5 Security management and system protection

Chapters §4 "System operation" and §5 "Monitoring, Logging & Audit".

1.6.6 Secure communication

Section §4.10 "Configuring secure network connections with *stunnel*" of this guide; and the *stunnel(1)* man page.

Section §3.10 "Setting up SSH" of this guide; and the *sshd(8)*, *ssh(1)*, and *sshd_config(5)* man pages.

1.7 Overview of security relevant events

The audit subsystem is intended to be the central interface for collecting and viewing the record of security relevant events. The events being monitored by default in the evaluated configuration include:

- All authentication done through the PAM library, including the identity and location (where available) of the user and the success or failure result.
- Use of *su(8)* to change identity. All actions done as part of a *su* session are marked in the audit record with the original user's login user ID.
- Adding, changing, or deleting users or groups.
- Changes and change attempts to the contents of security critical files.
- Changes to the access permissions or ownership of any files or IPC objects.
- Binding network ports and accepting connections.

Please refer to section §5 "Monitoring, Logging & Audit" for more information.

2 Installation

The evaluation covers a fresh installation of RHEL AS or WS, Version 3 Update 2, on one of the supported hardware platforms as defined in section §1.3.1 "Hardware requirements" of this guide.

On the platforms that support virtualization (VM) or secure logical partitioning (LPAR), other operating systems MAY be installed and active at the same time as the evaluated configuration. This is if (and only if) the VM or LPAR configuration ensures that the other operating systems cannot access data belonging to the evaluated configuration or otherwise interfere with its operation. Setting up this type of configuration is considered to be part of the operating environment and is not addressed in this guide.

On the other platforms, the evaluated configuration MUST be the only operating system installed on the server.

2.1 Supported hardware

You MAY attach the following peripherals without invalidating the evaluation results. Other hardware MUST NOT be installed in or attached to the system.

- Any storage devices and backup devices supported by the operating system (this includes hard disks, CD-ROM drives and tape drives).
- All Ethernet and Token Ring network adapters supported by the operating system. Modems, ISDN and other WAN adapters are not part of the evaluated environment.
- Any printers supported by the operating system.
- Operator console consisting of a keyboard, video monitor, and optionally mouse. Additionally, you MAY directly attach supported serial terminals (see section §4.8 "Using serial terminals" of this guide), but *not* modems, ISDN cards, or other remote access terminals.

Hot-pluggable hardware that depends on the dynamic loading of kernel modules is *not* supported. Examples of such unsupported hardware are USB and, IEEE1394/FireWire peripherals.

2.2 Selection of install options and packages

This section describes the detailed steps to be performed when installing the RHEL operating system on the target server.

All settings listed here are **REQUIRED** unless specifically declared otherwise.

1. It is **RECOMMENDED** that you disconnect all network connections until the post-install system configuration is finished. You **MAY** use a network if required for the installation (for example, zSeries hosts are usually installed using NFS, because they do not have a CD drive). If you do use a network, you **MUST** ensure that this network is secure, for example by directly connecting the new system to a standalone NFS server with no other network connections.
2. You **MUST** download the ISO images from the RedHat Network on a separate Internet-connected computer, and either burn CD-Rs from them, or make the contents available on a NFS file server. The download location https://rhn.redhat.com/network/software/download_isos_full.pxt contains links to the platform-specific images. You **MUST** use **Red Hat Enterprise Linux 3 Update 2**, either **AS** (Advanced Server) or **WS** (Workstation). Make sure that you are using the appropriate version for your platform:

xSeries	i386	AS or WS
eServer 325	x86_64	AS
pSeries	ppc	AS
iSeries	ppc	AS
zSeries	s390	AS

You **MUST** verify that the MD5 checksums of the image files are correct. Run `md5sum *.iso` to view the checksums for the downloaded images, and compare them with those shown in this list:

```
Red Hat Enterprise Linux 3 AS (i386) Update 2
Binary Disc 1    139M    240a6a59d7f2f1c4bb2d017016e46f5c
Binary Disc 2    619M    a45ce3753d0c9a599d4533c8ba4339d2
Binary Disc 3    633M    5bc126d663e681f874a15f55e05c070f
Binary Disc 4    168M    158f58ae5e1307415185e0b4863b40e4
```

```
Red Hat Enterprise Linux 3 WS (i386) Update 2
Binary Disc 1    131M    211692ca8fac5894f4f98d1a744f4c27
Binary Disc 2    619M    a45ce3753d0c9a599d4533c8ba4339d2
Binary Disc 3    633M    5bc126d663e681f874a15f55e05c070f
Binary Disc 4    168M    158f58ae5e1307415185e0b4863b40e4
```

```
Red Hat Enterprise Linux 3 AS (X86_64) Update 2
Binary Disc 1    142M    32f28eb619907faca2a94bb8e7658dae
Binary Disc 2    625M    2919142a9f3bb46572815bdb327baddc
Binary Disc 3    637M    be55fc8e56d3f1eb53d87f46f97b063b
Binary Disc 4    204M    2d7530d57c5c407bb2d820361a48bb18
```

```
Red Hat Enterprise Linux 3 AS (ppc) Update 2
Binary Disc 1    172M    5ed53761d9dfb69568cb7077a9f85f0e
Binary Disc 2    619M    d75085869102710c9803603d55c669af
Binary Disc 3    638M    ab819d81efced28099e3f6ebcf38de52
Binary Disc 4    39M     4089bad2a99ab560cf55ad88e40309da
```

```

Red Hat Enterprise Linux 3 AS (s390) Update 2
  Binary Disc 1    130M    df70a9aaa1a457306360c7d31ba74711
  Binary Disc 2    623M    6d6b3a93d38cfadc1553687bfa05387a
  Binary Disc 3    513M    b046f680a279479e3d9d63c53ffcdb42
  Binary Disc 4    376K    2a220f9dc346075c6037e73a40e657f0

```

3. Launch the installer program contained on the CD-ROM. The details of how to do this depend on the hardware platform, please refer to the installation guide that is part of the printed manual accompanying the CD.

For example:

- xSeries, eServer 325 (Opteron), pSeries: Insert the first CD and boot from CD-ROM.
 - zSeries, iSeries: Details depend on the operation mode (VM, LPAR or native). The process generally involves copying the installer onto the server and launching the installer using the host's management interface.
4. You MAY choose text-mode installation instead of the default graphical installation by entering `linux text` at the boot prompt.
You MAY also use a serial console to do a text-mode installation. To do so, connect a serial terminal (or a computer with terminal emulator software; such a computer MUST be appropriately secure) to the server's serial port, and boot from the RHEL CD. When the boot prompt is shown on the serial console, enter `linux text console=ttyS0` (use the appropriate name of the serial device if not using `ttyS0`) and press ENTER to start the installation.
 5. Running the CD **media test** for all installation CDs is RECOMMENDED.
 6. When the "Welcome" screen appears, press **Next**.
 7. Language Selection: choose **English (English)** to ensure that the messages shown during the installation match those described in this guide.
 8. Keyboard Configuration: You MAY change the **U.S. English** setting to match your keyboard.
 9. Mouse Configuration: You MAY change the **Mouse Selection** if the autodetected values are not appropriate, including choosing "No mouse" and using the keyboard only.
 10. Disk Partitioning Setup. Use **Manual partition with Disk Druid** to set up the partitions. For CAPP-compliant auditing, you MUST set up a separate partition for the directory `/var/log/`.

- Set up the **REQUIRED / (root)** and `/var/log` partitions, and as many additional mounted partitions as appropriate. `/var/log` **REQUIRES** at least 100 MB of space in order to be able to install and launch the audit system, but this does not include the additional space needed for saved audit logs, please refer to section §5.3 "Configuring the audit subsystem" of this guide for more information.

It is RECOMMENDED to also use separate partitions for `/var`, `/home` and `/tmp`. Some configurations (recognized automatically by the installation program) need a separate `/boot` partition. The following table shows a RECOMMENDED partitioning scheme together with minimum sizes for the partitions. Using more space is RECOMMENDED:

<code>/boot</code>	75 MB
<code>/</code>	1200 MB
<code>/tmp</code>	200 MB
<code>/home</code>	100 MB
<code>/var</code>	384 MB
<code>/var/log</code>	100 MB needed for install, >>1GB for use

- All mounted partitions MUST be of type **ext3** and **formatted**.
- Configuring a swap partition at least as large as the installed RAM is RECOMMENDED.

11. Boot Loader Configuration: Setting a **boot loader password** is RECOMMENDED. You MUST NOT add other operating systems to the configuration.
12. Network Configuration: Configure all installed network cards (zero or more) as appropriate for the platform. In the case of virtual network cards on zSeries or iSeries, these options are not available. The following options MUST be used for non-virtual network cards:
 - Use the **Edit** button to either assign a static IP address by unchecking the **Configure using DHCP** box and entering the **IP Address** and **Netmask**; or alternatively disable the card by unchecking the **Activate on boot** box.
 - You MUST NOT use DHCP for any active network card.
 - Enter a valid **hostname** (which is RECOMMENDED to be unique within your network) consisting of one or more alphanumeric components, separated by '.', and each matching the regular expression `[a-zA-Z][a-zA-Z0-9]*`
 - OPTIONAL: assign a **Gateway** address and **DNS** servers.
 - Modems and ISDN adapters MUST NOT be present.
13. Firewall Configuration: RECOMMENDED to select **No firewall** for the evaluated configuration, it is not needed on a hardened minimal installation and may cause conflicts with the use of *stunnel*(8) for secure connections. You MAY enable the firewall and choose a list of permitted ports.
14. Additional Language Support: RECOMMENDED to leave the **default language** set as **English (USA)** to ensure that system messages match those described in the documentation. (Note that users can individually override this setting.) You MAY add additional language support.
15. Time Zone Selection: RECOMMENDED to set the **Location** or **UTC Offset** as appropriate for the server's location, and RECOMMENDED to activate **System clock uses UTC**.
16. Set Root Password: Choose a **Root Password** according to the password policy (§6.3), and **Confirm** it.
17. Package Installation Defaults: Select **Customize the set of packages to be installed**. When using the graphical installer, put a check mark on the **Minimal** set of packages (last item, in the **Miscellaneous** group), this will deselect all the other package selections. If using the text-mode installer, you MUST manually deselect all packages by removing all check marks in the **Package Group Selection** dialog.
18. About to Install: This is the final confirmation to start the installation. Press **Next** to start the automated partitioning, formatting, and installation process. Insert additional disks if prompted to do so. For the x86 architecture, all the needed data is on the first CD-ROM, so you will not be asked to insert additional disks.
19. When the automated install is complete and the **Congratulations** screen appears, pressing **Exit** will reboot the system. It is RECOMMENDED that you now reconfigure the system to boot from the newly installed system only (typically the first hard disk) and disable all other boot methods such as CD-ROM, network boot (PXE) or floppy disk. If you choose not to do that, you MUST remove the installation CD-ROM from the drive before rebooting.
20. Wait for the freshly installed system to start, and verify that the issue message printed above the login prompt matches the installed system type and version. Then log in as "root" and proceed with the next section.

3 Secure initial system configuration

After the initial installation, the operating system is not yet in the evaluated configuration. The instructions in this section explain how to achieve that configuration.

After software upgrades or installation of additional packages, these steps MUST be re-done or at least re-checked to ensure that the configuration remains secure.

Log in as user 'root' on the system console for these steps.

3.1 Creating additional user accounts for administrators

The evaluated configuration disables direct "root" login over the network. All system administrators **MUST** log in using a non-root individual user ID, then use the `su(8)` command to gain superuser privileges for administrative tasks. This requires membership in the 'wheel' group of trusted users.

You **MUST** define at least one non-root user account with the `useradd(8)` command, and add this user account to the 'wheel' group. Note that the enhanced password quality checking mechanisms and the password expiry settings of the evaluated configuration are not active yet. You must manually set the password properties in accordance with the password policy.

Here is an example of creating an additional user account:

```
useradd -m -c "John Doe" -G wheel jdoe
passwd jdoe
chage -m 1 -M 60 -W 7 jdoe
```

Please refer to sections §4.7 "Managing user accounts" and §6.3 "Password policy" of this guide for more information on creating user accounts.

3.2 Installing required updates

You need to download several additional packages not included in Update 2 to set up the evaluated configuration. The packages are available at the following location:

```
ftp://partners.redhat.com/EAL3_RHEL3/U2/
```

Download the RPMs using an Internet-connected computer, and transfer the RPMs to the system being installed.

You **MUST** select the appropriate RPM packages for your architecture. The 64bit architectures support execution of both 64bit and 32bit binaries.

xSeries (i386)

This is a 32bit-only platform. Use ***.i686.rpm** variants of packages if available, ***.i386.rpm** or ***.noarch.rpm** otherwise.

eSeries 325 (Opteron/x86_64)

This system uses a 64bit kernel and 64bit userspace programs, and also supports running 32bit programs. Use the ***.x86_64.rpm** or ***.noarch.rpm** variants of packages. You may **OPTIONALLY** install the ***.i386.rpm** or ***.i686.rpm** variants of libraries (package names containing *-libs* or *-devel*) in addition to the 64bit versions.

iSeries or pSeries (ppc/ppc64)

These systems use a 64bit kernel, but the installed userspace programs are the 32bit variants. They support running 64bit programs as well. There are separate kernels available for the iSeries and pSeries platforms, use **ppc64iseries.rpm** or **ppc64pseries** as appropriate. Use the ***.ppc64.rpm** variants of *laus* and *laus-libs*, ***.ppc.rpm** or ***.noarch.rpm** for everything else. You may **OPTIONALLY** install the ***.ppc64.rpm** variants of libraries (package names containing *-libs* or *-devel*) in addition to the 32bit versions, and the ***.ppc.rpm** version of *laus-libs* in addition to the default 64bit version.

zSeries (s390)

The evaluated configuration uses a 31bit kernel running 31bit userspace programs. Install the ***.s390.rpm** or ***.noarch.rpm** variants only. Use of 64bit packages (***.s390x.rpm**) is **NOT** supported in the evaluated configuration.

