

Common Criteria EAL4+ Evaluated Configuration Guide for SUSE LINUX Enterprise Server on IBM Hardware

Klaus Weidner <klaus@atsec.com>

January 17, 2005; v1.20

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	6
1.1	Purpose of this document	6
1.2	How to use this document	6
1.3	What is a CC compliant system?	6
1.3.1	Hardware requirements	7
1.3.2	Software requirements	7
1.3.3	Environmental requirements	7
1.3.4	Operational requirements	7
1.4	Requirements for the system's environment	8
1.5	Requirements for the system's users	8
1.6	Overview of the system's security functions	9
1.6.1	Identification and authentication	9
1.6.2	Audit	9
1.6.3	Discretionary access control	9
1.6.4	Object reuse	9
1.6.5	Security management and system protection	10
1.6.6	Secure communication	10
1.7	Overview of security relevant events	10
2	Installation	10
2.1	Supported hardware	10
2.2	Selection of install options and packages	11
3	Secure initial system configuration	15
3.1	Prerequisites	16
3.1.1	Filesystem configuration	16
3.1.2	Getting required updates	16
3.2	Automated configuration of the system	18
3.3	Add and remove packages	20
3.4	Disable services	23
3.4.1	Disable usbfs	24
3.5	Remove SUID/SGID root settings from binaries	25
3.6	Update permissions for 'su'	25
3.7	Disable root login over the network	26
3.8	Setting up SSH	26
3.9	Setting up xinetd	27
3.10	Setting up FTP	27
3.11	Setting up Postfix	28
3.12	Setting up the audit subsystem	29
3.12.1	Setting up the audit configuration files	29
3.12.2	Starting auditd at boot as a system service	29
3.12.3	Starting auditd in fail-secure mode from init (OPTIONAL)	29
3.12.4	Ensuring that the PAM library is audit enabled	30
3.13	Introduction to Pluggable Authentication Module (PAM) configuration	30
3.14	Required Pluggable Authentication Module (PAM) configuration	32
3.14.1	/etc/pam.d/chage	32
3.14.2	/etc/pam.d/chfn	32
3.14.3	/etc/pam.d/chsh	33
3.14.4	/etc/pam.d/login	33
3.14.5	/etc/pam.d/other	34
3.14.6	/etc/pam.d/passwd	34
3.14.7	/etc/pam.d/sshd	34

3.14.8	/etc/pam.d/su	34
3.14.9	/etc/pam.d/useradd	35
3.14.10	/etc/pam.d/vsftpd	35
3.14.11	/etc/security/pam_pwcheck.conf	35
3.14.12	/etc/security/pam_unix2.conf	36
3.15	Setting up login controls	36
3.15.1	Maintaining <i>cracklib</i> dictionaries	38
3.16	Configuring the boot loader	38
3.16.1	GRUB boot loader configuration	38
3.16.2	Yaboot boot loader configuration	39
3.16.3	ZIPL boot loader configuration	39
3.16.4	iSeries kernel slots	40
3.17	Reboot and initial network connection	40
4	System operation	40
4.1	System startup, shutdown and crash recovery	41
4.2	Backup and restore	41
4.3	Gaining superuser access	42
4.4	Installation of additional software	42
4.5	Scheduling processes using cron and at	43
4.6	Mounting filesystems	44
4.7	Managing user accounts	45
4.8	Using serial terminals	47
4.9	SYSV shared memory and IPC objects	47
4.10	Configuring secure network connections with <i>stunnel</i>	47
4.10.1	Introduction	47
4.10.2	Creating an externally signed certificate	49
4.10.3	Creating a self-signed certificate	50
4.10.4	Activating the tunnel	51
4.10.5	Using the tunnel	53
4.10.6	Example 1: Secure SMTP delivery	53
4.10.7	Example 2: Simple web server	53
4.10.8	Example 3: system status view	54
4.11	The Abstract Machine Testing Utility (AMTU)	55
4.12	Setting the system time and date	55
5	Monitoring, Logging & Audit	56
5.1	Reviewing the system configuration	56
5.2	System logging and accounting	57
5.3	Configuring the audit subsystem	58
5.3.1	Intended usage of the audit subsystem	58
5.3.2	Selecting the events to be audited	58
5.3.3	Reading and searching the audit records	59
5.3.4	Starting and stopping the audit subsystem	59
5.3.5	Storage of audit records	60
5.3.6	Reliability of audit data	60
5.4	System configuration variables in <i>/etc/sysconfig</i>	60
5.4.1	<i>suseconfig</i>	60
5.4.2	<i>security</i>	61
5.4.3	<i>cron</i>	61
5.4.4	<i>language</i>	61
5.4.5	<i>backup</i>	61
5.4.6	<i>boot</i>	61
5.4.7	<i>displaymanager</i>	61

5.4.8	<i>kernel</i>	61
5.4.9	<i>clock</i>	61
5.4.10	<i>proxy</i>	62
5.4.11	<i>windowmanager</i>	62
5.4.12	<i>sysctl</i>	62
5.4.13	<i>java</i>	62
5.4.14	<i>mail</i>	62
5.4.15	<i>hardware</i>	62
5.4.16	<i>printer</i>	62
5.4.17	<i>news</i>	62
5.4.18	<i>console</i>	62
5.4.19	<i>keyboard</i>	63
5.4.20	<i>mouse</i>	63
5.4.21	<i>lvm</i>	63
5.4.22	<i>network</i>	63
5.4.23	<i>syslog</i>	63
5.4.24	<i>SuSEfirewall2</i>	63
5.4.25	<i>hotplug</i>	63
5.4.26	<i>ssh</i>	63
5.4.27	<i>postfix</i>	63
5.4.28	<i>bootloader</i>	63
5.4.29	<i>audit</i>	63
6	Security guidelines for users	64
6.1	Online Documentation	64
6.2	Authentication	64
6.3	Password policy	65
6.4	Access control for files and directories	66
6.5	Data import / export	67
7	Appendix	67
7.1	Online Documentation	67
7.2	Literature	68
7.3	The file <code>/etc/audit/audit.conf</code>	68
7.4	The file <code>/etc/audit/filter.conf</code>	69
7.5	The file <code>/etc/audit/filesets.conf</code>	78