The active kernel MUST be one of the listed *kernel* or *kernel-smp* packages. It is RECOMMENDED that you uninstall unused kernel packages, such as the uniprocessor kernel on a SMP machine. The development libraries (**-devel**) and additional non-default word size libraries as explained above are OPTIONAL. All other packages listed here are REQUIRED. You MUST verify the MD5 sums against the following list:

```
### i386
d61973789abea575dfdec8dd941b1446  amtu-0.1-3RHEL.i386.rpm
34edf418f9e2dc80e681b6ca38cb854a  at-3.1.8-48.ent.i386.rpm
00744d4d35e23b78fc6deb49a187c951
    kernel-2.4.21-15.0.2.EL.peterm.eal.3.i686.rpm
827f8140f520ff4677bbde524edd6e69
    kernel-smp-2.4.21-15.0.2.EL.peterm.eal.3.i686.rpm
cd31e226ffc751e4a490c4c46716c0b6
    kernel-source-2.4.21-15.0.2.EL.peterm.eal.3.i386.rpm
942d4e9687cc229ee7866e0810294ed9  laus-0.1-62RHEL3.i386.rpm
79ac3ab71e1917cc464dc2f3b1b9edd1  laus-devel-0.1-62RHEL3.i386.rpm
58e5e19bf98d287576ad7d0c50d3ff74  laus-libs-0.1-62RHEL3.i386.rpm
6789218233e890ce2c5005ce70abea99  passwd-0.68-3.1.i386.rpm
cd62831bd278541bac99f13edf3a93f9  shadow-utils-4.0.3-20.03.i386.rpm
af2f1ff1abb97e16b92b51b89ed69212  vixie-cron-3.0.1-75.i386.rpm

### iSeries
96b28d70770025a202a2c496b4106087  amtu-0.1-3RHEL.ppc.rpm
63935344db78a3e7173549fe4f7f57ef  at-3.1.8-48.ent.ppc.rpm
1d33d4613000f63c29426b8d82dc6560
    kernel-2.4.21-15.0.2.EL.peterm.eal.3.ppc64iseries.rpm
b54313d999396e92627b6d33c7c51058
    kernel-source-2.4.21-15.0.2.EL.peterm.eal.3.ppc64.rpm
8f17734994b40665c02f77f1d19ae9d5  laus-0.1-62RHEL3.ppc64.rpm
c6129b9d09330794a5f329acff7def6e  laus-devel-0.1-62RHEL3.ppc.rpm
91c6217bald27dbefee33452f31423d6  laus-devel-0.1-62RHEL3.ppc64.rpm
2db58f704ade9835b55077a7537546a5  laus-libs-0.1-62RHEL3.ppc.rpm
86dc0501ba9e8942010678806aea9e8c  laus-libs-0.1-62RHEL3.ppc64.rpm
1b1203051fe38a4c7d6714948b7cf348  passwd-0.68-3.1.ppc.rpm
e7edac75f05e107df39b82ade751f26f  shadow-utils-4.0.3-20.03.ppc.rpm
38bbec2cb8070032130a69d352dc70dd  vixie-cron-3.0.1-75.ppc.rpm

### pSeries
96b28d70770025a202a2c496b4106087  amtu-0.1-3RHEL.ppc.rpm
63935344db78a3e7173549fe4f7f57ef  at-3.1.8-48.ent.ppc.rpm
519d8724bdd022cc18465b8323b8e5c9
    kernel-2.4.21-15.0.2.EL.peterm.eal.3.ppc64pseries.rpm
b54313d999396e92627b6d33c7c51058
    kernel-source-2.4.21-15.0.2.EL.peterm.eal.3.ppc64.rpm
8f17734994b40665c02f77f1d19ae9d5  laus-0.1-62RHEL3.ppc64.rpm
c6129b9d09330794a5f329acff7def6e  laus-devel-0.1-62RHEL3.ppc.rpm
91c6217bald27dbefee33452f31423d6  laus-devel-0.1-62RHEL3.ppc64.rpm
2db58f704ade9835b55077a7537546a5  laus-libs-0.1-62RHEL3.ppc.rpm
86dc0501ba9e8942010678806aea9e8c  laus-libs-0.1-62RHEL3.ppc64.rpm
1b1203051fe38a4c7d6714948b7cf348  passwd-0.68-3.1.ppc.rpm
e7edac75f05e107df39b82ade751f26f  shadow-utils-4.0.3-20.03.ppc.rpm
38bbec2cb8070032130a69d352dc70dd  vixie-cron-3.0.1-75.ppc.rpm
```

```
### s390
```

```

fe1b09787097174edf594af8a10fccb9  amt-0.1-3RHEL.s390.rpm
491ca763be4d1d72bd8b1f9482063623  at-3.1.8-48.ent.s390.rpm
d96c577a3391feb227018d65f3c114b4  kernel-2.4.21-15.0.2.EL.peterm.eal.3.s390.rpm
4bf34cc066affc96937af09cdc7d6bfe  kernel-source-2.4.21-15.0.2.EL.peterm.eal.3.s390.rpm
73de7e7a4a57b15f9feaf81d62f41a9f  laus-0.1-62RHEL3.s390.rpm
298b8ae4714d9fbca2c279452a13a0ee  laus-devel-0.1-62RHEL3.s390.rpm
ce7cb4b2bb5789992889047794dbb70d  laus-libs-0.1-62RHEL3.s390.rpm
bf1c3950d090d33386e9ef7f2e4f91c7  passwd-0.68-3.1.s390.rpm
074298576ec409b7923509a12983442a  shadow-utils-4.0.3-20.03.s390.rpm
b2feeb0dd26a9c791cc342066357bc0b  vixie-cron-3.0.1-75.s390.rpm

### x86_64
480a11df2d2f02df60641588f61fd9e9  amt-0.1-3RHEL.x86_64.rpm
c7ca492b0266ae83f92f09a6d4b94ab0  at-3.1.8-48.ent.x86_64.rpm
77836e69f9cc29e1220dae6e3d3e42e0
    kernel-2.4.21-15.0.2.EL.peterm.eal.3.x86_64.rpm
8f4b4f77243e68dfba2feee8c37f6631
    kernel-smp-2.4.21-15.0.2.EL.peterm.eal.3.x86_64.rpm
b9bcb2fe3030052aafad5988d1945802
    kernel-source-2.4.21-15.0.2.EL.peterm.eal.3.x86_64.rpm
59d191b83642cd3642d173f4041b6d66  laus-0.1-62RHEL3.x86_64.rpm
79ac3ab71e1917cc464dc2f3b1b9edd1  laus-devel-0.1-62RHEL3.i386.rpm
5538d44c8fe6da06ce02d0c787c58eec  laus-devel-0.1-62RHEL3.x86_64.rpm
58e5e19bf98d287576ad7d0c50d3ff74  laus-libs-0.1-62RHEL3.i386.rpm
276582e5f9eea630400b32003cbfcfdc  laus-libs-0.1-62RHEL3.x86_64.rpm
932c7d5e07e3a172f0248ab06d405410  passwd-0.68-3.1.x86_64.rpm
8cbdcd09a6945a09ca8f7255f4cbbdb  shadow-utils-4.0.3-20.03.x86_64.rpm
14ce77fd0c0def8147ec109b152dae9  vixie-cron-3.0.1-75.x86_64.rpm

```

When using the automated configuration, copy these RPM files into the directory `/root/rpms/` of the system being installed, the installer will then handle the upgrade automatically. You **MUST** download the current version of the *eal3-certification* RPM package to use the automated configuration as described in the next section.

If installing manually, use the `rpm(8)` command to install and upgrade the downloaded packages:

```
rpm -Uvh *.rpm
```

Opteron/x86.64 only: you **MUST** verify that the program `/sbin/pam.tally` is a 64-bit binary, the installer may have unpacked the 32-bit version instead. Run the command `file /sbin/pam.tally` and check the output. If you have the 32-bit version of `/sbin/pam.tally` installed, you must replace it with the 64-bit version from the *pam* RPM package. Unpack it from the first RHEL Update 2 CD using the following command:

```

cd /
rpm2cpio /mnt/cdrom/RedHat/RPMS/pam-0.75-54.x86_64.rpm \
| cpio -idv '*/sbin/*'

```

3.3 Automated configuration of the system

The *eal3-certification* package **SHOULD** be installed initially to achieve the evaluated configuration. This RPM package contains EAL3 specific configuration files, and the script `/etc/audit/rhel-eal3.bash` that sets up the evaluated configuration.

It is **RECOMMENDED** that you use the `rhel-eal3.bash` script to reset the configuration to its initial state after any updates, but you **MAY** also perform the steps listed here manually.

Install the certification RPM with the following command:

```
rpm -Uvh eal3-certification*.rpm
```

Run the following command to view a summary of the supported options:

```
/etc/audit/rhel-eal3.bash -h
```

It is **RECOMMENDED** that you uninstall all unused kernel packages, such as the uniprocessor kernel on a SMP machine, before running the script. The script will upgrade the installed kernel package(s) to the required version, and if you have multiple packages, the wrong one might be activated due to the upgrade order. You **MAY** also manually upgrade the kernel package (and test it) before running the script.

You will need to specify a directory containing the required update packages (this is `/root/rpms/` by default), and also a directory or media containing the RHEL3 Update2 RPM packages. Specify these with the `--rpm-path` parameter, with the update packages listed first. For example:

```
/etc/audit/rhel-eal3.bash --rpm-path '/root/rpms /mnt/cd*'
```

If the RHEL3 Update2 RPM packages are stored on an NFS file server instead of on CD-R media, specify the path to the RPMS directory as in the following example, using the appropriate path to the mounted directory instead of `/mnt/SERVER/U2/`:

```
/etc/audit/rhel-eal3.bash \
--rpm-path '/root/rpms /mnt/SERVER/U2/RedHat/RPMS/'
```

You **MAY** also add the `--add-optional` flag to automatically install optional packages (useful for testing).

You **MAY** use the `-a` flag to automate the install and have it run without prompting. This is intended for people who are familiar with the process; if running it for the first time you **SHOULD** let it run interactively and verify the actions as described in this guide.

You **MUST** answer all questions asked by the script that are not marked as "optional" with `y` to achieve the evaluated configuration.

WARNING: The `rhel-eal3.bash` script will reboot the system as the final step in the process, as described in the manual instructions in section §3.19 "Reboot and initial network connection" of this guide. On zSeries, it will run the `zipl` boot configuration tool (with no arguments) before rebooting.

If the script has completed successfully, the remaining steps in this chapter were done automatically; you **MAY** skip ahead to section §4 "System operation" of this guide.

3.4 Configuring filesystem parameters

You **MUST** add the mount option `acl` in the file `/etc/fstab` for all ext3 file systems. You **MAY** also add the option `user_xattr`. Multiple options are separated with commas (not "comma space"), for example `acl,user_xattr`.

Edit `/etc/fstab` and replace the `defaults` option specification (fourth column) with `acl` for all file systems with type `ext3` (third column). Then, run `mount MOUNTPOINT -o remount` for each of the mount points (second column).

For more information, please refer to section §4.6 "Mounting filesystems" of this guide.

3.5 Add and remove packages

The minimal system that was initially installed does not contain all packages required for the evaluated configuration, and some of the initially installed packages need to be removed.

In the following lists, the suffix `/cross` indicates a package using the non-default word size. For example, the default "glibc" package on Opteron is named `glibc-*.x86_64.rpm`, while "glibc/cross" refers to `glibc-*.i686.rpm`. Please refer to section §3.2 "Installing required updates" for more information. The following table shows the mappings:

# Architecture	default	/cross
xSeries	i386 or i686	[not applicable]
iSeries, pSeries	ppc	ppc64
eServer 325	x86_64	i386 or i686
zSeries (31bit)	s390	[not applicable]

The evaluated configuration consists of exactly the following packages:

One or both of the following kernel packages:

```
kernel
kernel-smp
```

Packages installed on all architectures:

```
acl amtu ash aspell at attr authconfig autofs basesystem bash
bc beecrypt bind-utils binutils bzip2 bzip2-libs chkconfig
comps coreutils cpio cpp cracklib cracklib-dicts crontabs cups
cups-libs curl cvs cyrus-sasl cyrus-sasl-gssapi cyrus-sasl-md5
cyrus-sasl-plain db4 dev devlabel dhclient dialog diffutils
dos2unix dosfstools dump e2fsprogs ed eject elfutils
elfutils-libelf elinks ethtool expat fbset file filesystem
findutils finger fontconfig freetype ftp gawk gdbm gettext
glib glib2 glibc glibc-common glibc-headers glibc-kernheaders
gmp gnupg gpm grep groff gzip hdparm hesiod hotplug htmlview
hwdata info initscripts iproute ipsec-tools iptables
iptables-ipv6 iputils irda-utils isdn4k-utils jfsutils jwhois
kbd kernel kernel-utils krb5-libs krb5-workstation krbafs
krbafs-utils kudzu laus laus-libs less lftp lha libacl libattr
libcap libgcc libgcj libjpeg libpng libstdc++ libtermcap
libtiff libtool-libs libuser libwvstreams libxml2 lockdev
logrotate logwatch losetup lslk lsof lvm m4 mailcap mailx make
MAKEDEV man man-pages mdadm mgetty mingetty minicom mkinitrd
mktemp modutils mount mtools mtr mt-st nano nc ncompress
ncurses netconfig netdump netpbm netpbm-progs net-tools newt
nfs-utils nscd nss_ldap ntsysv openldap openssh
openssh-clients openssh-server openssl pam pam_krb5
pam_passwdqc pam_smb parted passwd patch pax pciutils pcre
pdksh perl perl-DateManip perl-Filter perl-HTML-Parser
perl-HTML-Tagset perl-libwww-perl perl-URI pinfo popd portmap
postfix ppp prelink procmail procs psacct psmisc pspell
pyOpenSSL python python-optik pyxf86config quota raidtools
rdate rdist readline redhat-config-mouse
redhat-config-network-tui redhat-config-securitylevel-tui
```

```
redhat-logos redhat-lsb redhat-menus redhat-release rhnlib
rhpl rmt rootfiles rpm rpmdb-redhat rpm-python rp-pppoe rsh
rsync schedutils sed setarch setserial setup setuptool
shadow-utils sharutils slang slocate specsps star stunnel sudo
symlinks sysklogd sysreport SysVinit talk tar tcl tcp_wrappers
tcpdump tcsh telnet termcap tetex tetex-fonts tetex-latex tftp
time tk tmpwatch traceroute tzdata unix2dos unzip up2date
usbutils usermode utempter util-linux vconfig vim-common
vim-minimal vixie-cron vsftpd wget which wireless-tools words
wvdial XFree86-libs XFree86-libs-data XFree86-Mesa-libGL
xinetd ypbind yp-tools zip zlib
```

additional packages on x86 (xSeries):

```
apmd eject fbset grub hdparm kbd kernel-pcmcia-cs laus minicom
mkbootdisk pyxf86config redhat-config-mouse rp-pppoe setserial
syslinux usbutils wireless-tools
```

additional packages on x86_64 (eServer 325 (Opteron)):

```
eject fbset grub hdparm kbd laus minicom pyxf86config
redhat-config-mouse rp-pppoe setserial syslinux usbutils
wireless-tools
```

additional packages on ppc (pSeries) and ppc (iSeries):

```
eject fbset hdparm kbd laus-libs/cross laus/cross minicom
ppc64-utils pyxf86config redhat-config-mouse rp-pppoe
setserial usbutils wireless-tools yaboot
```

additional packages on s390 (zSeries):

```
laus s390utils
```

In addition to these packages, certain additional software from the RHEL CDs MAY be installed without invalidating the evaluated configuration. The rules described in section §4.4 "Installation of additional software" of this guide MUST be followed to ensure that the security requirements are not violated.

The following packages are examples of tolerated packages that MAY be added to the system according to these rules. Note that the software contained in these packages is not intended to be used with 'root' privileges, but the presence of the packages does not invalidate the evaluated configuration. The `rhel-eal3.bash` script does not remove these packages if they are installed on the system, and MAY be used to install them automatically by specifying the `--add-optional` parameter to the command line. The example OPTIONAL packages are:

```
autoconf automake bison cracklib/cross db4/cross
eal3-certification-doc expect expect-devel flex gcc gcc-c++
glib/cross glibc-devel glibc-devel/cross glibc/cross
kernel-source/any krb5-devel laus-devel laus-devel/cross
laus-libs/cross libacl/cross libattr-devel libattr/cross
libgcc/cross libstdc++-devel libstdc++-devel/cross
libstdc++/cross netpbm netpbm-progs openldap-clients
openssl-devel pam-devel pam-devel/cross pam/cross
perl-Digest-HMAC perl-Digest-SHA1 rpm-build strace tetex
tetex-fonts tetex-latex texinfo zlib-devel zlib-devel/cross
zlib/cross
```

The next steps involve installing selected packages from the distribution CD-ROMs. Due to dependency issues, the RECOMMENDED method is to first copy all needed RPMs to a temporary directory, and then installing them all in one step using `rpm -Uvh *.rpm`.

The *rhel-eal3.bash* script handles the package selection and installation automatically, and will prompt for the installation media as necessary. After installation, the package selection is again verified, and the script will indicate which packages are still missing or the wrong version. In this case, verify that the needed RPM packages are available in the locations specified, and that they are the correct versions and for the correct architecture.

If you are performing this step manually, first create a temporary directory to store the RPM files:

```
mkdir /root/rpms
```

Copy all the missing package files to that directory. This step is very time consuming when done manually, the RECOMMENDED method is to use the *rhel-eal3.bash* script to do this automatically. The following shows an example of the manual method, this needs to be repeated until all missing packages are copied:

```
# Get list of currently installed packages
rpm -qa | sort | less

# Search for one of the missing packages
find /mnt/cd* -name 'vsftpd*'

# Copy missing packages from the installation media
cp /mnt/cdrom/RedHat/RPMS/vsftpd-1.2.0-4.i386.rpm /root/rpms/

# Repeat these steps for the other missing packages
[...]
```

Once the packages are all copied, install them all in one step with the following single command:

```
# Install all packages
rpm -Uvh /root/rpms/*.rpm
```

This MAY result in the following expected warning and error messages that are harmless:

```
warning: user bhcompile does not exist - using root
warning: group bhcompile does not exist - using root
error: %postun(laus-0.1-54RHEL3) scriptlet failed, exit status 255
```

Any other errors or warnings indicate that the installation is invalid and needs to be redone.

Now you can remove the temporary directory with the following command:

```
rm -rf /root/rpms
```

3.6 Disable services

Note: The system runlevel as specified in the 'initdefault' entry in */etc/inittab* MUST remain at the default setting of '3' for these steps to be valid.

The following services are REQUIRED for runlevel 3:

```
atd           # the 'at' daemon
audit         # the audit daemon
crond         # vixie-cron
irqbalance    # configures SMP IRQ balancing
kudzu         # new device discovery
network       # network interface configuration
random        # random numbers
syslog        # system logging
```

The following services are OPTIONAL for runlevel 3:

```
cups          # print subsystem
gpm           # console mouse management
mdmonitoring  # software raid monitoring
postfix       # SMTP MTA
rawdevices    # Raw partition management (eg. for Oracle)
sshd          # Secure Shell
vsftpd        # FTP server
xinetd        # Internet Services
```

You MUST ensure that all REQUIRED services are active. You MAY enable or disable services from the OPTIONAL list as suitable for your configuration. All other services MUST be deactivated.

Use *chkconfig* *SERVICENAME* *off* to disable a service, and *chkconfig* *SERVICENAME* *on* to enable it. The following command lists the active services:

```
chkconfig --list | grep "3:on" | sort
```

Make sure that the audit subsystem is activated. If *auditd* is not running, all logins are automatically disabled in the evaluated configuration as required by CAPP.

3.7 Remove SUID/SGID root settings from binaries

Use of the SUID bit on binaries (to run with root privileges, a.k.a. "setuid bit") MUST be limited to those shown in the following list:

```
/bin/ping
/bin/su
/usr/bin/at
/usr/bin/chage
/usr/bin/chfn
/usr/bin/chsh
/usr/bin/crontab
/usr/bin/gpasswd
/usr/bin/newgrp
/usr/bin/passwd
```

The other binaries that were installed with the SUID bit set **MUST** have this bit removed. Administrators can still run these binaries normally, but they are not available for ordinary users.

There are also a number of SGID files on the system that are needed:

```
/usr/sbin/postdrop    # group "maildrop"
/usr/sbin/postqueue    # group "maildrop"
/usr/sbin/utempter     # group "tty"
```

Similarly, the SGID bit **MUST NOT** be used to give group "root" privileges to any binary.

Generate a list of all SUID/SGID programs on the system by running the following command:

```
find / -not -fstype ext3 -prune -o \
    -type f \( -perm -4000 -o -perm -2000 \) \
    -print
```

Then, for each file in this list that is not one of the permitted SUID or SGID programs, run the command `chmod -s FILE` to remove the SUID and SGID bits. When done, re-run the *find* command to verify that the list matches the permitted programs.

3.8 Update permissions for 'su'

The 'su' binary **MUST** be restricted to members of the trusted 'wheel' group. This will be enforced both with PAM configuration (configured later) and the binary's permissions.

```
chgrp wheel /bin/su
chmod 4710 /bin/su
```

You **MUST** have at least one user account other than 'root' configured to be a member of the 'wheel' group, otherwise system administration will **ONLY** be possible from the system console.

3.9 Disable root login over the network

Login from the network with user ID 0 ('root') **MUST NOT** be permitted over the network. Administrators **MUST** use an ordinary user ID to log in, and then use the `/bin/su -` command to switch identities. For more information, refer to section §4.3 "Gaining superuser access" of this guide.

It is **RECOMMENDED** that you remind administrators of this by adding the following alias to the bash configuration file `/etc/profile` that disables the pathless 'su' command:

```
alias su="echo \"Always use '/bin/su -' (see Configuration Guide)\""
```

This alias can be disabled for the root user in `/root/.bashrc`:

```
unalias su
```

The restriction for direct root logins is enforced through two separate mechanisms. For network logins using ssh, the `PermitRootLogin no` entry in `/etc/ssh/sshd.config` **MUST** be set (see next section). Console and serial terminal logins use the `pam_securetty.so` PAM module in the `/etc/pam.d/login` file that verifies that the terminal character device used is listed in the file `/etc/securetty`.

The file `/etc/securetty` **MUST NOT** be changed from the secure default settings. The original contents are the following:

```
console
vc/1
vc/2
vc/3
vc/4
vc/5
vc/6
vc/7
vc/8
vc/9
vc/10
vc/11
tty1
tty2
tty3
tty4
tty5
tty6
tty7
tty8
tty9
tty10
tty11
```

3.10 Setting up SSH

SSH protocol version 1 **MUST** be disabled. It has known security deficiencies.

The ssh client **MUST NOT** be set up SUID root (the SUID bit was removed in the post-install configuration). This prevents the use of some authentication methods normally supported by OpenSSH, but does not affect the evaluated configuration that uses password authentication exclusively.

The SSH Server **MUST** be configured to reject attempts to log in as root.

The permitted authentication mechanisms are per-user (nonempty) passwords and per-user DSS public key authentication. All other authentication methods **MUST** be disabled.

The setting `PAMAuthenticationViaKbdInt` **MUST** be disabled, since this would otherwise circumvent the disabled root logins over the network.

This results in the following option set for the SSH daemon that **MUST** be set in `/etc/ssh/sshd.config`:

```
Protocol 2
Ciphers 3des-cbc
SyslogFacility AUTHPRIV
PermitRootLogin no
RSAAuthentication no
PubkeyAuthentication no
IgnoreRhosts yes
RhostsRSAAuthentication no
HostbasedAuthentication no
PasswordAuthentication yes
PermitEmptyPasswords no
ChallengeResponseAuthentication no
KerberosAuthentication no
```

```
PAMAuthenticationViaKbdInt no
X11Forwarding no
Subsystem sftp /usr/libexec/openssh/sftp-server
```

All other options MUST NOT be changed from the defaults or from those settings specified here. Specifically, you MUST NOT add other authentication methods (AFS, Kerberos, host-based) to those permitted here.

3.11 Setting up xinetd

The *xinetd* super server is not used in the evaluated configuration, but MAY be used to start non-root network processes. The file */etc/xinetd.conf* contains default settings, these can be overridden by service-specific entry files stored in the directory */etc/xinetd.d/*.

The log method and the data that is to be logged are defined using the *defaults* entry in the */etc/xinetd.conf* file. The RECOMMENDED settings are:

```
defaults
{
    instances          = 60
    log_type           = SYSLOG authpriv
    log_on_success     = HOST PID EXIT DURATION
    log_on_failure     = HOST ATTEMPT
    cps                = 25 30
}

includedir /etc/xinetd.d
```

The *xinetd.conf*(5) man page contains more information on *xinetd* and configuration examples.

3.12 Setting up FTP

The evaluated configuration OPTIONALLY includes FTP services. Note that FTP does not provide support for encryption, so this is only RECOMMENDED for anonymous access to non-confidential files. If you do not specifically need FTP, it is RECOMMENDED that you disable the *vsftpd*(8) service.

Use the *chkconfig*(8) command to control the FTP service:

```
# Activate FTP service
chkconfig vsftpd on

# Disable FTP service
chkconfig vsftpd off
```

The *vsftpd* service uses several additional configuration files. In */etc/vsftpd/vsftpd.conf* the configuration of the ftp daemon is specified. In addition, the file */etc/vsftpd.ftputers* is used for access control. Users listed in that file can NOT log in via FTP. This file initially contains all system IDs and the root user. It can be augmented with other IDs according to the local needs, but the *root* entry MUST NOT be removed. The *ftputers* file is not checked by the ftp daemon itself but by a PAM module. Please see section §3.16 "Required Pluggable Authentication Module (PAM) configuration" of this guide for details.

The setup of */etc/vsftpd/vsftpd.conf* depends on the local needs. Please refer to *vsftpd.conf*(5) for details.

The default configuration permits only anonymous FTP. This setting is only suitable for distribution of public files for which no read access control is needed. It is RECOMMENDED disabling anonymous FTP if you do not need this functionality with the following */etc/vsftpd/vsftpd.conf* setting:

```
anonymous_enable=NO
```

You MAY enable FTP authentication for local user accounts. The corresponding setting in */etc/vsftpd/vsftpd.conf* is:

```
local_enable=YES
```

It is RECOMMENDED to use the more secure alternatives *sftp*(1) or *scp*(1) to copy files among users, and to use FTP only for legacy applications that do not support this alternative.

3.13 Setting up Postfix

The default settings of the postfix MTA are in accordance with the EAL3 requirements. An alias MUST be set up for root in */etc/aliases*, as postfix will not deliver mail while running with UID 0. Specify one or more user names of administrators to whom mail addressed to *root* will be forwarded.

Please see *postfix*(1), *master*(8), and the documentation in */usr/share/doc/postfix** for details.

3.14 Setting up the audit subsystem

This section describes only the initial setup and default configuration of the audit subsystem. Please refer to section §5.3 "Configuring the audit subsystem" of this guide for information about how it works and what changes MAY be made to the configuration.

Since failures in the audit subsystem may result in an unusable system, ensure that the audit subsystem is properly activated:

```
# Create the audit device if necessary
ls -l /dev/audit
mknod -m 600 /dev/audit c 10 224

# ensure that auditd gets launched
chkconfig audit on
```

3.14.1 Installing the packages needed for auditing

The required packages have already been installed in the previous step described in section §3.2 "Installing required updates" of this guide. This section describes the further changes that need to be made to reach the initial state of the evaluated configuration.

The audit subsystem consists of the following packages:

kernel-*

The kernels include the audit modifications, including the driver *drivers/audit/** and the required hooks in the rest of the kernel.

laus, laus-libs

Contains the userspace Linux Auditing Subsystem (LAuS) programs including *auditd*(8), *aucat*(8) and *augrep*(8), the *liblaus.so* shared library, the */etc/rc.d/init.d/audit* startup script, the configuration in */etc/sysctl.conf*, the */lib/security/pam.laus.so* PAM module and the corresponding man pages. The corresponding development libraries and headers are in the *laus-devel* RPM, which is not installed as part of the evaluated configuration.

at, cron, shadow-utils

These packages contain audit-enabled versions of the trusted programs, which will generate audit records for security relevant events.

3.14.2 Setting up the audit configuration files

It is RECOMMENDED to add the following settings to the file `/etc/sysctl.conf`:

```
dev.audit.max-messages = 1024
dev.audit.paranoia = 0
dev.audit.attach-all = 0
dev.audit.allow-suspend = 1
```

The `rhel-eal3.bash` script automatically sets up this configuration.

The appendix of this guide lists the RECOMMENDED content of the audit configuration files. The *laus* package by default installs these files with the RECOMMENDED contents:

```
/etc/audit/audit.conf
/etc/audit/filter.conf
/etc/audit/filesets.conf
```

3.14.3 Starting auditd at boot as a system service

The evaluated configuration runs `auditd` as a standard daemon service launched as part of the normal startup sequence, this is activated with the following command:

```
chkconfig audit on
```

3.14.4 Starting auditd in fail-secure mode from init (OPTIONAL)

Running `auditd` as a system service is the standard and recommended method, other system components such as `cron` and `atd` are also launched in this way.

However, if `auditd` is killed or unexpectedly terminates, audit messages will be lost until the administrator restarts the service. This failure mode does not violate CAPP requirements, because only the `sysadmin` can kill the audit daemon. The only failure mode addressed by CAPP concerns running out of disk space, and that is handled directly by `auditd`. Any other abnormal termination would indicate a serious bug that should be investigated, reported and fixed.

If you want to ensure that an instance of `auditd` will always be running even in case of these unusual failure modes, you MAY set up an alternative configuration and launch `auditd` via the `init` daemon.

To do this, disable the *audit* system service, then create an entry in the file `/etc/inittab` and activate it:

```
chkconfig audit off
echo "au:35:respawn:/sbin/auditd -F" >> /etc/inittab
init q
```

This operating mode ensures that an instance of `auditd` will always be running, because `init` will automatically restart `auditd` immediately if it terminates for any reason. If `init` cannot restart `auditd` in this way, it will generate a *syslog* warning message and temporarily deactivate the *inittab* entry for five minutes.

3.15 Introduction to Pluggable Authentication Module (PAM) configuration

The PAM subsystem is responsible for maintaining passwords and other authentication data. Because this is a security-critical system, understanding how it works is very important. In addition to the *pam(8)* manual page, full documentation is available in */usr/share/doc/pam-*/txts/* and includes "*The Linux-PAM System Administrator's Guide*" (*pam.txt*) as well as information for writing PAM applications and modules. Detailed information about modules is available in */usr/share/doc/pam-*/txts/README.pam_** as well as manual pages for individual modules, such as *pam_stack(8)*.

The PAM configuration is stored in the */etc/pam.d/* directory. Note that the documentation refers to a file */etc/pam.conf* that is not used by RHEL (PAM was compiled to ignore this file if the */etc/pam.d/* directory exists).

Each service (application) that uses PAM for authentication uses a *service-name* to determine its configuration, stored in the */etc/pam.d/SERVICE_NAME* file. The special *service-name* OTHER (case insensitive) is used for default settings if there are no specific settings.

The configuration file for the service contains one entry for each module, in the format:

```
module-type    control-flag    module-path    args
```

Comments MAY be used extending from '#' to the end of the line, and entries MAY be split over multiple lines using a backslash at the end of a line as a continuation character.

The *module-type* defines the type of action being done. This can be one of four types:

auth

Authenticates users (determines that they are who they claim to be). It can also assign credentials, for example additional group memberships beyond those specified through */etc/passwd* and */etc/groups*. This additional functionality MUST NOT be used.

account

Account management not related to authentication, it can also restrict access based on time of day, available system resources or the location of the user (network address or system console).

session

Manages resources associated with a service by running specified code at the start and end of the session. Typical usage includes logging and accounting, and initialization such as auto mounting a home directory.

password

Used for updating the password (or other authentication token), for example when using the *passwd(1)* utility to change it.

The *control-flag* specifies the action that will be taken based on the success or failure of an individual module. The modules are stacked (executed in sequence), and the *control-flags* determine which final result (success or failure) will be returned, thereby specifying the relative importance of the modules.

Stacked modules are executed in the order specified in the configuration file.

The *control-flag* can be specified as either a single keyword, or alternatively with a more elaborate syntax that allows greater control. RHEL uses only the single keyword syntax by default.

The following keywords control how a module affects the result of the authentication attempt:

required

If this module returns a failure code, the entire stack will return failure. The failure will be reported to the application or user only after all other modules in the stack have been run, to prevent leakage of information (for example, ask for a password even if the entered username is not valid).

requisite

Same as **required**, but return failure immediately not executing the other modules in the stack. Can be used to prevent a user from entering a password over an insecure connection.

sufficient

Return success immediately if no previous **required** modules in the stack have returned failure. Do not execute succeeding modules.

optional

The return code of this module is ignored, except if all other modules in the stack return an indeterminate result (PAM_IGNORE).

The *module-path* specifies the filename of the module to be run (relative to the directory */lib/security/*, and the optional *args* are passed to the module - refer to the module's documentation for supported options.

3.16 Required Pluggable Authentication Module (PAM) configuration

You **MUST** restrict authentication to services that are explicitly specified. The 'other' fallback **MUST** be disabled by specifying the *pam_deny.so* module for each *module-type* in the 'other' configuration. This ensures that access decisions within the PAM system are handled only by the service specific PAM configuration.

Note that RHEL uses the *pam_stack(8)* module to unify commonly used configuration options within single files, rather than having redundant information in multiple files. You **MUST** verify that the shared settings are applicable to services that use *pam_stack*, and keep in mind that a change to the shared file will affect several services.

You **MUST** add the *pam_wheel.so* module to the 'auth' *module-type* configuration for the 'su' service to restrict use of *su(1)* to members of the 'wheel' group.

You **MUST** add the *pam_tally.so* module to the *auth* and *account module-type* configurations of *login*, *sshd* and *vsftpd*. This ensures that accounts are disabled after several failed login attempts. The *pam_tally.so* module is used in the *auth* stack to increment a counter in the file */var/log/lastlog*, and in the *account* stack to either deny login after too many failed attempts, or to reset the counter to zero after successful authentication. The evaluated configuration uses a lockout after six failed attempts, corresponding to the setting *deny=6*, you **MAY** decrease the number for stricter enforcement. Be aware that this can be used in denial-of-service attacks to lock out legitimate users. Please refer to section §4.7 "Managing user accounts" of this guide for more information.