1 Introduction

1.1 Purpose of this document

The SUSE LINUX Enterprise Server (SLES) 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 zSeries model z800, z890, z900, or z990
    executes in VM 5.1 LPAR
```

```
IBM iSeries (SF2) model 520
    machine type 9406
    executes in OS/400 5.3 LPAR
```

```
IBM pSeries (SF2) model 520
    machine type 9111
    executes in static LPAR
```

```
IBM xSeries model 335
    machine type 8676
```

```
IBM eServer model e325 (based on AMD64/Opteron processor)
    machine type 8835
```

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.

For more information about these requirements, please refer to section §1.4 "Requirements for the system's environment" of this guide.

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 SLES, networked in a non-hostile network, with a well-managed and non-hostile user community. It is not intended to address the needs of a directly 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.

If other systems are connected to the network they **MUST** be configured and managed by the same authority using an appropriate security policy not conflicting with the security policy of the target of evaluation. All links from this network to untrusted networks (such as the Internet) need to be protected by appropriate measures like carefully configured firewall systems that prevent attacks from the untrusted networks.

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 SLES 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 §3.13 "Introduction to Pluggable Authentication Module (PAM) configuration", §3.14 "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 §3.8 "Setting up SSH" of this guide; and the *sshd(8)*, *ssh(1)*, *sshd_config(5)* man pages.

vsftpd

Section §3.10 "Setting up FTP" of this guide; and the *vsftpd(8)*, *vsftpd.conf(5)* man pages.

su

Sections §3.6 "Update permissions for 'su'", §4.3 "Gaining superuser access" of this guide; and the *su(8)* man page.

1.6.2 Audit

Sections §3.12 "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 SLES 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*(8) man page.

Section §3.8 "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 SLES 9, 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.

USB keyboards and mice *MAY* be attached, as some of the supported hardware platforms would otherwise not have supported console input devices. If a USB keyboard or mouse is used, it *MUST* be connected before booting the operating system, and *NOT* added later to a running system. Other hot-pluggable hardware that depends on the dynamic loading of kernel modules *MUST NOT* be attached. Examples of such unsupported hardware are USB and IEEE1394/FireWire peripherals other than mice and keyboards.

2.2 Selection of install options and packages

This section describes the detailed steps to be performed when installing the SLES 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 when using a NFS file server instead of CD-ROMs). 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. Verify that the installation CD is an authentic SUSE distribution CD for SLES 9. The original CD is shipped in a sealed sleeve with the label "SUSE LINUX ENTERPRISE SERVER Installation". If using downloaded ISO images, you *MUST* verify that the MD5 checksums of the image files are correct. Run `md5sum *.iso` to view the checksums for the downloaded images (note that "RC5" is the name of the final release), and compare them with those shown in this list:

cc419d86f3f5ff99395ca4de9d967600	SLES-9-i386-RC5-CD1.iso
86e97184aae42ba6013ea7460372ffe5	SLES-9-i386-RC5-CD2.iso
f880b3ba92fc43add18259c9437f648d	SLES-9-i386-RC5-CD3.iso
bc7b88f34a8142bacbdd4d1fddd3fc50	SLES-9-i386-RC5-CD4.iso
7844c76fc9f39a2af9ef6751ec18af60	SLES-9-i386-RC5-CD5.iso
9e0fdd835e52f53906dff110515eb002	SLES-9-i386-RC5-CD6.iso
bafc5da257b993a4bf8064674cc1d5b5	SLES-9-ppc-RC5-CD1.iso
61432643f16855bb24f802cbd4d0efdf	SLES-9-ppc-RC5-CD2.iso
cebc4110b4f214a5b6721d716c63bcf0	SLES-9-ppc-RC5-CD3.iso
bf1e29329f2f465d775d6f57920d453a	SLES-9-ppc-RC5-CD4.iso
7515fb38f36a206c7ea98c31df63ebab	SLES-9-ppc-RC5-CD5.iso
0b82bd2547969605051428f250917a76	SLES-9-ppc-RC5-CD6.iso
1fb270a7123895abc846a377c3d1dfdd	SLES-9-x86-64-RC5-CD1.iso
1579c82c85fe71b4cc68685d638c883d	SLES-9-x86-64-RC5-CD2.iso
5330a150eb51542d378dbe281f2315a2	SLES-9-x86-64-RC5-CD3.iso
a9d5b0dddd45f378f97ea9bca2b12b11	SLES-9-x86-64-RC5-CD4.iso
dc6de0d488e4dbf6aef322f7aadba7dd	SLES-9-x86-64-RC5-CD5.iso
2324e164faab72df89e4078b6a7c1aec	SLES-9-x86-64-RC5-CD6.iso

337478017119d6a76577938b3a0f0ff0	SLES-9-s390x-RC5-CD1.iso
04e7b9d8629e1d973230be67b22dc9b7	SLES-9-s390x-RC5-CD2.iso
1a5b5f0d15b4e182d2b23da478725933	SLES-9-s390x-RC5-CD3.iso
cee521cbd1ac6c2e3980da99bd4b652e	SLES-9-s390x-RC5-CD4.iso
7fd515da7f5cdc2bd916ee432da30fd4	SLES-9-s390x-RC5-CD5.iso
74301159228f91a4da4a2bb78cb607aa	SLES-9-s390x-RC5-CD6.iso

You **MUST** use **SUSE LINUX Enterprise Server 9**. Make sure that you are using the appropriate version for your platform, refer to section §1.3.1 "Hardware requirements" of this guide for the list of supported hardware and the corresponding version needed.

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 pressing the F2 key at the boot prompt, or add the option `console=tty1`.
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 SLES CD. When the boot prompt appears, add the option `console=ttyS0` (use the appropriate name of the serial device if not using `ttyS0`) and press ENTER to start the installation.
 5. Accept the **end user license** agreement.
 6. **Select your language:** choose **English (US)** to ensure that the messages shown during the installation match those described in this guide.
 7. If prompted (due to having Linux installed already), choose **New installation**.
 8. Next is the **Installation settings** dialog. Change the settings shown by clicking on the blue headings, or alternatively by choosing the corresponding items from the :

Mode

Choose **New installation**

Keyboard layout

RECOMMENDED: set to match the attached keyboard

Mouse

OPTIONAL: set to match the attached mouse. A mouse is not needed for the evaluated configuration.

Partitioning

You **MUST** use specific settings for the evaluated configuration, using ext3 file systems with ACL support and including a separate `/var/log/` partition (for CAPP-compliant auditing). Select either **Base partition setup on this proposal** or **Create custom partition setup**.

- Configuring a swap partition at least as large as the installed RAM is **RECOMMENDED**.
- 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 need a separate */boot* partition. This is usually recognized automatically by the installation program. For pSeries machines, you **MUST** create a partition of type 41 and at least 2MB in size for boot information, and you do **NOT** 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**:

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

- Set the file system type of all partitions to **ext3**, then choose **Fstab Options** and turn on the **Access Control Lists** check mark. You **MAY** activate the additional options "No access time", "Mount read-only", and "Extended User Attributes" as required.

Software

Choose **Minimum system** (or "Minimum graphical system (without KDE)" if "Minimum system" is not offered as an option), and confirm the choice. Extra packages will be removed during the following hardening steps.

Select **Detailed selection** and add the following packages to the selection. This is easiest when **Filter** is set to **Search**, then you can enter (part of) the package names in the search field and add a check mark to the package in the search result.

The packages marked as **OPTIONAL** are services that are part of the evaluated configuration but **MAY** be omitted if you do not need them for your system. Packages containing documentation files or viewers that this document refers to are marked as **RECOMMENDED**, but you **MAY** omit them.

The installer will automatically choose an appropriate kernel (single processor or SMP) based on the detected hardware. You **MAY** override this choice and choose a compatible kernel from the following list manually:

```

kernel-default
kernel-smp
kernel-ppc64
kernel-iseries64
kernel-s390x

### REQUIRED packages
laus                # The Linux Audit System
laus-64bit          # ONLY for ppc64 (pSeries, iSeries) systems
pam-laus            # Audit-enabled version of the PAM libraries
star                # Data archival tool with ACL support

### RECOMMENDED packages
texinfo             # Info documentation viewer
man-pages           # Manual pages
howtoenh            # how-to documentation (HTML format)
sles-admin_en       # Administrator Manual

### OPTIONAL packages
lprng               # Print spooler
xinetd              # Xinetd (only used for vsftpd)
vsftpd              # FTP daemon (needs xinetd)
stunnel             # set up encrypted SSL tunnels
```

Booting

keep default (no other OS is permitted on the server).

Time zone

- RECOMMENDED: keep hardware clock time as **UTC**.
- RECOMMENDED: set the time zone as appropriate for the server location.

Language

choose **English (US)** to ensure that the messages shown during the installation match those described in this guide.

- To start the installation: press the **Accept** and **Yes, install** buttons.
- Installation will proceed. Insert the CDs as prompted by the installer.
- The installer will reboot to continue running on the installed 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.
- The installer will continue in text mode, confirm the explanatory text about this.
- Password for "root", the administrator**
 - choose according to the password policy (§6.3)
 - in **Expert Options**, set Password Encryption: **MD5**
- 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.
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 when using a NFS file server instead of CD-ROMs). 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.
You **MUST NOT** install, connect or configure modems or ISDN adapters.
Use the **Change...** menu to configure the **Network interfaces**.
For each network card, select **Change...**, then **Edit** the network settings. The following options **MUST** be used for non-virtual network cards:
 - Use **Static address setup** for each card, and configure an appropriate **IP Address** and **Subnet mask**. You **MUST NOT** use DHCP.
 - Select the **Host name and name server** dialog, and make the following changes:
 - Disable the **Change host name via DHCP** check box.
 - Disable the **Update name servers via DHCP** check box.
 - RECOMMENDED: set the system's **Host name**.
 - OPTIONAL: configure **Name server** and **Domain search** entries as required.
 - In the **Routing** dialog, configure the **Default gateway** and/or static routes in the routing table as required. You **MAY** enable IP forwarding.
 - Use the **Next** button to continue.
- In the **Test Internet Connection** dialog, select **No, Skip This Test**. Use the **Next** button to continue.
- In the **Service Configuration** dialog, you **MAY** change the **CA Management** settings if needed. The **OpenLDAP Server** **MUST** be disabled.
- In the **User Authentication Method** dialog, select the **Local** authentication method.

18. In the **Add a New Local User** dialog, create an account for one of the administrators (RECOMMENDED: use the real name of the person doing the installation).
 - Fill out the **First name**, **Last name**, **User login** and **Password** fields. The password **MUST** be chosen as described in section §6.3 "Password policy" of this guide.
 - It is RECOMMENDED to activate **Receive System Mail** for administrators.
 - You **MUST NOT** activate "Auto Login".
 - Open the **Password settings** dialog, and edit the settings according to the parameters described in section §3.15 "Setting up login controls" of this guide:

Issue warning how many days before password expiration?	5
How many days after password expires is the login usable?	-1
Maximum number of days for the same password	60
Minimum number of days for the same password	1

The "Expiration date" **MAY** be left blank. Close the dialog.
 - Open the **Details** dialog, select the **Additional group** pane and add membership in the group **trusted** for this administrator. Close the dialog.
 - You **MAY** select **User Management** to create additional administrator accounts at this time, but it is RECOMMENDED to do this later, after setup of the evaluated configuration has been completed.
 - Use the **Next** button to continue.
19. Confirm the **Release Notes** dialog by selecting **Next**.
20. In the **Hardware Configuration** dialog, you **MUST NOT** enable 3D Acceleration.
21. Confirm the **Installation Completed** dialog to start the system.
22. 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.

- Install the required updates, including post-SLES9 updates.
- Either install the *certification-sles-ibm-eal4* RPM and run the *sles-eal4* script according to the instructions in section §3.2 "Automated configuration of the system" of this guide; or do the required actions manually.
- Reboot.

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 Prerequisites

3.1.1 Filesystem configuration

CD/DVD devices **MUST** be accessed using the `iso9660` filesystem type. Using the `subfs` automounter is **NOT** permitted in the evaluated configuration. See also section §4.6 "Mounting filesystems" of this guide, specifically that writable removable media are **NOT** permitted in the evaluated configuration.

Skip this section if the system does not have any removable media storage devices.

You **MUST** edit the `/etc/fstab` file and update the filesystem type and options. In each line mentioning the filesystem type `subfs`, change the filesystem type to `iso9660` and the options to `ro,nodev,nosuid,noauto`. The following example shows a sample `/etc/fstab` line:

```
/dev/cdrom /media/cdrom iso9660 ro,nodev,nosuid,noauto 0 0
```

Note that the device name is hardware dependent, and may be `/dev/dvd`, `/dev/cdrecorder` or similar. Floppy disks usually do not contain `iso9660` file systems, it is **RECOMMENDED** to delete or disable entries referring to floppy disks.

Repeat these steps for each entry if you have multiple CD/DVD/floppy drives installed.

After modifying the `/etc/fstab` file, re-mount the filesystems to activate the changes by running the following commands:

```
umount /media/*
mount -a
```

3.1.2 Getting required updates

You **MAY** make files available to other SLES systems in the secure network and use the YAST2 online update mechanism to retrieve the files from this local mirror, but you **MUST NOT** connect to the Internet from the target system at this time.

The ISO images and packages are available from the SUSE maintenance web:

```
http://portal.suse.com/
```

Access is restricted to registered users, use your user name and password to log in. Make sure that you have registered your SLES9 copy using the enclosed registration key.

Navigate to the section **Patch Support Database**, and use the **Search** function to locate the files needed. Set the search scope to **Fulltext search in PSDB**. As search terms, use either the package name (such as "amtu"), and/or the keywords "IBM EAL4" for a list of packages related to the evaluation.

Note that the `yast2-installation` package uses the architecture specification `noarch` that usually indicates a platform independent package, but the file content is different for each platform despite the identical file name. Make sure that you use the correct one as listed in the table below.

# i386 (xSeries)	
1e404439355ec3fa3f3a86b31409a9fb	aaa_base-9-29.13.i586.rpm
c42d123078ee453d1cbf088ad30eace1	amtu-0.1-1.6.i586.rpm
8162ab5c488c487ac67eccb53b4dfa96	glibc-locale-2.3.3-98.31.i586.rpm
933a87f09e9bac7b7bdbc122b0fd57e4	kernel-default-2.6.5-7.111.30.i586.rpm
44301f5fc37934f585ae83c54c1c9c4a	kernel-smp-2.6.5-7.111.30.i586.rpm


```

73b31c4f482d74e82f3c33fe9cb2286b liby2util-2.9.25-0.2.i586.rpm
35b7e69c3e1d980103ee8a9326b45949 mkinitrd-1.0-199.53.i586.rpm
d809dc2216dd3f5b950e828484b21f89 net-tools-1.60-543.6.i586.rpm
81005198981cc98ab06ffe682dfc0e8c openssl-0.9.7d-15.13.i586.rpm
190208b5bcf932f26357ef34cf63d9b1 permissions-2004.7.30-0.2.i586.rpm
c92b8ad4412433fcde888783c241e06d release-notes-9.1-8.40.i586.rpm
b47cf6b61f2a3cf7cc2f2194a89a8cc0 submount-0.9-33.6.i586.rpm
ce45be851071e95c3bec774072e90793 timezone-2.3.3-98.31.i586.rpm
4b8df352876407a90953bfef91f3cf36 yast2-installation-2.9.89-0.2.noarch.rpm
6d97c8dfea265c3ff5bd287b759748c2 yast2-network-2.9.59-0.2.i586.rpm
b1f5e00e5320db1316db898256594449 zlib-1.2.1-70.6.i586.rpm
fc2dacfc98567293f07931659b8b7553 glibc-2.3.3-98.31.i686.rpm
282fc463469682e87d8875facc6f38e3 glibc-devel-2.3.3-98.31.i686.rpm

```

```
# ppc (pSeries or iSeries)
```

```

978985ddb08d2fb39d474170e02a96ed aaa_base-9-29.13.ppc.rpm
161d6fe66a5e67adclb85e46c3cf36e7 amtu-0.1-1.6.ppc.rpm
3935a5c4ca97523fdbc3507f529b1538 glibc-2.3.3-98.31.ppc.rpm
f30d37efdddf4d5b5f3178c454bdc578 glibc-devel-2.3.3-98.31.ppc.rpm
d6dc2358403a2b884414eedfd955caa9 glibc-locale-2.3.3-98.31.ppc.rpm
39df6febb6e572dcd39dc5e0d324c0cc kernel-pseries64-2.6.5-7.111.30.ppc.rpm
b446846fe66150d04cfb8d7bcfd0e0e6 liby2util-2.9.25-0.2.ppc.rpm
d178667fb7708d11c78a9eb215ded73e mkinitrd-1.0-199.53.ppc.rpm
12634d3f5e11b822cb38d53f03ce33cb net-tools-1.60-543.6.ppc.rpm
85afdea4c5fdea9672373a5199181e16 openssl-0.9.7d-15.13.ppc.rpm
4b5e9579e4e0dd6af90d6b39dbb7b9cd permissions-2004.7.30-0.2.ppc.rpm
f63cd55fe244b2758355b09313a91f4d release-notes-9.1-8.40.ppc.rpm
712f96fc47762f5621f15463091f7276 submount-0.9-33.6.ppc.rpm
cblacc1d2a10776233e5589484adf11d timezone-2.3.3-98.31.ppc.rpm
7a264da90a234507ab476d4c3d8a65f3 yast2-installation-2.9.89-0.2.noarch.rpm
0dbaaac7c22b63b27b6a3bc446826410 yast2-network-2.9.59-0.2.ppc.rpm
d1e11182ff5d3312e7cca6b488e0976b zlib-1.2.1-70.6.ppc.rpm

```