You **MUST** use the *pam_passwdqc.so* password quality checking module to ensure that users will not use easily-guessable passwords.

The system supports many other PAM modules apart from the ones shown here. In general, you **MAY** add PAM modules that add additional restrictions. You **MUST NOT** weaken the restrictions through configuration changes of the modules shown here or via additional modules. Also, you **MUST NOT** add PAM modules that provide additional privileges to users (such as the *pam_console.so* module).

You **MUST NOT** run the *authconfig(8)* tool to modify the authentication configuration.

Following are the pam configuration files:

3.16.1 /etc/pam.d/system-auth

This file contains common settings that are shared by multiple services using authentication. The *pam_passwdqc.so* module is configured to enforce the minimum password length of 8 characters. Note that the *pam_passwdqc.so* module is not part of a default installation, it was added previously as described in section §3.5 "Add and remove packages" of this guide.

The *pam_tally* module **MUST** be used to block the user after 5 failed login attempts.

The *remember* option to *pam_unix.so* prevents users from reusing old passwords. Hashes of old passwords are stored in the file */etc/security/opasswd*. Note that this file **MUST** exist, otherwise users cannot change passwords. Use the following commands to create it:

```
touch /etc/security/opasswd
chmod 600 /etc/security/opasswd
```

The file */etc/pam.d/system-auth* **MUST** be set up with the following content:

```
##PAM-1.0
#
# This file is auto-generated.
# User changes will be destroyed the next time authconfig is run.
#
# ... so don't run authconfig
#
# pam.d/system-auth - PAM master configuration for EAL3/CAPP compliance
#           see the Evaluated Configuration Guide for more info
#

auth        required      pam_tally.so onerr=fail no_magic_root
auth        required      pam_env.so
auth        required      pam_unix.so likeauth nullok

account     required      pam_unix.so
account     required      pam_tally.so deny=5 reset no_magic_root

password    required      pam_passwdqc.so min=disabled,disabled,16,12,8 \
                        random=42
password    required      pam_unix.so nullok use_authtok md5 \
                        shadow remember=7

session     required      pam_limits.so
session     required      pam_unix.so
```

3.16.2 /etc/pam.d/login

This file configures the behavior of the *login* program. It allows root login only for terminals configured in */etc/securetty*. If the file */etc/nologin* is present, then only root can log in.

The *pam_laous.so* module is by default configured to be *optional* instead of *required*, which assumes that all terminals available for login are in physically secure locations and accessible only for authorized administrators. This permits administrators to log in on the console even if the audit subsystem is not available. If any serial terminals are attached and available for arbitrary users, you **MUST** specify the *pam_laous.so* module to be *required* to ensure the CAPP-compliant fail-secure operating mode that disables login if audit is not working. Please refer to section §4.8 "Using serial terminals" of this guide for more information.

```
##PAM-1.0
#
# pam.d/login - PAM login configuration for EAL3/CAPP certification
#           see the Evaluated Configuration Guide for more info
#
```

```
# If serial terminals are in use, pam_laush.so MUST be changed to be 'required'
# for CAPP-complaint fail-secure auditing. The default 'optional' setting
# assumes that all terminals are in physically secure locations.
#

auth      required      pam_securetty.so
auth      required      pam_stack.so service=system-auth
auth      required      pam_nologin.so

account   required      pam_stack.so service=system-auth

password  required      pam_stack.so service=system-auth

session   required      pam_stack.so service=system-auth
#session  required      pam_laush.so # fail-secure mode
session   optional      pam_laush.so # requires physically secure terminals
```

3.16.3 /etc/pam.d/other

This configuration applies for all PAM usage for which no explicit service is configured. It will log and block any attempts.

```
##PAM-1.0
#
# pam.d/other - PAM other configuration for EAL3/CAPP compliance
#               see the Evaluated Configuration Guide for more info
#

auth      required      pam_warn.so
auth      required      pam_deny.so

account   required      pam_warn.so
account   required      pam_deny.so

password  required      pam_warn.so
password  required      pam_deny.so

session   required      pam_warn.so
session   required      pam_deny.so
```

3.16.4 /etc/pam.d/sshd

This file configures the PAM usage for SSH. This is similar to the *login* configuration. The *securetty* entry is not applicable to network logins, and the *pam_laush.so* module MUST be configured to prevent network login if the audit system is not available. Note that *pam_laush.so* MUST run in the *account* stack, it does not work in the *account* or *auth* stacks due to the OpenSSH privilege separation mechanism.

```
##PAM-1.0
#
# pam.d/sshd - pam.d/sshd configuration for EAL3/CAPP compliance
#               see the Evaluated Configuration Guide for more info
#
```

```

auth      required    pam_stack.so service=system-auth
auth      required    pam_nologin.so

account   required    pam_stack.so service=system-auth
account   required    pam_laas.so detach

password  required    pam_stack.so service=system-auth

session   required    pam_stack.so service=system-auth

```

3.16.5 /etc/pam.d/su

This file configures the behavior of the 'su' command. Only users in the trusted 'wheel' group can use it to become 'root', as configured with the *pam_wheel* module.

```

#%PAM-1.0
#
# pam.d/su - PAM su configuration from EAL3/CAPP compliance
#           see the Evaluated Configuration Guide for more info
#

auth      sufficient  pam_rootok.so
auth      required    pam_wheel.so use_uid
auth      required    pam_stack.so service=system-auth

account   required    pam_stack.so service=system-auth

password  required    pam_deny.so

session   required    pam_stack.so service=system-auth
session   optional    pam_xauth.so

```

The *password* branch is disabled because forcing the root user to change the root password is not desired for this program,

3.16.6 /etc/pam.d/vsftpd

This file configures the authentication for the FTP daemon. With the *listfile* module, users listed in */etc/vsftpd.ftpusers* are denied FTP access to the system. Note that the setting is relevant only for authentication of incoming connections, and does not prevent local users from using the *ftp(1)* client to access other machines on the network.

```

#%PAM-1.0
#
# pam.d/vsftpd - vsftpd configuration for EAL3/CAPP compliance
#           see the Evaluated Configuration Guide for more info
#

auth      required    pam_listfile.so item=user sense=deny \
                    file=/etc/vsftpd.ftpusers onerr=succeed
auth      required    pam_stack.so service=system-auth
auth      required    pam_shells.so

```

```

account    required    pam_stack.so service=system-auth
account    required    pam_laas.so detach

password   required    pam_deny.so

session    required    pam_stack.so service=system-auth

```

pam_deny.so is used in the password stack because the FTP protocol has no provisions for changing passwords.

3.17 Configuring default account properties

The file */etc/login.defs* defines settings that will be used by user management tools such as *useradd*(8). The file is not used during the authentication process itself.

The password aging settings defined in this file are used when creating users and when changing passwords, and stored in the user's */etc/shadow* entry. Note that only the */etc/shadow* entries are considered during authentication, so changes in */etc/login.defs* will not retroactively change the settings for existing users.

The `PASS_MIN_LEN` setting has no effect in the evaluated configuration, the relevant settings are instead configured using the `min=` parameter to *pam_passwdqc.so* in the */etc/pam.d/system-auth* file.

```

#   Directory where mailboxes reside, _or_ name of file, relative to the
#   home directory.  If you _do_ define both, MAIL_DIR takes precedence.
#   QMAIL_DIR is for Qmail
#
# The setting is used only when creating or deleting users, and has
# no effect on the mail delivery system. MAY be changed as required.
#
#QMAIL_DIR      Maildir
#MAIL_FILE      .mail
MAIL_DIR        /var/spool/mail

# Password aging controls:
#
#   PASS_MAX_DAYS   Maximum number of days a password may be used.
#   PASS_MIN_DAYS   Minimum number of days allowed between password changes.
#   PASS_MIN_LEN    Minimum acceptable password length.
#   PASS_WARN_AGE   Number of days warning given before a password expires.
#
PASS_MAX_DAYS   60   # MAY be changed, must be <= 60
PASS_MIN_DAYS   1    # MAY be changed, 0 < PASS_MIN_DAYS < PASS_MAX_DAYS
PASS_WARN_AGE   7    # MAY be changed
PASS_MIN_LEN    5    # no effect in the evaluated configuration

#
# Min/max values for automatic uid selection in useradd
#
# MAY be changed, 100 < UID_MIN < UID_MAX < 65535
#
UID_MIN         500
UID_MAX         60000

```

```

#
# Min/max values for automatic gid selection in groupadd
#
# MAY be changed, 100 < GID_MIN < GID_MAX < 65535
#
GID_MIN          500
GID_MAX          60000

#
# If defined, this command is run when removing a user.
# It should remove any at/cron/print jobs etc. owned by
# the user to be removed (passed as the first argument).
#
# MAY be activated as described in the "Managing user accounts"
# section of the ECG.
#
#USERDEL_CMD      /usr/sbin/userdel_local

#
# If useradd should create home directories for users by default
# On RH systems, we do. This option is ORed with the -m flag on
# useradd command line.
#
# MAY be changed.
#
CREATE_HOME      yes

```

3.18 Configuring the boot loader

You MUST set up the server in a secure location where it is protected from unauthorized access. Even though that is sufficient to protect the boot process, it is RECOMMENDED to configure the following additional protection mechanisms:

- Ensure that the installed system boots exclusively from the disk partition containing RHEL, and not from floppy disks, USB drives, CD-ROMs, network adapters, or other devices.
- Ensure that this setting cannot be modified, for example by using a BootProm/BIOS password to protect access to the configuration.

3.18.1 GRUB boot loader configuration

The GRUB boot loader is used on the xSeries and eServer 325 (Opteron) platforms. It is highly configurable, and permits flexible modifications at boot time through a special-purpose command line interface. Please refer to the *grub(8)* man page or run `info grub` for more information.

- Use the `password` command in */boot/grub/menu.lst* to prevent unauthorized use of the boot loader interface. Using md5 encoded passwords is RECOMMENDED, run the command *grub-md5-crypt* to generate the encoded version of a password.
- Protect all menu entries other than the default RHEL boot with the `lock` option, so that the boot loader will prompt for a password when the user attempts to boot from other media (such as a floppy) or sets other non-default options for the boot process. To implement this, add a line containing just the keyword `lock` after the `title` entry in the */boot/grub/menu.lst* file.

- Remove group and world read permissions from the grub configuration file if it contains a password by running the following command:

```
chmod 600 /boot/grub/menu.lst
```

All changes to the configuration take effect automatically on the next boot, there is no need to re-run an activation program.

The following example of the */boot/grub/menu.lst* configuration file shows RECOMMENDED settings:

```
default=0
timeout=10
splashimage=(hd0,0)/boot/grub/splash.xpm.gz
password --md5 $1$04711/$H/JW2MYeugX6Y1h3v.1Iz0
title Red Hat Enterprise Linux AS (2.4.21-15.EL)
    lock
    root (hd0,0)
    kernel /boot/vmlinuz-2.4.21-15.EL ro root=LABEL=/
    initrd /boot/initrd-2.4.21-15.EL.img
```

Note that the configuration shown here might not be exactly the configuration used on the installed system, depending on the kernel options needed for the hardware.

3.18.2 Yaboot boot loader configuration

Yaboot is used on the pSeries machines, it is an OpenFirmware-based boot loader, and can be reconfigured at boot time from a specialized command line.

Yaboot and GRUB are very similar, both support MD5-encrypted passwords specified in the configuration file.

You need to re-run the *ybin(8)* tool when you have modified the configuration file, this is however not necessary if you replace a kernel and keep all path names unchanged.

Please refer to the *yaboot.conf(5)* and *ybin(8)* manual pages and the yaboot HOWTO for more information:

<http://penguinppc.org/projects/yaboot/doc/yaboot-howto.shtml>

3.18.3 ZIPL boot loader configuration

The ZIPL boot loader is used on the zSeries mainframe when the system is set up using the VM virtualization layer. In this context, "booting" refers to the initial program load (IPL) done from the CP command prompt, which affects only a single specific Linux instance (a.k.a. "partition", which refers to the running system and not the disk partition in this context).

Configuration of the VM system is beyond the scope of this document. You MUST ensure that the configuration settings and virtual devices used are only accessible to the authorized administrators. Do NOT use unencrypted 3270 sessions for console access on insecure networks.

ZIPL writes a boot record on the virtual disk (DASD) used by this Linux instance, this boot record then proceeds to load and run the Linux kernel itself. The *zip1* command must be re-run after any kernel or boot argument modifications. Please refer to the *zip1(8)* man page for more information.

The following example shows a typical */etc/zip1.conf* file:

```
[defaultboot]
default=ipl

[ipl]
target=/boot/zipl
image=/boot/kernel/image
ramdisk=/boot/initrd
parameters="dasd=0200 root=/dev/dasda1"
```

3.18.4 iSeries kernel slots

Similar to zSeries, the iSeries hosts use an initial program load (IPL) system to load and initialize a virtual Linux instance. There is no boot loader program on the Linux side, the host platform's boot loader is configured through device drivers accessed via virtual files in the */proc/* file system.

The system supports multiple kernel slots. Usually, slot A contains the production kernel, and slot B is reserved for experimental kernels. The default boot image is selected via the */proc/iSeries/mf/side* virtual file.

The kernel slot may contain either just a plain kernel (file name "vmlinux" or similar), or a combined kernel plus initial root disk (file name "vmlinitrd" or similar). Use the combined kernel+initrd if available to ensure that all necessary modules will be available for booting.

There are usually several different kernels and/or kernel+initrd files in */boot/*, be careful to use the right file based on the version number information contained in the file name.

Here is a sample session to copy an image to kernel slot B, and activate it:

```
dd if=/boot/vmlinitrd of=/proc/iSeries/mf/B/vmlinux bs=4k
cat /proc/cmdline > /proc/iSeries/mf/B/cmdline
echo "B" > /proc/iSeries/mf/side
```

For more information, please refer to:

http://www-1.ibm.com/servers/eserver/iseries/linux/tech_faq.html

3.19 Reboot and initial network connection

After all the changes described in this chapter have been done, you **MUST** reboot the system to ensure that all unwanted tasks are stopped, and that the running kernel, modules and applications all correspond to the evaluated configuration.

Please make sure that the boot loader is configured correctly for your platform. On zSeries, remember to run the *zipl(8)* tool to write the boot record.

The system will then match the evaluated configuration. The server **MAY** then be connected to a secure network as described above.

4 System operation

To ensure that the systems remains in a secure state, special care **MUST** be taken during system operation.

4.1 System startup, shutdown and crash recovery

Use the *shutdown*(8), *halt*(8) or *reboot*(8) programs as needed to shut down or reboot the system.

When powered on (or on initial program load of the logical partition on a host system), the system will boot into the RHEL operating system. If necessary (for example after a crash), a filesystem check will be performed automatically. In rare cases manual intervention is necessary, please refer to the *e2fsck*(8) and *debugfs*(8) documentation for details in this case.

In case a nonstandard boot process is needed (such as booting from floppy disk or CD-ROM to replace a defective hard drive), interaction with the boot loader and/or the host's management system can be used to modify the boot procedure for recovery.

For example, on xSeries you can use the following grub commands to launch a shell directly from the kernel, bypassing the normal init/login mechanism:

```
# view the current grub configuration
grub> cat (hd0,1)/boot/grub/menu.lst

# manually enter the modified settings
grub> kernel (hd0,1)/boot/vmlinuz root=/dev/sda1 init=/bin/sh
grub> initrd (hd0,1)/boot/initrd
grub> boot
```

Please refer to the relevant documentation of the boot loader, as well as the RHEL administrator guide, for more information.

4.2 Backup and restore

Whenever you make changes to security-critical files, you MAY need to be able to track the changes made and revert to previous versions, but this is not required for compliance with the evaluated configuration.

The *star*(1) archiver is RECOMMENDED for backups of complete directory contents, please refer to section §6.5 "Data import / export" of this guide. Regular backups of the following files and directories (on removable media such as tapes or CD-R, or on a separate host) are RECOMMENDED:

```
/etc/
/var/spool/cron/
/var/spool/at/
```

Depending on your site's audit requirements, also include the contents of */var/log/* in the backup plan. In that case, the automatic daily log file rotation needs to be disabled or synchronized with the backup mechanism, refer to sections §5.2 "System logging and accounting" and §5.3 "Configuring the audit subsystem" of this guide for more information.