```
# s390x (zSeries)
```

```

a14b96a90018025aad127084c7a78a0e aaa_base-9-29.13.s390x.rpm
9b6fbc0c12f582a0cd048387e27a8100 amtu-0.1-1.6.s390x.rpm
1cae6398518595f0d10a9c6462bafcb2 glibc-2.3.3-98.31.s390x.rpm
35364600ca5905fb04ea29949b0a4281 glibc-devel-2.3.3-98.31.s390x.rpm
a454a0ad0c960b92a009152c006db11d glibc-locale-2.3.3-98.31.s390x.rpm
aef2fd9a36f260f0391dba014b9c803e kernel-s390x-2.6.5-7.111.30.s390x.rpm
7e31e85f4ada7b3063b41361a19f761f liby2util-2.9.25-0.2.s390x.rpm
caa677ddf98093d19c24e601d6885596 mkinitrd-1.0-199.53.s390x.rpm
cb69a31c593a3d2bfaf79d0f2d4225e7 net-tools-1.60-543.6.s390x.rpm
fbc0980ebaaf460c56c34429e9616656 openssl-0.9.7d-15.13.s390x.rpm
df671cdee08f9e4ec68a6f96f0cb16ca permissions-2004.7.30-0.2.s390x.rpm
302c0eda4a2d3597e03c3ea9aee58dc7 release-notes-9.1-8.40.s390x.rpm
0310000e93fe2b855a121a27c307307c submount-0.9-33.6.s390x.rpm
359e83353992994a2e4383dfaca3607a timezone-2.3.3-98.31.s390x.rpm
76d4ec70c1c412e2a460f547a2c82f4d yast2-installation-2.9.89-0.2.noarch.rpm
709c541b5e42113fb94441aa99cel1183 yast2-network-2.9.59-0.2.s390x.rpm
518da232b7bb612f006cbb36987026a6 zlib-1.2.1-70.6.s390x.rpm

```

```
# x86_64 (eServer 325 Opteron)
```

```

6236dab482a758371d670ae8386d887d  aaa_base-9-29.13.x86_64.rpm
48ea711088732f1d52991f37a6eea69a  amtu-0.1-1.6.x86_64.rpm
5ad0c82db8b431bd978813668ff86955  glibc-2.3.3-98.31.x86_64.rpm
69f562d6c703a7659d44e4cb0a9c6d91  glibc-devel-2.3.3-98.31.x86_64.rpm
578edfb6a863effdc058f7f06a955cad  glibc-locale-2.3.3-98.31.x86_64.rpm
8b4fe7985786def7caee0c61c98ffd26  kernel-default-2.6.5-7.111.30.x86_64.rpm
60611e610708616e53f8a63ea62fb9db  kernel-smp-2.6.5-7.111.30.x86_64.rpm
b40f1fd909a9f229ecf0946919b0b9e7  liby2util-2.9.25-0.2.x86_64.rpm
84635bc562acd8111ae5d810c757189b  mkinitrd-1.0-199.53.x86_64.rpm
8460b9aefa8a191eceb963eff2e89315  net-tools-1.60-543.6.x86_64.rpm
dfb4d48ac27cd320815c18b29fb95025  openssl-0.9.7d-15.13.x86_64.rpm
1bac0b65c7430d6805b4ceeaf8f27096  permissions-2004.7.30-0.2.x86_64.rpm
517c33b6f2a6f10525c7aa2ae976124d  release-notes-9.1-8.40.x86_64.rpm
905b0c0781cd95aa6e65f1b1a7563785  submount-0.9-33.6.x86_64.rpm
89acb3aceccc81e5f3182c628fab3574  timezone-2.3.3-98.31.x86_64.rpm
42201f9664e09e16c819de29251030fb  yast2-installation-2.9.89-0.2.noarch.rpm
07e5d54cde527cb2d3af2cb96aelc550  yast2-network-2.9.59-0.2.x86_64.rpm
ab54e347e562bf3c3a5026d9aa62f9f4  zlib-1.2.1-70.6.x86_64.rpm

```

The `glibc-devel` packages are **OPTIONAL**, but if you do install the development library you **MUST** use the updated version.

You **MUST** use exactly one of the kernels from the following list for the evaluated configuration, as appropriate for your hardware:

```

kernel-default
kernel-smp
kernel-pseries64
kernel-s390x

```

All other packages (other than `glibc-devel` and the kernel) are **REQUIRED**.

Note that both the iSeries and pSeries systems use the kernel named `kernel-pseries64` in the evaluated configuration.

You **MUST** verify the MD5 sums against those shown in the list. Run the following command to display those of the downloaded files:

```
md5sum *.rpm
```

Download the CD-ROM images and RPMs using an Internet-connected computer, and transfer them to the system being installed, for example using a CD-R disk.

Do **NOT** install the downloaded packages yet.

It is **RECOMMENDED** that you store the RPM packages in a separate directory (this guide uses `/root/rpm/`) to avoid confusion with similarly named packages from installation CDs.

3.2 Automated configuration of the system

The `certification-sles-ibm-eal4` package **SHOULD** be installed initially to achieve the evaluated configuration. This RPM package contains EAL4 specific configuration files, updates to the online manuals, and scripts that set up the evaluated configuration.

Install the RPM as follows:

```
rpm -Uvh /root/rpm/certification-sles-ibm-eal4*.noarch.rpm
```

Please check the file `/usr/share/doc/packages/certification-sles-ibm-eal4/README-eal4.txt` from the *certification-sles-ibm-eal4.rpm* for the latest errata information.

The automated installation depends on having the correct versions available for those packages that **MUST** be updated or added to the evaluated configuration.

You **MUST** use the versions of the packages downloaded from the maintenance web exactly as listed in section §3.1.2 "Getting required updates" of this guide. Be careful to pay close attention that the version number and architecture name matches exactly for each package.

Copy (or move) the update RPMs from your download directory to the location expected by the script. You **MUST** include all **REQUIRED** packages and **MAY** add the **OPTIONAL** packages as well. Use the following command:

```
cp /root/rpm/*.rpm /usr/lib/eal4/rpm/
```

Verify that the directory now contains all packages listed as **REQUIRED** in section §3.1.2 "Getting required updates", and that they match the architecture you are installing on. Use the following command to show the directory contents:

```
ls -l /usr/lib/eal4/rpm/
```

You **MAY** copy multiple kernels into the directory, the script will install only the appropriate one.

The *certification-sles-ibm-eal4.rpm* package contains a setup script that implements the evaluated configuration when run. You **MAY** add the `-a` switch to run the script automatically, but be aware that this will change the configuration with no prompting. Run it with no arguments to use the default interactive mode (with prompts for confirmation before making changes):