You MUST protect the backup media from unauthorized access, because the copied data does not have the access control mechanisms of the original file system. Among other critical data, it contains the secret keys used by the *SSH* and *stunnel* servers, as well as the */etc/shadow* password database. Store the backup media at least as securely as the server itself.

A RECOMMENDED method to track changes is to use a version control system. RCS is easy to set up because it does not require setting up a central repository for the changes, and you can use shell scripting to automate the change tracking. RCS is not included in the evaluated configuration, see *rcsintro*(1) in the *rcs* RPM package for more information. Alternatively, you can manually create backup copies of the files and/or copy them to other servers using *scp*(1).

4.3 Gaining superuser access

System administration tasks require superuser privileges. Since directly logging on over the network as user 'root' is disabled, you **MUST** first authenticate using an unprivileged user ID, and then use the `su` command to switch identities. Note that you **MUST NOT** use the 'root' rights for anything other than those administrative tasks that require these privileges, all other tasks **MUST** be done using your normal (non-root) user ID.

You **MUST** use exactly the following `su(1)` command line to gain superuser access:

```
/bin/su -
```

This ensures that the correct binary is executed irrespective of `PATH` settings or shell aliases, and that the root shell starts with a clean environment not contaminated with the starting user's settings. This is necessary because the `.profile` shell configuration and other similar files are writable for the unprivileged ID, which would allow an attacker to easily elevate privileges to root if able to subvert these settings.

Administrators **MUST NOT** add any directory to the root user's `PATH` that are writable for anyone other than 'root', and similarly **MUST NOT** use or execute any scripts, binaries or configuration files that are writable for anyone other than 'root', or where any containing directory is writable for a user other than 'root'.

4.4 Installation of additional software

Additional software packages **MAY** be installed as needed, provided that they do not conflict with the security requirements.

Any additional software added is not intended to be used with superuser privileges. The administrator **MUST** use only those programs that are part of the original evaluated configuration for administration tasks, except if the administrator has independently ensured that use of the additional software is not a security risk.

Administrators **MAY** add scripts to automate tasks as long as those only depend on and run programs that are part of the evaluated configuration.

The security requirements for additional software are:

- Kernel modules other than those provided as part of the evaluated configuration **MUST NOT** be installed or loaded. You **MUST NOT** load the `tux` kernel module (the in-kernel web server is not supported). You **MUST NOT** add support for non-ELF binary formats or foreign binary format emulation that circumvents system call auditing. You **MUST NOT** activate `knfsd` or export NFS file systems.
- Device special nodes **MUST NOT** be added to the system.
- SUID root or SGID root programs **MUST NOT** be added to the system. Programs which use the SUID or SGID bits to run with identities other than 'root' **MAY** be added.
- The content, permissions, and ownership of all existing filesystem objects (including directories and device nodes) that are part of the evaluated configuration **MUST NOT** be modified. Files and directories **MAY** be added to existing directories provided that this does not violate any other requirement.
- Programs automatically launched with 'root' privileges **MUST NOT** be added to the system. Exception: processes that *immediately* and *permanently* switch to a non privileged identity on launch are permitted, for example by using `su USERID -c LAUNCH.COMMAND` in the startup file, or alternatively by using the `setgroups(2)`, `setgid(2)` and `setuid(2)` system calls in a binary. (`seteuid(2)` etc. are insufficient.)

Automatic launch mechanisms are:

- Entries in `/etc/inittab`
- Executable files or links in `/etc/rc.d/init.d/` and its subdirectories

- Entries in */etc/xinetd.conf*
- Scheduled jobs using *cron* (including entries in */etc/cron** files) or *at*

Examples of programs that usually do not conflict with these requirements and MAY be installed are compilers, interpreters, network services running with non-root rights, and similar programs. The requirements listed above MUST be verified in each specific case.

4.5 Scheduling processes using *cron* and *at*

The *cron*(8) program schedules programs for execution at regular intervals. Entries can be modified using the *crontab*(1) program - the file format is documented in the *crontab*(5) manual page.

You MUST follow the rules specified for installation of additional programs for all entries that will be executed by the 'root' user. Use non-root *crontab* entries in all cases where 'root' privileges are not absolutely necessary.

The *at*(1) and *batch*(1) programs execute a command line at a specific single point of time. The same rules apply as for jobs scheduled via *cron*(8). Use *atq*(1) and *atrm*(1) to manage the scheduled jobs.

Errors in the non interactive jobs executed by *cron* and *at* are reported in the system log files in */var/log/*, and additionally via e-mail to the user who scheduled it.

Permission for users to schedule jobs with *cron* and *at* is controlled through the following *allow* and *deny* files:

```
/etc/at.allow
/etc/at.deny
/etc/cron.allow
/etc/cron.deny
```

The *allow* file has precedence if it exists, then only those users whose usernames are listed in it are permitted to use the service. If it does not exist, the *deny* file is used instead and all users who are *not* listed in that file can use the service. Note that the contents of these files are only relevant when the scheduling commands are executed, and changes have no effect on already scheduled commands.

In the RHEL distribution, the *allow* files do not exist, and *deny* files are used to prevent system-internal IDs and/or guest users from using these services. By default, the evaluated configuration permits everybody to use *cron* and *at*.

It is RECOMMENDED to restrict the use of *cron* and *at* to human users and disallow system accounts from using these mechanisms. For example, the following commands add all system accounts other than root to the *deny* files:

```
awk -F: '{if ($3>0 && $3<100) print $1}' /etc/passwd >/etc/at.deny
chmod 600 /etc/at.deny
cp /etc/at.deny /etc/cron.deny
```

Administrators MAY schedule jobs that will be run with the privileges of a specified user by editing the file */etc/crontab* with an appropriate username in the sixth field. Entries in */etc/crontab* are not restricted by the contents of the *allow* and *deny* files.

You MAY create */etc/at.allow* and/or */etc/cron.allow* files to explicitly list users who are permitted to use these services. If you do create these files, they MUST be owned by the user 'root' and have file permissions 0600 (no access for group or others).

4.6 Mounting filesystems

If any filesystems need to be mounted in addition to those set up at installation time, appropriate mount options **MUST** be used to ensure that mounting the filesystem does not introduce capabilities that could violate the security policy.

A new file system can be integrated as part of the evaluated configuration, for example by installing an additional hard disk, under the following conditions:

- The device is protected against theft or manipulation in the same way as the server itself, for example by being installed inside the server.
- One or more new, empty, file systems in EXT3 format are created on it.
- The file systems are mounted using the `acl` option, for example with the following setting in the `/etc/fstab` file:

```
/dev/sdc1 /home2 ext3 acl 1 2
```

Existing files and directories **MAY** then be moved onto the new file systems.

- If a device containing a file system is ever removed from the system, the device **MUST** be stored within the secure server facility, or alternatively **MUST** be destroyed in a way that the data on it is reliably erased.

Alternatively, media **MAY** be accessed without integrating them into the evaluated configuration, for example CD-ROMs.

The following mount options **MUST** be used if the filesystems contain data that is not part of the evaluated configuration:

```
ro,nodev,nosuid
```

Adding the `noexec` mount option to avoid accidental execution of files or scripts on additional mounted filesystems is **RECOMMENDED**.

Note that these settings do not completely protect against malicious code and data, you **MUST** also verify that the data originates from a trustworthy source and does not compromise the server's security. Specifically, be aware of the following issues:

- Even unprivileged programs and scripts can contain malicious code that uses the calling user's rights in unintended ways, such as corrupting the user's data, introducing trojan horses in the system, attacking other machines on the network, revealing confidential documents, or sending unsolicited commercial e-mail ("spam").
- Data on the additional filesystem **MUST** have appropriate access rights to prevent disclosure to or modification by unauthorized users. Be aware that imported data may have been created using user names and permissions that do not match your system's security policies.
- You **MUST NOT** write data on removable file systems such as floppy disks, since it cannot be adequately protected by the system's access control mechanisms after being removed from the system. Please refer to section §4.2 "Backup and restore" of this guide for more information regarding non-filesystem-based backup.

Each new file system **MUST** be mounted on an empty directory that is not used for any other purpose. It is **RECOMMENDED** using subdirectories of `/mnt` for temporary disk and removeable storage media mounts.

For example:

```
# mount /dev/cdrom /mnt/cdrom -t iso9660 -o ro,nodev,nosuid,noexec
```

You MAY also add an equivalent configuration to */etc/fstab*, for example:

```
/dev/cdrom /mnt/cdrom iso9660 ro,noauto,nodev,nosuid,noexec 0 0
```

You MUST NOT include the *user* flag, ordinary users are not permitted to mount filesystems. This is also enforced by the deletion of the SUID bit on the *mount* command.

4.7 Managing user accounts

Use the *useradd*(8) command to create new user accounts, then use the *passwd*(1) command to assign an initial password for the user. Alternatively, if the user is present when the account is created, permit them to choose their own password. Refer to the manual pages for *useradd*(8) and *passwd*(1) for more information.

If you assign an initial password for a new user, you MUST transfer this initial password in a secure way to the user, ensuring that no third party gets the information. For example, you can tell the password to a user personally known to you. If this is not possible, you MAY send the password in written form in a sealed letter. This applies also when you set a new password for a user in case the user has forgotten the password or it has expired. You need to advise the user that he MUST change this initial password when he first logs into the system and select his own password in accordance with the rules defined in section §6.3 "Password policy" of this guide.

You MUST NOT use the *-p* option to *useradd*(8), specifying a password in that way would bypass the password quality checking mechanism.

The temporary password set by the administrator MUST be changed by the user as soon as possible. Use the *chage*(8) command with the *-d* option to set the last password change date to a value where the user will be reminded to change the password. The RECOMMENDED value is based on the settings in */etc/login.defs* and is equivalent to today's date plus *PASS_WARN_AGE* minus *PASS_MAX_DAYS*.

Example:

```
useradd -m -c "John Doe" jdoe
passwd jdoe
chage -d $(date +%F -d "53 days ago") jdoe
```

The *-m* option to *useradd*(8) creates a home directory for the user based on a copy of the contents of the */etc/skel/* directory. Note that you MAY modify some default configuration settings for users, such as the default *umask*(2) setting or time zone, by editing the corresponding global configuration files:

```
/etc/profile
/etc/bashrc
/etc/csh.cshrc
```

If necessary, you MAY reset the user's password to a known value using *passwd USER*, and entering the new password. You cannot recover the previously used password, since the hash function used is not reversible.

You MAY use the *usermod*(8) command to change a user's properties. For example, if you want to add the user 'jdoe' to the *wheel* group, you could use the following:

```
# List the groups the user is currently a member of:
groups jdoe

# Add the additional group
usermod -G $(su jdoe -c groups | sed 's/ /,/g'),wheel jdoe
```

Users MAY be locked out (disabled) using `passwd -l USER`, and re-enabled using `passwd -u USER`.

The `pam_tally.so` PAM module enforces automatic lockout after excessive failed authentication attempts, as described in section §3.16 "Required Pluggable Authentication Module (PAM) configuration" of this guide. Use the program `pam_tally` to view and reset the counter if necessary, as documented in the file `/usr/share/doc/pam-*/txts/README.pam_tally`. Note that the `pam_tally` mechanism does not *prevent* password guessing attacks, it only prevents *use* of the account after such an attack has been detected. Therefore, you **MUST** assign a new password for the user before reactivating an account. For example:

```
# view the current counter value
pam_tally --user jdoe

# set new password, and reset the counter
passwd jdoe
pam_tally --user jdoe --reset
```

The `chage(1)` utility MAY be used to view and modify the expiry settings for user accounts. Unprivileged users are able to view but not modify their own expiry settings.

The `userdel(8)` utility removes the user account from the system, but does not remove files outside the home directory (and the mail spool file), or kill processes belonging to this user. Use `kill` (or reboot the system) and `find` to do so manually if necessary, for example:

```
# Which user to delete?
U=jdoe

# Lock user account, but don't remove it yet
passwd -l $U

# Kill all user processes, repeat if needed (or reboot)
kill -9 `ps -la --User $U|awk '{print $4}'`

# Recursively remove all files and directories belonging to user
# (Careful - this may delete files belonging to others if they
# are stored in a directory owned by this user.)
find / -depth \( ! -fstype ext3 -prune -false \) \
    -o -user $U -exec rm -rf {} \;

# Remove cron and at jobs
crontab -u $U -r
find /var/spool/at -user $U -exec rm {} \;

# Now delete the account
userdel $U
```

If you need to create additional groups or modify existing groups, use the `groupadd(8)`, `groupmod(8)` and `groupdel(8)` commands.

Group passwords are NOT supported in the evaluated configuration, and have been disabled by removing the SUID bit from the `newgrp(8)` program. You **MUST NOT** re-enable this feature and **MUST NOT** use `passwd(1)` with the `-g` switch or the `gpasswd(1)` command to set group passwords.

4.8 Using serial terminals

You MAY attach serial terminals to the system. They are activated by adding an entry in the file */etc/inittab* for each serial terminal that causes *init*(8) to launch an *agetty*(8) process to monitor the serial line. *agetty* runs *login*(1) to handle user authentication and set up the user's session.

If you use serial terminals and require the CAPP-compliant fail-safe audit mode, you MUST ensure that the file */etc/pam.d/login* is configured to require the *pam.laus.so* module in the session stack. Please refer to section §3.16.2 “*/etc/pam.d/login*” of this guide for more information about the needed PAM configuration.

For example, adding the following line to */etc/inittab* activates a VT102-compatible serial terminal on serial port */dev/ttyS1*, communicating at 19200 bits/s:

```
S1:3:respawn:/sbin/agetty 19200 ttyS1 vt102
```

The first field MUST be an unique identifier for the entry (typically the last characters of the device name). Please refer to the *agetty*(8) and *inittab*(5) man pages for further information about the format of entries.

You MUST reinitialize the *init* daemon after any changes to */etc/inittab* by running the following command:

```
init q
```

4.9 SYSV shared memory and IPC objects

The system supports SYSV-compatible shared memory, IPC objects, and message queues. If programs fail to release resources they have used (for example, due to a crash), the administrator MAY use the *ipcs*(8) utility to list information about them, and *ipcrm*(8) to force deletion of unneeded objects. Note that these resources are also released when the system is rebooted.

For additional information, please refer to the *msgctl*(2), *msgget*(2), *msgrcv*(2), *msgsnd*(2), *semctl*(2), *semget*(2), *semop*(2), *shmat*(2), *shmctl*(2), *shmdt*(2), *shmget*(2) and *ftok*(3) manual pages.

4.10 Configuring secure network connections with *stunnel*


4.10.1 Introduction

The *stunnel* program is a flexible and secure solution for setting up encrypted network connections, enabling the use of strong encryption even for applications that are not able to use encryption natively. *stunnel* uses the OpenSSL library for its encryption functions, and the corresponding *openssl*(1) command line tool for key management.

Stunnel has three main operating modes:

- Accept incoming SSL-encrypted TCP connections, and run a specific program to handle the request.
This is similar to how *xinetd* launches programs, and any program compatible with *xinetd* can also be used for this purpose. It must read and write the communication data on the *stdin* and *stdout* file descriptors and stay in the foreground. *stunnel* also supports switching user and group IDs before launching the program.
- Open a SSL connection to a remote SSL-capable TCP server, and copy data to and from *stdin* and *stdout*.
- Bind a TCP port to accept incoming unencrypted connections, and forward data using SSL to a prespecified remote server.

The following diagram shows a sample usage scenario:



In this scenario, neither the client nor the server have administrator privileges, they are running as normal user processes. Also, the client and server do not support encryption directly.

stunnel makes a secure communication channel available for the client and server. On the client, *stunnel* is accepting connections on TCP port 82. The client connects to this port on the local machine using normal unencrypted TCP, *stunnel* accepts the connection, and opens a new TCP connection to the *stunnel* server running on the remote machine. The *stunnel* instances use cryptographic certificates to ensure that the data stream has not been intercepted or tampered with, and then the remote *stunnel* opens a third TCP connection to the server, which is again a local unencrypted connection.

Any data sent by either the client or server is accepted by the corresponding *stunnel* instance, encrypted, sent to the other *stunnel*, decrypted and finally forwarded to the receiving program. This way, no modifications are required to the client and server.

To set up a secure connection compliant with the evaluated configuration, you **MUST** start the *stunnel* server(s) with administrator rights, and you **MUST** use a TCP port in the administrator-reserved range 1-1023 to accept incoming connections. A corresponding client which connects to the server **MAY** be started by any user, not just administrators.

stunnel **MAY** also be used by non-administrative users to receive encrypted connections on ports in the range 1024-65536. This is permitted, but it is outside of the scope of the evaluated configuration and not considered to be a trusted connection.

Any network servers and clients other than the trusted programs described in this guide (*stunnel*, *sshd*, *vsftpd*, *postfix* and *cupsd*) **MUST** be run using non-administrator normal user identities. Programs run from *stunnel* **MUST** be switched to a non-root user ID by using the *setuid* and *setgid* parameters in the */etc/stunnel/*.conf* configuration files.

It is **RECOMMENDED** configuring any such servers to accept connections only from machine-local clients, either by binding only the *localhost* IP address 127.0.0.1, or by software filtering inside the application. This ensures that the only encrypted connections are possible over the network. Details on how to do this depend on the software being used and are beyond the scope of this guide.

Please refer to the *stunnel*(8) and *openssl*(1) man pages for more information.

4.10.2 Creating an externally signed certificate

It is strongly **RECOMMENDED** that you have your server's certificate signed by an established Certificate Authority (CA), which acts as a trusted third party to vouch for the certificate's authenticity for clients. Please refer to the

openssl(1) and *req*(1) man pages for instructions on how to generate and use a certificate signing request.

Create the server's private key and a certificate signing request (CSR) with the following commands:

```
touch /etc/stunnel/stunnel.pem

chmod 400 /etc/stunnel/stunnel.pem

openssl req -newkey rsa:1024 -nodes \
    -keyout /etc/stunnel/stunnel.pem -out /etc/stunnel/stunnel.csr
```

You will be prompted for the information that will be contained in the certificate. Most important is the "Common Name", because the connecting clients will check if the hostname in the certificate matches the server they were trying to connect to. If they do not match, the connection will be refused, to prevent a 'man-in-the-middle' attack.

Here is a sample interaction:

```
Generating a 1024 bit RSA private key
.....++++++
.....++++++
writing new private key to '/etc/stunnel/stunnel.pem'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [PL]:US
State or Province Name (full name) [Some-State]:TX
Locality Name (eg, city) []:Austin
Organization Name (eg, company) [Stunnel Developers Ltd]:Example Inc.
Organizational Unit Name (eg, section) []:
Common Name (FQDN of your server) []:www.example.com
Common Name (default) []:localhost
```

The file */etc/stunnel/stunnel.pem* will contain both the certificate (public key) and also the secret key needed by the server. The secret key will be used by non-interactive server processes, and cannot be protected with a passphrase. You **MUST** protect the secret key from being read by unauthorized users, to ensure that you are protected against someone impersonating your server.

Next, send the generated CSR file */etc/stunnel/stunnel.csr* (*not* the private key) to the CA along with whatever authenticating information they require to verify your identity and your server's identity. The CA will then generate a signed certificate from the CSR, using a process analogous to `openssl req -x509 -in stunnel.csr -key CA-key.pem -out signed-cert.pem`.

When you receive the signed certificate back from the CA, append it to the file */etc/stunnel/stunnel.pem* containing the private key using the following command:

```
echo >> /etc/stunnel/stunnel.pem
cat signed-cert.pem >> /etc/stunnel/stunnel.pem
```

Make sure that the resulting file contains no extra whitespace or other text in addition to the key and certificate, with one blank line separating the private key and certificate:

```

-----BEGIN RSA PRIVATE KEY-----
MIICXQIBAAKBgQCzF3ezbZFLjgv1YHNXnBnI8jmeQ5MmkvdNw9XkLnA2ONKQmvPQ
[ ... ]
4tjzwTFxPKYvAW3DnXxRAkAvaf1mbc+GTMoAiepXPVfqSpW2Qy5r/wa04d9phD5T
oUNbDU+ezu0Pana7mmmvG3Mi+BuqwlQ/iU+G/qrG6VGj
-----END RSA PRIVATE KEY-----

-----BEGIN CERTIFICATE-----
MIIC1jCCAj+gAwIBAgIBADANBgkqhkiG9w0BAQQFADBXMQswCQYDVQQGEwJQTDET
[ ... ]
bIbYKL6Q1kE/vhGmRXcXQrZzkfu8sgJv7JsDpoTpAdUnmvssUY0bchqFo4Hhzkvs
U/whL2/8RFv5jw==
-----END CERTIFICATE-----

```

You MAY distribute the original signed certificate (*signed-cert.pem* in this example) to clients, it does not contain any confidential information. *Never* distribute the file containing the private key, that is for use by the `stunnel` server only.

When using externally signed certificates, you **MUST** use the option `CAspath` in `stunnel` client configuration files along with the setting `verify=2` or `verify=3` to enable the clients to verify the certificate.

4.10.3 Creating a self-signed certificate

Alternatively, you MAY use a self-signed certificate instead of one signed by an external CA. This saves some time and effort when first setting up the server, but each connecting client will need to manually verify the certificate's validity. Experience shows that most users will not do the required checking and simply click "OK" for whatever warning dialogs that are shown, resulting in significantly reduced security. Self-signed certificates can be appropriate for controlled environments with a small number of users, but are not recommended for general production use.

Create a self-signed host certificate with the following commands:

```

cd /usr/share/ssl/certs
make stunnel.pem
mv stunnel.pem /etc/stunnel/
chmod 400 /etc/stunnel/stunnel.pem

```

The secret key contained in this file **MUST** be kept secret.

You MAY extract the public certificate from this file for distribution to clients. Make sure you do not accidentally distribute the secret key:

```

cd /etc/stunnel
sed '1,/END/d' < stunnel.pem > signed-cert.pem

```

The client has no independent way to verify the validity of a self-signed certificate, each client **MUST** manually verify and confirm the validity of the certificate.

One method is to give a copy of the self-signed certificate to the client (using a secure transport mechanism, not e-mail), and import it into the client directly. The `stunnel` client uses the `CAsfile` option for this purpose.

Alternatively, many client programs (not `stunnel`) can interactively import the certificate when connecting to the server. The client will display information about the server's certificate including an MD5 key fingerprint. You need to compare this fingerprint with the original fingerprint of the server's certificate.

Run the following command on the server to display the original certificate's fingerprint:

```
openssl x509 -fingerprint -in /etc/stunnel/stunnel.pem
```

Most clients will store the certificate for future reference, and will not need to do this verification step on further invocations.

4.10.4 Activating the tunnel

In the evaluated configuration, you **MUST** use the RC4-SHA cipher suite as defined in the SSL v3 protocol, also known as SSL_RSA_WITH_RC4_128_SHA (SC1.8).

You **MUST** specify the cipher list in all *stunnel* client and server configuration files:

```
ciphers = RC4-SHA
```

For a service or tunnel that will only be used temporarily, simply launch the *stunnel* program from the command line and specify an appropriate configuration file. The tunnel will be available for multiple clients, but will not be started automatically after a reboot. To shut down the tunnel, search for the command line in the `ps ax` process listing, and use the `kill(1)` command with the PID shown for the *stunnel* process.

The **RECOMMENDED** method is to use two separate configuration files, one for server definitions (incoming connections use SSL), and one for client definitions (outgoing connections use SSL). More complex configurations will require additional configuration files containing individual service-specific settings. You **MUST** use the **REQUIRED** settings in all *stunnel* configuration files.

Use the following content for the file `/etc/stunnel/stunnel-server.conf`:

```
### /etc/stunnel/stunnel-server.conf
#
# The following settings are REQUIRED for CAPP compliance when used
# as a server, see ECG. File names MAY be changed as needed.
cert = /etc/stunnel/stunnel.pem
ciphers = RC4-SHA
#
# User and group ID MUST NOT be "root", but MAY be changed as needed.
setuid = nobody
setgid = nobody
#
# The following settings are RECOMMENDED
debug = 6
output = /var/log/stunnel-server.log
pid =
foreground = yes
#
# Individual service definitions follow
```

Use the following content for the file `/etc/stunnel/stunnel-client.conf`:

```
### /etc/stunnel/stunnel-client.conf
#
# The following settings are REQUIRED for CAPP compliance when used
# as a client, see ECG. File names MAY be changed as needed. You
# MAY use CApth instead of CAfile for externally signed certificates.
CAfile = /etc/stunnel/signed-cert.pem
```

```

ciphers = RC4-SHA
client = yes
verify = 2
#
# User and group ID MUST NOT be "root", but MAY be changed as needed.
setuid = nobody
setgid = nobody
#
# The following settings are RECOMMENDED
debug = 6
output = /var/log/stunnel-client.log
pid =
foreground = yes
#
# Individual service definitions follow

```

The RECOMMENDED launch method for *stunnel*(8) is via the *init*(8) process. This requires adding new entries to */etc/inittab*, the tunnels will be re-launched automatically whenever they are terminated, as well as after a reboot. The following are the RECOMMENDED */etc/inittab* entries:

```

ts:3:respawn:/usr/sbin/stunnel /etc/stunnel/stunnel-server.conf
tc:3:respawn:/usr/sbin/stunnel /etc/stunnel/stunnel-client.conf

```

Make sure you use the option `foreground = yes` in the configuration file when running from *init* (otherwise *init* will misinterpret the backgrounded server as having died and will try to restart it immediately, causing a loop), and use the `output` option to redirect the output to a log file.

4.10.5 Using the tunnel

If the client program supports SSL encryption, it will be able to communicate with the *stunnel* service directly. You will need to verify and accept the server's certificate if the client cannot recognize it as valid according to its known certification authorities.

If the client program does not support SSL directly, you can use *stunnel* as a client, or indirectly by setting up a proxy that allows the client to connect to an unencrypted local TCP port.

WARNING: The *stunnel* client does *not* verify the server's certificate by default. You **MUST** specify either `verify = 2` or `verify = 3` on the client command line to switch on certificate verification.

You **MAY** also activate client certificate verification in the server's configuration file, so that the server can verify the client's identity as well.

As described in the previous section, you **MUST** use the option `ciphers = RC4-SHA` in the configuration file to ensure that the cipher selection supported in the evaluated configuration will be used.

4.10.6 Example 1: Secure SMTP delivery

Normal SMTP e-mail delivery is not encrypted, but most mail clients support the enhanced SMTPS protocol that uses SSL encryption. The protocol itself is unchanged other than being encrypted.

stunnel can easily be used as a proxy to receive SMTPS connections on the standard port expected by clients (465/tcp), and then forward the data to the mail server listening on the SMTP port (25/tcp). The mail server configuration does not need to be modified to support encryption of incoming mail.

To implement SSL support for incoming mail, add the following service definition to the */etc/stunnel/stunnel-server.conf* configuration:

```
[inbound_mail]
accept = 465
connect = 127.0.0.1:25
```

4.10.7 Example 2: Simple web server

The following shell script acts as a simple web server, reading requests from standard input and writing HTTP/HTML to standard output:

```
cat > /usr/local/sbin/webserver_test <<__EOF__
#!/bin/sh
# Simple web server, can be run via stunnel or xinetd
#
# read and discard client data
dd bs=65536 count=1 >/dev/null 2>&1
#
# Send HTTP header
echo -e "HTTP/1.0 200\r"
echo -e "Content-type: text/html\r"
echo -e "\r"
#
# Send HTML output
echo "<html>"
echo "<h1>Test Page</h1>"
date
echo "<h2>Memory usage</h2>"
echo "<pre>"
free
echo "</pre>"
echo "</html>"
__EOF__

chmod +x /usr/local/sbin/webserver_test
```

Add the following entry to the `/etc/stunnel/stunnel-server.conf` configuration to make this service available using the encrypted HTTPS protocol:

```
[webserver_test]
accept = 443
exec = /usr/local/sbin/webserver_test
TIMEOUTclose = 0
```

Then, use a SSL-capable web browser to connect to port 443:

```
elinks https://localhost/
```

4.10.8 Example 1: system status view

This example shows how to combine *stunnel* client and server definitions to implement an encrypted tunnel for applications that do not themselves support encryption.

First, on the server machine, set up a *stunnel* server definition that accepts SSL connections on TCP port 444, and reports memory usage statistics for the server to connecting clients. Add the following service definition to the `/etc/stunnel/stunnel-server.conf` configuration:

```
[free]
accept = 444
exec = /usr/bin/free
execargs = free
```

Then, on the client machine, add the following entry to the `/etc/stunnel/stunnel-client.conf` configuration, using the server's IP address instead of "127.0.0.1":

```
[free]
accept  = 81
connect = 127.0.0.1:444
```

On the client machine, connect to the local *stunnel* proxy by running the following command as a normal user:

```
telnet localhost 81
```

This will open an unencrypted TCP connection to the client's local port 81, then *stunnel* builds an encrypted tunnel to the server's port 444 and transfers the decrypted data (in this case, the "free" output) back to the client. All unencrypted connections are machine local, and the data transferred over the network is encrypted.

4.11 The Abstract Machine Testing Utility (AMTU)

The security of the operating system depends on correctly functioning hardware. For example, the memory subsystem uses hardware support to ensure that the memory spaces used by different processes are protected from each other.

The Abstract Machine Testing Utility (AMTU) is distributed as an RPM, and was installed previously as described in section §3.5 "Add and remove packages" of this guide.

To run all supported tests, simply execute the `amtu` program:

```
amtu
```

A successful run is indicated by the following output:

```
Executing Memory Test...
Memory Test SUCCESS!
Executing Memory Separation Test...
Memory Separation Test SUCCESS!
Executing Network I/O Tests...
Network I/O Controller Test SUCCESS!
Executing I/O Controller - Disk Test...
I/O Controller - Disk Test SUCCESS!
Executing Supervisor Mode Instructions Test...
Privileged Instruction Test SUCCESS!
```

The program will return a nonzero exit code on failure, which MAY be used to automatically detect failures of the tested systems and take appropriate action.

Please refer to the *amtu*(8) man page for more details.

4.12 Setting the system time and date

You **MUST** verify periodically that the system clock is sufficiently accurate, otherwise log and audit files will contain misleading information. When starting the system, the time and date are copied from the computer's hardware clock to the kernel's software clock, and written back to the hardware clock on system shutdown.

All internal dates and times used by the kernel, such as file modification stamps, use universal time (UTC), and do not depend on the current time zone settings. Userspace utilities usually adjust these values to the currently active time zone for display. Note that text log files will contain ASCII time and date representations in local time, often without explicitly specifying the time zone.

The `date(1)` command displays the current time and date, and can be used by administrators to set the software clock, using the argument `mmddHHMMyyyy` to specify the numeric month, day, hour, minute and year respectively. For example, the following command sets the clock to May 1st 2004, 1pm in the local time zone:

```
date 050113002004
```

The `hwclock(8)` can query and modify the hardware clock on supported platforms, but is not available in virtual environments such as z/VM or LPAR. The typical use is to copy the current value of the software clock to the hardware clock. Note that the hardware clock **MAY** be running in either local time or universal time, as indicated by the *UTC* setting in the `/etc/sysconfig/clock` file. The following command sets the hardware clock to the current time using UTC:

```
hwclock -u -w
```

Use the command `tzselect(8)` to change the default time zone for the entire system. Note that users **MAY** individually configure a different time zone by setting the `TZ` environment variable appropriately in their shell profile, such as the `$HOME/.bashrc` file.

5 Monitoring, Logging & Audit

5.1 Reviewing the system configuration

It is **RECOMMENDED** that you review the system's configuration at regular intervals to verify if it still agrees with the evaluated configuration. This primarily concerns those processes that may run with 'root' privileges.

The permissions of the device files `/dev/*` **MUST NOT** be modified.