```
/usr/lib/eal4/bin/sles-eal4
```

When running the script in interactive mode, you **MUST** permit it to make each change unless the step is clearly documented to be **OPTIONAL**.

It is **RECOMMENDED** that you use the `sles-eal4` script to configure the system, but you **MAY** also perform the steps listed in sections §3.3 to §3.17 of this guide manually instead.

If the script fails with an error message, verify that all the steps listed in section §3.1 "Prerequisites" of this guide have been followed. If the message indicates that the PAM library is not audit enabled, please follow the instructions in section §3.12.4 "Ensuring that the PAM library is audit enabled" of this guide to reinstall the library. Then re-run the script.

The *certification-sles-ibm-eal4* RPM contains the following EAL4 specific configuration files:

```
/etc/permissions.eal4
```

We **RECOMMEND** that you also use the `sles-eal4` script to reset the configuration to its initial state after any updates, but you **MAY** also do this manually.

WARNING: The `sles-eal4` script will reboot the system as the final step in the process, as described in the manual instructions in section §3.17 "Reboot and initial network connection". Remember to remove any CD-ROM from the drive and/or configure the system to boot from hard disk only.

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.

3.3 Add and remove packages

The minimal install still contains some packages that **MUST** be removed for the evaluated configuration. Use `rpm -q` to get a list of installed packages, and `rpm -e PACKAGE_NAME . . .` to remove all packages EXCEPT those listed here.

Some packages are listed as RECOMMENDED or OPTIONAL in section §2.2 "Selection of install options and packages". If you did not select all of those, some of the following packages will not be present on your system.

The evaluated configuration including all RECOMMENDED and OPTIONAL packages consists of the following packages:

all architectures:

aaa_base	netcfg
aaa_skel	openldap2-client
acl	openslp
amtu	openssh
ash	openssl
at	pam-laui
attr	pam-modules
autoyast2-installation	parted
bash	pciutils
bc	pcre
bzip2	perl
certification-sles-ibm-eal4	perl-Config-Crontab
core-release	permissions
coreutils	popt
cpio	postfix
cracklib	procps
cron	psmisc
curl	pwdutils
cyrus-sasl	readline
db	release-notes
device-mapper	resmgr
devs	rpm
dialog	scsi
diffutils	sed
dosfstools	sitar
e2fsprogs	sles-admin_en
ed	sles-release
evms	star
file	stunnel
filesystem	submount
fillup	suse-build-key
findutils	sysconfig
gawk	syslogd
gdbm	sysvinit
glibc	tar
glibc-locale	tcpd
gpg	telnet
gpm	terminfo
grep	texinfo
groff	timezone
gzip	udev

hdparm	usbutils
heimdal-lib	utempter
hfsutils	util-linux
hotplug	vim
howtoenh	vsftpd
hwinfo	w3m
info	wget
insserv	xinetd
iproute2	yast2
iputils	yast2-bootloader
kbd	yast2-core
ksymoops	yast2-country
laus	yast2-inetd
ldapcpplib	yast2-installation
less	yast2-ldap
libacl	yast2-ldap-client
libattr	yast2-mail-aliases
libgcc	yast2-mouse
libselinux	yast2-ncurses
libstdc++	yast2-network
libxcrypt	yast2-online-update
libxml2	yast2-packagemanager
liby2util	yast2-packager
logrotate	yast2-pam
lprng	yast2-perl-bindings
lsof	yast2-runlevel
lukemftp	yast2-security
m4	yast2-storage
mailx	yast2-sysconfig
man	yast2-theme-SuSELinux
man-pages	yast2-trans-en_US
mingetty	yast2-transfer
mkinitrd	yast2-update
mktemp	yast2-users
module-init-tools	yast2-x11
ncurses	yast2-xml
net-tools	zlib
netcat	

additional on x86:

either the "kernel-default" or the "kernel-smp" kernel

grub
isapnp
lilo

additional on x86_64 (Opteron):

either the "kernel-default" or the "kernel-smp" kernel

bzip2-32bit	heimdal-lib-32bit	ncurses-32bit
cracklib-32bit	irqbalance	openldap2-client-32bit
cyrus-sasl-32bit	isapnp	openssl-32bit
db-32bit	laus-32bit	pcre-32bit

e2fsprogs-32bit	libacl-32bit	perl-32bit
file-32bit	libattr-32bit	popt-32bit
gdbm-32bit	libselinux-32bit	readline-32bit
glibc-32bit	libxcrypt-32bit	resmgr-32bit
glibc-locale-32bit	libxml2-32bit	utempter-32bit
grub	lilo	zlib-32bit

additional on ppc (pSeries):

bzip2-64bit	libacl-64bit	pcre-64bit
cracklib-64bit	libattr-64bit	pdisk
cyrus-sasl-64bit	libgcc-64bit	perl-64bit
db-64bit	libselinux-64bit	popt-64bit
e2fsprogs-64bit	libstdc++-64bit	ppc64-utils
file-64bit	libxcrypt-64bit	readline-64bit
gdbm-64bit	libxml2-64bit	resmgr-64bit
glibc-64bit	lilo	utempter-64bit
heimdal-lib-64bit	ncurses-64bit	zlib-64bit
kernel-pseries64	openldap2-client-64bit	
laus-64bit	openssl-64bit	

additional on ppc (iSeries):

baselibs-64bit	kernel-pseries64	openldap2-client-64bit
bzip2-64bit	laus-64bit	openssl-64bit
cracklib-64bit	libacl-64bit	pcre-64bit
cyrus-sasl-64bit	libattr-64bit	pdisk
db-64bit	libgcc-64bit	perl-64bit
e2fsprogs-64bit	libselinux-64bit	popt-64bit
file-64bit	libstdc++-64bit	ppc64-utils
gdbm-64bit	libxcrypt-64bit	readline-64bit
glibc-64bit	libxml2-64bit	resmgr-64bit
glibc-locale-64bit	lilo	utempter-64bit
heimdal-lib-64bit	ncurses-64bit	zlib-64bit

additional on s390 (zSeries):

bzip2-32bit	kernel-s390x	openssl-32bit
cracklib-32bit	laus-32bit	pcre-32bit
cyrus-sasl-32bit	libacl-32bit	perl-32bit
db-32bit	libattr-32bit	popt-32bit
e2fsprogs-32bit	libgcc-32bit	readline-32bit
file-32bit	libselinux-32bit	resmgr-32bit
gdbm-32bit	libxcrypt-32bit	s390-tools
glibc-32bit	libxml2-32bit	utempter-32bit
glibc-locale-32bit	ncurses-32bit	zlib-32bit
heimdal-lib-32bit	openldap2-client-32bit	

The *pam* package will be listed in the RPM database as being installed, but all of its files were overwritten by the *pam-laas* package. You MUST NOT try to uninstall, reinstall or update the *pam* package.

In addition to these packages, certain additional software from the SLES CDs MAY be installed without invalidating the evaluated configuration. The rules described in the section §4.4 "Installation of additional software" 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 `sles-ea14` script does not remove these packages if they are installed on the system:

<code>autoconf</code>	<code>libattr-devel</code>	<code>perl-HTML-Tagset</code>
<code>automake</code>	<code>libgcc</code>	<code>perl-IO-Stty</code>
<code>binutils</code>	<code>libstdc++-devel</code>	<code>perl-IO-Tty</code>
<code>binutils-32bit</code>	<code>libstdc++-devel-32bit</code>	<code>perl-Mon</code>
<code>binutils-64bit</code>	<code>libstdc++-devel-64bit</code>	<code>perl-Net-SNMP</code>
<code>cpp</code>	<code>make</code>	<code>perl-Net_SSLeay</code>
<code>cvs</code>	<code>ncurses-devel</code>	<code>perl-Tie-IxHash</code>
<code>expect</code>	<code>ncurses-devel-32bit</code>	<code>perl-Time-Period</code>
<code>flex</code>	<code>ncurses-devel-64bit</code>	<code>perl-TimeDate</code>
<code>flex-32bit</code>	<code>openssl-devel</code>	<code>perl-Tk</code>
<code>flex-64bit</code>	<code>openssl-devel-32bit</code>	<code>perl-URI</code>
<code>gcc</code>	<code>openssl-devel-64bit</code>	<code>perl-gettext</code>
<code>gcc-64bit</code>	<code>pam</code>	<code>perl-libwww-perl</code>
<code>gcc-c++</code>	<code>pam-devel</code>	<code>strace</code>
<code>glib</code>	<code>pam-laui-devel</code>	<code>strace-32bit</code>
<code>glibc-devel</code>	<code>patch</code>	<code>strace-64bit</code>
<code>glibc-devel-32bit</code>	<code>perl-Convert-BER</code>	<code>tcl</code>
<code>glibc-devel-64bit</code>	<code>perl-Crypt-DES</code>	<code>tcl-32bit</code>
<code>kernel-source</code>	<code>perl-DateManip</code>	<code>tcl-64bit</code>
<code>laus-devel</code>	<code>perl-Digest-HMAC</code>	<code>tcl-devel</code>
<code>laus-devel-32bit</code>	<code>perl-Digest-SHA1</code>	<code>tcsh</code>
<code>laus-devel-64bit</code>	<code>perl-Expect</code>	<code>tk</code>
<code>laus-devel-64bit</code>	<code>perl-HTML-Parser</code>	

3.4 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
audit
coldplug
cron
hwscan
network
random
rpmconfigcheck
syslog
```

The following services are OPTIONAL for runlevel 3:

```
hotplug
kbd
lpd
postfix
sshd
xinetd
```

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 *inserv ServiceName* to activate a service, and *inserv -r ServiceName* to deactivate it.

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. If it is missing, create the link with *inserv audit*.

3.4.1 Disable usbfs

The *usbfs* file system is not permitted in the evaluated configuration and **MUST** be disabled. Note that the only permitted USB devices are keyboards and mice connected at boot, and these also work without *usbfs* for the supported hardware. Please refer to sections §2.1 "Supported hardware" and §4.6 "Mounting filesystems" of this guide for more information.

usbfs is activated in three different places in two startup files, use the following command to verify the current content:

```
grep -n 'mount.*usb' /etc/init.d/kbd /etc/hotplug/usb.rc
```

Here is the output of the *grep* command before modification:

```
/etc/init.d/kbd:88:
    mount -n -t usbfs usbfs /proc/bus/usb >/dev/null 2>&1
/etc/hotplug/usb.rc:181:
    mount /proc/bus/usb
/etc/hotplug/usb.rc:185:
    mount -t usbfs usbfs /proc/bus/usb
/etc/hotplug/usb.rc:296:
    umount /proc/bus/usb >/dev/null 2>&1
```

Either use a text editor to edit the files manually (replacing each instance of *mount* (but not *umount*) with the string *: mount*), or use the following automated method:

```
perl -pi.bak -e 'next if /umount/; s/mount.*usb/: $&/' \
    /etc/init.d/kbd /etc/hotplug/usb.rc
```

Make sure that you use the "colon space" sequence *: '* to disable the commands and not the hashmark *'#'* comment character, you will get syntax errors if only comment lines are left inside a conditional branch. You do not need to change the *umount* command.

After the modification, the file content **MUST** be as follows, verify by re-running the *grep* command:

```
/etc/init.d/kbd:88:
    : mount -n -t usbfs usbfs /proc/bus/usb >/dev/null 2>&1
/etc/hotplug/usb.rc:181:
    : mount /proc/bus/usb
/etc/hotplug/usb.rc:185:
    : mount -t usbfs usbfs /proc/bus/usb
/etc/hotplug/usb.rc:296:
    umount /proc/bus/usb >/dev/null 2>&1
```

You **MUST NOT** manually mount the *usbfs* file system.

3.5 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/lpq
/usr/bin/lpr
/usr/bin/lprm
/usr/bin/lpstat
/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.

The SLES permission mechanism MUST be used to set permission bits appropriately. First make sure that no SUID/SGID programs are present on the system:

```
find / \( ! -fstype ext3 -prune -false \) -o \
  -type f \( -perm -4000 -o -perm -2000 \) \
  -exec chmod u-s,g-s {} \; -print
```

Make sure that `/etc/sysconfig/security` has the following two variables set:

```
CHECK_PERMISSIONS=set
PERMISSION_SECURITY="eal4"
```

Then run `chkstat -set /etc/permissions.eal4` to set the needed SUID and SGID bits.

3.6 Update permissions for 'su'

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

```
chgrp trusted /bin/su
chmod 4750 /bin/su
```

When running the `chkstat` command as described above, this will be configured automatically.