In particular, review settings in the following files and directories to ensure that the contents and permissions have not been modified:

```
/etc/at.allow  
/etc/at.deny  
/etc/audit/*  
/etc/cron.allow  
/etc/cron.d/*  
/etc/cron.daily/*  
/etc/cron.deny  
/etc/cron.hourly/*  
/etc/cron.monthly/*  
/etc/cron.weekly/*  
/etc/crontab  
/etc/group  
/etc/gshadow
```

```

/etc/hosts
/etc/inittab
/etc/ld.so.conf
/etc/login.defs
/etc/modules.conf
/etc/pam.d/*
/etc/passwd
/etc/rc.d/init.d/*
/etc/securetty
/etc/security/opasswd
/etc/shadow
/etc/ssh/ssh_config
/etc/ssh/sshd_config
/etc/stunnel/*
/etc/sysconfig/*
/etc/vsftpd.ftpusers
/etc/vsftpd/vsftpd.conf
/etc/xinetd.conf

/var/log/audit.d/*
/var/log/faillog
/var/log/lastlog
/var/spool/at/*
/var/spool/cron/*

```

Use the command `lastlog` to detect unusual patterns of logins.

Also verify the output of the following commands (run as 'root'):

```

atq
crontab -l
find / \( -perm -4000 -o -perm -2000 \) -ls
find / \( -type f -o -type d -o -type b \) -perm -0002 -ls

find /bin /boot /etc /lib /sbin /usr \
    ! -type l \( ! -uid 0 -o -perm +022 \)

```

5.2 System logging and accounting

System log messages are stored in the `/var/log/` directory tree in plain text format, most are logged through the `syslogd(8)` and `klogd(8)` programs, which MAY be configured via the `/etc/syslog.conf` file.

The `logrotate(8)` utility, launched from `/etc/cron.daily/logrotate`, starts a fresh log file every week or when they reach a maximum size and automatically removes or archives old log files. You MAY change the configuration files `/etc/logrotate.conf` and `/etc/logrotate.d/*` as required.

In addition to the `syslog` messages, various other log files and status files are generated in `/var/log` by other programs:

File	Source
<code>audit.d</code>	Directory for LAuS logs
<code>boot.msg</code>	Messages from system startup
<code>lastlog</code>	Last successful log in (see <code>lastlog(8)</code>)
<code>vsftpd.log</code>	Transaction log of the VSFTP daemon

localmessages	Written by syslog
mail	Written by syslog, contains messages from the MTA (postfix)
messages	Written by syslog, contains messages from su and ssh
news/	syslog news entries (not used in the evaluated configuration)
warn	Written by syslog
wtmp	Written by the PAM subsystem, see <code>who(1)</code>

Please see *syslog(3)*, *syslog.conf(5)* and *syslogd(8)* man pages for details on syslog configuration.

The *ps(1)* command can be used to monitor the currently running processes. Using `ps faux` will show all currently running processes and threads.

5.3 Configuring the audit subsystem

The audit subsystem implements a central monitoring solution to keep track of security relevant events, such as changes and change attempts to security critical files.

This is accomplished through two separate mechanisms. All system calls are intercepted, and the kernel writes the parameters and return value to the audit log for those calls that are marked as security relevant in the filter configuration. In addition, some trusted programs contain audit-specific code to write audit trails of the actions they are requested to perform.

Please see *auditd(8)*, *laus(7)*, *auditd.conf(5)*, *aucat(8)* and *augrep(8)* for details.

5.3.1 Intended usage of the audit subsystem

The Controlled Access Protection Profile (CAPP) specifies the auditing capabilities that a compliant system must support. The evaluated configuration described here is based on these requirements.

WARNING: Some of the CAPP requirements may conflict with your specific requirements for the system. For example, a CAPP-compliant system **MUST** disable logins if the audit subsystem is not working. Please ensure that you are aware of the consequences if you enable auditing.

CAPP is designed for a multiuser system, with multiple unique users who maintain both shared and private resources. The auditing features are intended to support this mode of operation with a reliable trail of security-relevant operations. It is less useful for a pure application server with no interactive users.

Please be aware that the auditing subsystem will, when activated, cause some slowdown for applications on the server. The impact depends on what the application is doing and how the audit subsystem is configured. As a rule of thumb, applications that open a large number of separate files are most affected, and CPU-bound programs should not be measurably affected. You will need to balance the performance requirements against your security needs when deciding if and how you want to use auditing.

5.3.2 Selecting the events to be audited

You **MAY** make changes to the set of system calls and events that are to be audited. CAPP requires that the system has the *capability* to audit security relevant events, but it is up to you to choose how you want to use these capabilities. It is acceptable to turn off system call auditing completely even in an evaluated configuration, for example on a pure application server with no interactive users on the system.

The configuration file */etc/audit/filter.conf* by default contains a suggested setup for a typical multiuser system, all access to the security relevant files (as configured in */etc/audit/filter.conf* and */etc/audit/filesets.conf*) is audited, along with other security relevant events such as system reconfiguration.

You **MAY** selectively disable and enable auditing for specific events or users as required by setting up predicates and filters in the *filter.conf* file. The following excerpt from the default configuration is an example:

```

predicate is-non-root-uid = !eq(0);
filter not-root-user = is-non-root-uid(login-uid);

tag "Open_Denied"
syscall open = denied(result)
                && (( not-root-user || effectivenonroot )
                && is-sysdir(arg0));

```

Please refer to the *audit-filter(5)* man page for more details.

5.3.3 Reading and searching the audit records

Use the *aucat(8)* and *augrep(8)* tools to retrieve information from the audit logs. The information available for retrieval depends on the active filter configuration. If you modify the filter configuration, it is **RECOMMENDED** keeping a dated stamped copy of the applicable configuration with the log files for future reference.

For example:

```

# view the last 100 audit records
aucat | tail -100

# view all successful PAM authentications
augrep -e TEXT -U AUTH_success

# all actions recorded for a specified login UID (this includes
# actions done by this user with a different effective UID,
# for example, via SUID programs or as part of a "su" session)
augrep -l kw

# file removals
augrep -e SYSCALL -S unlink

```

Of course, you can use other tools such as plain *grep(1)* or scripting languages such as *awk(1)*, *python(1)* or *perl(1)* to further analyze the text output generated by the low-level audit tools.

5.3.4 Starting and stopping the audit subsystem

The audit subsystem is only active when all of the following conditions are met:

- The *audit.o* kernel module must be loaded.
- The audit daemon *auditd* must be running.
- Processes are attached to the audit subsystem by explicitly launching them with the *aurun(8)* wrapper program; starting them from an interactive shell session that used the *pam.laus.so* PAM module when logging in; or when syscall auditing is enabled globally for all processes (setting *dev.audit.attach-all=1* in */etc/sysctl.conf*).

If the audit daemon is terminated, no audit events are generated until it is restarted. To avoid lost audit records when you have modified the filter configuration, you **MUST** use the command *auditd -r* to re-load the filters.

WARNING: *auditd -r* will *not* reload */etc/audit/audit.conf*, it only reloads the filter configuration file. To activate changes to this configuration file you **MUST** restart the audit daemon:

```
/etc/rc.d/init.d/audit restart
```

You **MUST NOT** attempt to reload the configuration by sending *auditd* a *HUP* signal or by running `/etc/rc.d/init.d/audit reload`, because that will not write the required audit record showing the reconfiguration. You **MUST** use one of the two restart methods described above.

If the audit module is unloaded with *rmmmod*, all processes are detached permanently from the audit subsystem. They can only be re-attached when using the *AUDIT_ATTACH_ALL=1* option in */etc/sysconfig/audit*.

5.3.5 Storage of audit records

The **REQUIRED** operating mode for the audit records is "bin mode" ("bin" as in bucket), using several preallocated files of constant size for the audit records. *auditd* will write data to the first file until it is filled, then switch to the next one re-using each one in turn in a round-robin fashion.

Each time a bin is filled, *auditd* will launch the configured notification program to process the file. The default configuration saves a copy of each filled file before re-using the storage. If the notification program exits with a failure status, for example, due to lack of disk space, *auditd* will then take the configured action, by default setting the message queue size to zero and thereby blocking all processes that try to write new records. These audited processes will sleep until *auditd* resumes processing (typically once disk space has been freed by the administrator), then they will be woken up by the kernel and proceed running normally.

You **MAY** instead configure round-robin reuse of the files without saving, to keep the disk space used by the audit logs constant. To do that, remove the `"-S /var/log/audit.d/save.%u"` option in */etc/audit/audit.conf*. In this configuration, you have access to a fixed amount of historical audit data, but any new events will cyclically overwrite old data. A user could exploit this mechanism by intentionally generating a large number of irrelevant entries to wipe out the previously generated records. The default configuration uses four files of only 20 MiB size each. You **SHOULD** increase these numbers in */etc/audit/audit.conf* according to available disk space, your organizational requirements, and the system's usage patterns to ensure that a sufficient amount of historic audit data will be saved.

5.3.6 Reliability of audit data

By default, the audit records are written using the normal Linux filesystem buffering, which means that information may be lost in a crash because it has not been written to the physical disk yet. Any applications that read the records while the system is running will always get the most current data out of the buffer cache, even if it has not yet been committed to disk, so this does not affect normal operation. If you want to ensure that *auditd* always forces a disk write for each record, you **MAY** set the `"sync = yes;"` option in */etc/audit/audit.conf*, but be aware that this will result in significantly reduced performance and high strain on the disk.

The audit record files are *not* protected against a malicious administrator, and are not intended for an environment where the administrators are not trustworthy.

5.4 System configuration variables in */etc/sysconfig*

The system uses various files in */etc/sysconfig* to configure the system. Most files in this directory tree contain variable definitions in the form of shell variables that are either read by the rc scripts at system boot time or are evaluated by other commands at runtime. Note that changes will not take effect until the affected service is restarted or the system is rebooted.

6 Security guidelines for users

6.1 Online Documentation

The system provides a large amount of online documentation, usually in text format. Use the `man` program to read entries in the online manual, for example:

```
man ls
man man
```

to read information about the `ls` and `man` commands respectively. You can search for keywords in the online manual with the `apropos(1)` utility, for example:

```
apropos password
```

When this guide refers to manual pages, it uses the syntax `ENTRY(SECTION)`, for example `ls(1)`. Usually you do not need to provide the section number, but if there are several entries in different sections, you can use the optional `-S` switch and pick a specific one.

Some programs provide additional information GNU 'texinfo' format, use the `info` program to read it, for example:

```
info diff
```

Additional information, sorted by software package, can be found in the `/usr/share/doc/*/` directories. Use the `less(1)` pager to read it, for example:

```
/usr/share/doc/bash*/FAQ
```

Many programs also support a `--help`, `-?` or `-h` switch you can use to get a usage summary of supported command-line parameters.

A collection of How-To documents in HTML format can be found under `/usr/share/doc/howto/en/html` if the optional `howtoenh` package is installed.

Please see `/usr/share/doc/howto/en/html/Security-HOWTO` for security information. The HTML files can be read with the `w3m` browser.

The RHEL server documentation is also available in electronic form in the directories `/usr/share/doc/rhel*`.

Note that this Configuration Guide has precedence over other documents in case of conflicting recommendations.

6.2 Authentication

You **MUST** authenticate (prove your identity) before being permitted to use the system. When the administrator created your user account, he or she will have assigned a user name and default password, and provided that information for you along with instructions how to access the system.

Logging in to the system will usually be done using the Secure Shell (SSH) protocol, alternatively a serial terminal may be available. Use the `ssh` command to connect to the system unless instructed otherwise by the administrator, for example:

```
ssh jdoe@172.16.0.1
```

The *ssh(1)* manual page provides more information on available options. If you need to transfer files between systems, use the *scp(1)* or *sftp(1)* tools.

If this is the first time you are connecting to the target system, you will be prompted if you want to accept the host key. If the administrator has provided a key fingerprint for comparison, verify that they match, otherwise type *yes* to continue. You **MUST** immediately change your initially assigned password with the *passwd(1)* utility.

You **MUST NOT** under any circumstances attempt to log in from an insecure device, such as a public terminal or a computer belonging to a friend. Even if the *person* owning the computer is trustworthy, the *computer* may not be due to having been infected with malicious code. Always remember that the device you are typing your password into has the ability to save and re-use your authentication information, so you are in effect giving the computer you are using the right to do any and all actions in your name. Insecure handling of authentication information is the leading cause for exploits of otherwise secure systems, and SSH can only protect the information during transit, and offers no protection at all against an insecure end point.

When you log out from the system and leave the device you have used for access (such as a terminal or a workstation with terminal emulation), you **MUST** ensure that you have not left information on the screen or within an internal buffer that should not be accessible to another user. You should be aware that some terminals also store information not displayed on the terminal (such as passwords, or the contents of a scrollbar buffer). Nevertheless this information may be extractable by the next user unless the terminal buffer has been cleared. Safe options include completely shutting down the client software used for access, powering down a hardware terminal, or clearing the scrollbar buffer by switching among virtual terminals in addition to clearing the visible screen area.

If you ever forget your password, contact your administrator who will be able to assign a new password.

You **MAY** use the *chsh(1)* and *chfn(1)* programs to update your login shell and personal information if necessary. Not all settings can be changed this way, contact your administrator if you need to change settings that require additional privileges.

6.3 Password policy

All users, including the administrators, **MUST** ensure that their authentication passwords are strong (hard to guess) and handled with appropriate security precautions. The password policy described here is designed to satisfy the requirements of the evaluated configuration. If your organization already has a password policy defined, your administrator **MAY** refer you to that policy if it is equivalently strong.

You **MUST** change the initial password set by the administrator when you first log into the system. You **MUST** select your own password in accordance with the rules defined here. You **MUST** also change the password if the administrator has set a new password, for example if you have forgotten your password and requested the administrator to reset the password.

Use the *passwd(1)* program to change passwords. It will first prompt you for your old password to confirm your identity, then for the new password. You need to enter the new password twice, to catch mistyped passwords.

The *passwd(1)* program will automatically perform some checks on your new password to help ensure that it is not easily guessable, but you **MUST** nevertheless follow the requirements in this chapter.

Note that the administrators **MUST** also ensure that their own passwords comply with this password policy, even in cases where the automatic checking is not being done, such as when first installing the system.

- Your password **MUST** be a minimum of 8 characters in length. More than 8 characters **MAY** be used (it is **RECOMMENDED** to use more than 8, best is to use passphrases), and all characters are significant.
- Use at least one character each from the following sets for passwords:

```
Lowercase letters: abcdefghijklmnopqrstuvwxyz
Uppercase letters: ABCDEFGHIJKLMNOPQRSTUVWXYZ
Digits:           0123456789
Punctuation:     !"#%&'()*+,-./:;<=>?[\]^_`{|}~
```

- You **MUST NOT** base the password on a dictionary word, your real name, login name, or other personal details (such as dates, names of relatives or pets), or names of real people or fictional characters.
- Instead of a password, you **MAY** use a passphrase consisting of multiple unrelated words (at least three) joined with random punctuation characters. Such a passphrase **MUST** have a length of at least 16 characters.
- You **MUST NOT** use a simple alphabetic string, palindrome or combinations of adjacent keyboard keys.
- When you choose a new password, it **MUST NOT** be a simple variation or permutation of a previously used one.
- You **MUST NOT** write the password on paper or store it on electronic devices in unprotected form. Storage in a secure location (such as an envelope in a safety deposit box, or encrypted storage on an electronic device) **MAY** be acceptable, contact your administrator first to ensure that the protection is strong enough to make password recovery infeasible for the types of attackers the system is intended to protect against.
- The password is for you and you only. A password is like a toothbrush - you do not want to share it with anybody, even your best friend. You **MUST NOT** disclose your password to anybody else, or permit anybody else to use the system using your identity.

Note that administrators will never ask you for your password, since they do not need it even if they are required to modify settings affecting your user account.

- You **MUST NOT** use the same password for access to any systems under external administration, including Internet sites. You **MAY** however use the same password for accounts on multiple machines within one administrative unit, as long as they are all of an equivalent security level and under the control of the same administrators.
- You **MUST** inform the administrator and select a new password if you have reason to believe that your password was accidentally disclosed to a third party.
- If the system notifies you that your password will expire soon or has expired, choose a new one as instructed. Contact your administrator in case of difficulty.

A **RECOMMENDED** method of generating passwords that fits these criteria while still being easy to memorize is to base it on letters of words in a sentence (NOT a famous quotation), including capitalization and punctuation and one or two variations. Example:

```
"Ask not for whom the bell tolls."
=> An4wtbt.
```

```
"Password 'P'9tw;ciSd' too weak; contained in RHEL documentation"
=> P'9tw;ciRd
```

6.4 Access control for files and directories

Linux is a multiuser operating system. You can control which other users will be able to read or modify your files by setting the Unix permission bits and user/group IDs, or (if more precise control is needed) by using POSIX-style access control lists (ACLs).

Note that the administrators ('root') are able to override these permissions and access all files on the system. Use of encryption is **RECOMMENDED** for additional protection of sensitive data.

The 'umask' setting controls the permissions of newly created files and directories and specifies the access bits that will be *removed* from new objects. Ensure that the setting is appropriate, and never grant write access to others by default. The umask **MUST** include at least the 002 bit (no write access for others), and the **RECOMMENDED** setting is 027 (read-only and execute access for the group, no access at all for others).

Do not set up world-writable areas in the filesystem - if you want to share files in a controlled manner with a fixed group of other users (such as a project group), please contact your administrator and request the creation of a user group for that purpose.

Always remember that **you** are responsible for the security of the data you create and use. Choose permissions that match the protection goals appropriate for the content, and that correspond to your organization's security policy. Access to confidential data **MUST** be on a need-to-know basis, do not make data world-readable unless the information is intended to be public.

Whenever you start a program or script, it will execute with your access rights. This implies that a malicious program would be able to read and modify all files that you have access to. Never execute any code that you have received from untrustworthy sources, and do not run commands that you do not understand.

Programs can be configured to run with the access rights of the program file's owner and/or group instead of the rights of the calling user. This is the SUID/SGID mechanism, which utilities such as *passwd(1)* use to be able to access security-critical files. You could also create your own SUID/SGID programs via *chmod(1)*, but **DO NOT** do that unless you fully understand the security implications - you would be giving away *your* access privileges to whoever launches the SUID program. Please refer to the "Secure Programming HOWTO" in the unlikely case that you need to create such a program, there you will find explanations of the many aspects that must be considered, such as the risk of unintended shell escapes, buffer overflows, resource exhaustion attacks and many other factors. Note that SUID root programs **MUST NOT** be added to the evaluated configuration, the only permitted use of the SUID bit is for setting non-root user IDs.

Please refer to the *chmod(1)*, *umask(2)*, *chown(1)*, *chgrp(1)*, *acl(5)*, *getfacl(1)*, and *setfacl(1)* manual pages for information, or any of the many available books covering Linux security (cf. Appendix 'Literature'), or ask your system administrator for advice.

6.5 Data import / export

The system comes with various tools to archive data (*tar*, *star*, *cpio*). If ACLs are used, then only *star* **MUST** be used to handle the files and directories as the other commands do not support ACLs. The options *-H=exustar -acl* must be used with *star*.

Please see the *star(1)* man page for more information.

7 Appendix

7.1 Online Documentation

If there are conflicting recommendations in this guide and in one of the sources listed here, the Configuration Guide has precedence concerning the evaluated configuration.

"Red Hat Enterprise Linux 3 Installation Guide for the x86, Itanium and AMD64 Architectures",
</usr/share/doc/rhel-ig-x8664-multi-en-3/index.html>

"Red Hat Enterprise Linux 3 Installation Guide for the IBM eServer iSeries and IBM eServer pSeries Architectures",
</usr/share/doc/rhel-ig-ppc-multi-en-3/index.html>

"Red Hat Enterprise Linux 3 Installation Guide for the IBM S/390 and IBM eServer zSeries Architectures",
</usr/share/doc/rhel-ig-s390-multi-en-3/index.html>

"Red Hat Enterprise Linux 3 System Administration Guide", </usr/share/doc/rhel-sag-en-3/index.html>

"Red Hat Enterprise Linux 3 Reference Guide", </usr/share/doc/rhel-rg-en-3/index.html>

"Red Hat Enterprise Linux 3 Security Guide", </usr/share/doc/rhel-sg-en-3/index.html>

David A. Wheeler, "Secure Programming for Linux and Unix HOWTO", /usr/share/doc/howto/en/html_single/Secure-Programs-HOWTO.html, <http://tldp.org/HOWTO/Secure-Programs-HOWTO/>
 Kevin Fenzi, Dave Wreski, "Linux Security HOWTO", /usr/share/doc/howto/en/html_single/Security-HOWTO.html, <http://www.linuxsecurity.com/docs/LDP/Security-HOWTO/>

7.2 Literature

Ellen Siever, Stephen Spainhour, Stephen Figgins, & Jessica P. Hekman, "Linux in a Nutshell, 3rd Edition", O'Reilly 2000, ISBN 0596000251

Simson Garfinkel, Gene Spafford, Alan Schwartz, "Practical Unix & Internet Security, 3rd Edition", O'Reilly 2003, ISBN 0596003234

Aeleen Frisch, "Essential System Administration, 3rd Edition", O'Reilly 2002, ISBN 0596003439

Daniel J. Barrett, Richard Silverman, "SSH, The Secure Shell: The Definitive Guide", O'Reilly 2001, ISBN 0596000111

David N. Blank-Edelman, "Perl for System Administration", O'Reilly 2000, ISBN 1565926099

Shelley Powers, Jerry Peek, Tim O'Reilly, Mike Loukides, "Unix Power Tools, 3rd Edition", O'Reilly 2002, ISBN 0596003307

W. Richard Stevens, "Advanced Programming in the UNIX(R) Environment", Addison-Wesley 1992, ISBN 0201563177

Linda Mui, "When You Can't Find Your UNIX System Administrator", O'Reilly 1995, ISBN 1565921046

7.3 The file `/etc/audit/audit.conf`

```
# kernel interface
device-file = "/dev/audit";

# filter config
filter-config = "/etc/audit/filter.conf";

# Standard output method is bin mode.
#
output {
    mode           = bin;
    num-files      = 4;
    file-size      = 20M;
    file-name      = "/var/log/audit.d/bin";
    notify         = "/usr/sbin/audbin -S /var/log/audit.d/save.%u -C";

    # The following symlink is created whenever we switch to
    # a new bin.
    current        = "/var/log/audit";

    sync          = no;

# uncomment these to cause audit records to be
# flushed to the disk after sync-after records
# are written to the log
```

```

#      sync                = yes;
#      sync-after          = 16;
#      error {
#          action {
#              type = suspend;
#          };
#      };
};

# Alternatively, write to /var/log/audit in normal
# append mode
# output {
#     mode                = append;
#     file-name            = "/var/log/audit";
#     sync                = yes;
# };

# Alternative output
# output {
#     mode                = stream;
#     command              = "/usr/local/sbin/send_to_syslog"
# };

# Disk usage thresholds.
# These thresholds are checked at regular intervals when
# append mode is used.
# (bin mode doesn't require these checks as the bin files
# are preallocated).
threshold disk-space-low {
    space-left = 10M;
    action {
        type = syslog;
        facility = security;
        priority = warning;
    };
    action {
        type = notify;
        command = "/usr/local/bin/page-admin";
    };
    action {
        type = audit;
        event = AUDIT_disklow;
    };
};

threshold disk-full {
    space-left = 20K;
    action {
        type = syslog;
        facility = security;
        priority = crit;
    };
    action {

```

```

        type = audit;
        event = AUDIT_diskfull;
    };
};

```

7.4 The file `/etc/audit/filter.conf`

[illegible]

```

#
# Predicates for audit-tags
predicate      is-o-creat          = mask(O_CREAT, O_CREAT);
predicate      is-ipc-remove       = eq(IPC_RMID);
predicate      is-ipc-setperms     = eq(IPC_SET);
predicate      is-ipc-creat        = mask(IPC_CREAT, IPC_CREAT);
predicate      is-auditdevice      = prefix("/dev/audit");
predicate      is-cmd-set-auditid  = eq(AUIOCSETAUDITID);
predicate      is-cmd-set-loginid  = eq(AUIOCLOGIN);

#
# Misc filters
filter         is-root            = is-null(uid);
filter         is-setuid          = is-null(dumpable);
filter         syscall-failed     = is-negative(result);
filter         syscall-addr-succeed = is-lower-1024(result);
predicate      is-af-packet       = eq(AF_PACKET);
predicate      is-af-netlink      = eq(AF_NETLINK);
predicate      is-sock-raw        = eq(SOCK_RAW);

#
# Include filesets.
#
include "filesets.conf";

#
# "Secret" files should not be read by everyone -
# we also log read access to these files
#
predicate      is-secret = prefix(@secret-files);

#
# All regular files owned by a system uid are deemed sensitive
#
predicate      is-system-file = is-system-uid(file-uid)
                                && !prefix("/var")
                                && !is-world-writable(file-mode);

#
# Define ioctls we track
#
set            sysconf-ioctls = {
    SIOCADDLDCI,
    SIOCADMULTI,
    SIOCADDRT,
    SIOCBONDCHANGEACTIVE,
    SIOCBONDENSLAVE,
    SIOCBONDRELEASE,
    SIOCBONDSETHWADDR,
    SIOCDAEP,
    SIOCDELDCI,
    SIOCDELMULTI,
    SIOCDELRT,

```

```

        SIOCDIFADDR,
        SIOCDRARP,
        SIOCETHTOOL,
        SIOCGIFBR,
        SIOCSARP,
        SIOCSIFADDR,
        SIOCSIFBR,
        SIOCSIFBRDADDR,
        SIOCSIFDSTADDR,
        SIOCSIFENCAP,
        SIOCSIFFLAGS,
        SIOCSIFHWADDR,
        SIOCSIFHWBROADCAST,
        SIOCSIFLINK,
        SIOCSIFMAP,
        SIOCSIFMEM,
        SIOCSIFMETRIC,
        SIOCSIFMTU,
        SIOCSIFNAME,
        SIOCSIFNETMASK,
        SIOCSIFPFLAGS,
        SIOCSIFSLAVE,
        SIOCSIFTXQLEN,
        SIOCSMIIREG
    };
    predicate is-sysconf-ioctl      = eq(@sysconf-ioctls);

#
# System calls on file names
#
set    file-ops = {
        "mkdir", "rmdir", "unlink",
        "chmod",
        "chown", "lchown",
        "chown32", "lchown32",
    };

#
# General system related ops
#
set    system-ops = {
        swapon, swapoff,
        create_module, init_module, delete_module,
        sethostname, setdomainname,
    };

set    priv-ops = {
        "setuid",
        "setuid32",
        "seteuid",
        "seteuid32",
        "setreuid",
        "setreuid32",
    };

```

```

        "setresuid",
        "setresuid32",
        "setgid",
        "setgid32",
        "setegid",
        "setegid32",
        "setregid",
        "setregid32",
        "setresgid",
        "setresgid32",
        "setgroups",
        "setgroups32",
        "capset",
};

#
# Audit-Tags (only syscall related tags are handled here)
#

# define sets of syscalls related to audit-tags

# System calls for changing file modes
set    mode-ops = {
        "chmod",
        "fchmod",
};

# System calls for changing file owner
set    owner-ops = {
        "chown", "lchown",
        "chown32", "lchown32",
        "fchown",
};

# System calls doing file link operations
set    link-ops = {
        "link", "symlink",
};

# System calls for creating device files
set    mknod-ops = {
        "mknod",
};

# System calls for opening a file
set    open-ops = {
        "open",
};

# File renaming
set    rename-ops = {
        "rename",
};

```

```
# File truncation
set    truncate-ops = {
        "truncate", "truncate64",
        "ftruncate", "ftruncate64",
    };

# Unlink files
set    unlink-ops = {
        "unlink",
    };

# Deletion of directories
set    rmdir-ops = {
        "rmdir",
    };

# Mounting of filesystems
set    mount-ops = {
        "mount",
    };

# Unmounting of filesystems
set    umount-ops = {
        "umount",
        "umount2"
    };

# Changing user (-role)
set    userchange-ops = {
        "setuid",
        "setuid32",
        "seteuid",
        "seteuid32",
        "setreuid",
        "setreuid32",
        "setresuid",
        "setresuid32",
    };

# Execute another program
set    execute-ops = {
        "execve",
    };

# Set real user-ID
set    realuid-ops = {
        "setuid",
        "setuid32",
    };

# Set user-IDS in gernerall
set    setuserids-ops = {
```



```

        "setuid",
        "setuid32",
        "seteuid",
        "seteuid32",
        "setreuid",
        "setreuid32",
        "setresuid",
        "setresuid32",
};

# Set real group-ID
set    realgid-ops = {
        "setgid",
        "setgid32",
        "setgroups",
        "setgroups32",
};

# Set group-IDs in gernerall
set    setgroups-ops = {
        "setgid",
        "setgid32",
        "setegid",
        "setegid32",
        "setregid",
        "setregid32",
        "setresgid",
        "setresgid32",
        "setgroups",
        "setgroups32",
};

# Set other kind of privileges (capabilities)
set    privilege-ops = {
        "capset",
};

# Change system-time
set    timechange-ops = {
        "adjtimex",
        "stime",
        "settimeofday",
};

# bring sets and tags in conjunction

tag "FILE_mode"
syscall @mode-ops = always;

tag "FILE_owner"
syscall @owner-ops = always;

```

```
tag "FILE_link"
syscall @link-ops = always;

tag "FILE_mknod"
syscall @mknod-ops = always;

tag "FILE_create"
syscall open = is-o-creat(arg1);
tag "FILE_create"
syscall creat = always;

#tag "FILE_open"
#syscall @open-ops = always;

tag "FILE_open"
syscall @open-ops = (is-system-file(arg0) && !(is-rdonly(arg1)))
                    || is-secret(arg0);

tag "FILE_rename"
syscall @rename-ops = always;

tag "FILE_truncate"
syscall @truncate-ops = always;

tag "FILE_unlink"
syscall @unlink-ops = always;

tag "FS_rmdir"
syscall @rmdir-ops = always;

tag "FS_mount"
syscall @mount-ops = always;

tag "FS_umount"
syscall @umount-ops = always;

# I think owner changing doesnt make much sense
tag "MSG_owner"
syscall msgctl = is-ipc-setperms(arg1);

tag "MSG_mode"
syscall msgctl = is-ipc-setperms(arg1);

tag "MSG_delete"
syscall msgctl = is-ipc-remove(arg1);

tag "MSG_create"
syscall msgget = always;

tag "SEM_owner"
syscall semctl = is-ipc-setperms(arg2);
```

```
tag "SEM_mode"
syscall semctl = is-ipc-setperms(arg2);

tag "SEM_delete"
syscall semctl = is-ipc-remove(arg2);

tag "SEM_create"
syscall semget = always;

tag "SHM_owner"
syscall shmctl = is-ipc-setperms(arg1);

tag "SHM_mode"
syscall shmctl = is-ipc-setperms(arg1);

tag "SHM_delete"
syscall shmctl = is-ipc-remove(arg1);

tag "SHM_create"
syscall shmget = always;

tag "PRIV_userchange"
syscall @userchange-ops = always;

tag "PROC_execute"
syscall @execute-ops = always;

tag "PROC_realuid"
syscall @realuid-ops = always;

tag "PROC_auditid"
syscall ioctl = (is-auditdevice(arg0) && is-cmd-set-auditid(arg1));

tag "PROC_loginid"
syscall ioctl = (is-auditdevice(arg0) && is-cmd-set-loginid(arg1));

tag "PROC_setuserids"
syscall @setuserids-ops = always;

tag "PROC_realgid"
syscall @realgid-ops = always;

tag "PROC_setgroups"
syscall @setgroups-ops = always;

tag "PROC_privilege"
syscall @privilege-ops = always;

tag "SYS_timechange"
syscall @timechange-ops = always;
```

```
# not required by CAPP
syscall ipc = always;

syscall socket = is-af-packet(arg0) || is-sock-raw(arg1);
syscall ioctl = is-sysconf-ioctl(arg1);

#
# Special filters for process/termination
event process-exit = if-crash-signal(exitcode);

#
# Events we want to log unconditionally:
event network-config = always;
event user-message = always;
event process-login = always;
```

7.5 The file `/etc/audit/filesets.conf`

```
#
# This file contains file name sets etc used in the default
# audit filter configuration file.
#
# The syntax of this file is described in filter.conf(5).
#

#
# Set of files for which we track read access.
#
set          secret-files = {
    "/etc/shadow",
    "/etc/gshadow",
    "/var/log/audit",
    "/var/log/audit.d",
    "/var/log/audit.d/bin.0",
    "/var/log/audit.d/bin.1",
    "/var/log/audit.d/bin.2",
    "/var/log/audit.d/bin.3",
};
```