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

3.7 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/bash.bashrc.local` 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:

```
#
# This file contains the device names of tty lines (one per line,
# without leading /dev/) on which root is allowed to login.
#
tty1
tty2
tty3
tty4
tty5
tty6
# for devfs:
vc/1
vc/2
vc/3
vc/4
vc/5
vc/6
```

3.8 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`:

```

# Cryptographic settings. Disallow the obsolete (and
# insecure) protocol version 1, and hardcode a strong
# cipher.
Protocol 2
Ciphers 3des-cbc

# Configure password-based login. This MUST use the PAM
# library exclusively, and turn off the builtin password
# authentication code.
UsePAM yes
ChallengeResponseAuthentication yes
PasswordAuthentication no
PermitRootLogin no
PermitEmptyPasswords no

# No other authentication methods allowed
IgnoreRhosts yes
RhostsRSAAuthentication no
HostbasedAuthentication no
PubkeyAuthentication no
RSAAuthentication no
KerberosAuthentication no
GSSAPIAuthentication no

# Other settings, MAY change "X11Forwarding" to "yes"
X11Forwarding no
Subsystem sftp /usr/lib/ssh/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.

3.9 Setting up xinetd

The *xinetd* super server is used to start the FTP daemon. The defaults entry in the */etc/xinetd.conf* file specifies the log file and the data that is to be logged:

```

defaults
{
    log_type           = FILE /var/log/xinetd.log
    log_on_success     = PID HOST EXIT DURATION
    log_on_failure     = HOST ATTEMPT
    instances          = 2
}

```

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

3.10 Setting up FTP

The evaluated configuration includes 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.

The FTP server is started via *xinetd*, see *xinetd*(8). The following entry is the only active configuration entry in */etc/xinetd.conf*:

```
service ftp
{
    socket_type      = stream
    protocol         = tcp
    wait             = no
    user             = root
    server           = /usr/sbin/vsftpd
    instances        = UNLIMITED
}
```

The *vsftpd* service uses several additional configuration files. In */etc/vsftpd.conf* the configuration of the ftp daemon is specified. In addition, the file */etc/ftpusers* 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 *ftpusers* file is not checked by the ftp daemon itself but by a PAM module. Please see section §3.14 "Required Pluggable Authentication Module (PAM) configuration" of this guide for details.

The setup of */etc/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 therefore only suitable for distribution of public files for which no read access control is needed.

```
anonymous_enable=YES
local_enable=NO
```

It is RECOMMENDED disabling anonymous FTP if you do not need this functionality with the following */etc/vsftpd.conf* setting:

```
anonymous_enable=NO
```

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

3.11 Setting up Postfix

The default settings of the postfix MTA are in accordance with the EAL4 requirements. It is RECOMMENDED that you set up an alias for root in the */etc/aliases* file. Specify one or more user names of administrators to whom mail addressed to *root* will be forwarded.

For example, run the following commands (assuming you are starting from the default Postfix configuration) to forward root mail to user "jdoe":

```
echo "root: jdoe" >>/etc/aliases
newaliases
postfix reload
```

Please see *postfix*(1), *master*(8), *aliases*(5), *newaliases*(1), and the documentation in */usr/share/doc/packages/postfix/html/* for details.

3.12 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.

This section describes the further changes that MUST be made to reach the initial state of the evaluated configuration.

3.12.1 Setting up the audit configuration files

For all platforms, it is RECOMMENDED to use the following settings in the */etc/sysconfig/audit* file:

```
AUDIT_ALLOW_SUSPEND=1
AUDIT_ATTACH_ALL=0
AUDIT_MAX_MESSAGES=1024
AUDIT_PARANOIA=0
```

The *sles-ea14* 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
```

3.12.2 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:

```
insserv audit
```

3.12.3 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:

```
insserv -r audit
echo "au:35:respawn:/etc/init.d/audit inittab" >> /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.

3.12.4 Ensuring that the PAM library is audit enabled

The installation media contain two copies of the PAM library, the `pam` package contains the default version without audit capability, and the `pam-laui` package the audit-capable one. The set of files included in these packages is the same, but the file content differs.

In the initial installation, the package `pam-laui` was added to the list of packages to install. Due to installation order side effects, the wrong version of the PAM library may be active after the initial install.

You MUST verify that the active PAM library is audit enabled, it MUST link to the `liblaui.so` shared library. Use the following command to check:

```
grep laui_open `ldd /bin/login | awk '/libpam.so/ { print $3 }`
```

The expected output is the following line (the initial path is `lib64/` instead of `lib/` on some platforms):

```
Binary file /lib/libpam.so.0 matches
```

If the `grep` command produces no output, you MUST reinstall the `pam-laui` package from CD #2 of the installation media. Use the following command (using your system's architecture instead of `i586` as appropriate) to reinstall it:

```
# 'cd' to the directory containing the RPM file,
# then reinstall the package:
rpm --oldpackage --force --nodeps -Uhv pam-laui-0.77-4.3.i586.rpm
```

You MUST NOT reinstall the non-audit-enabled `pam` library.

3.13 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/packages/pam/text/`, 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/packages/pam/modules/README.pam_*`, as well as manual pages for individual modules, such as `pam_pwcheck(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 SLES (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. SLES 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.

3.14 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.

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 'trusted' 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 five failed attempts, corresponding to the setting *deny=5*, 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.

You MUST NOT modify other settings, specifically you MUST use the 'md5' and 'use_cracklib' options for the *pam_pwcheck.so* module.

The 'remember=XX' option must be added to the */etc/security/pam_pwcheck.conf* file to force users to create new passwords and not re-use ones that they had previously, i.e. to prevent users from simply alternating between two passwords when asked to change it due to expiration. XX is any number between 7 and 400.

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).

Following are the pam configuration files:

3.14.1 /etc/pam.d/chage

This file configures the access control for the *chage* command. It allows the use of *chage* only after the user's password has been entered or the calling user is 'root'.

```

#%PAM-1.0
# root is allowed to use chage without authentication
auth      sufficient    pam_rootok.so
auth      required      pam_unix2.so
account    required      pam_permit.so
password   required      pam_deny.so
session    required      pam_deny.so

```

3.14.2 /etc/pam.d/chfn

This file configures the access control for the *chfn* command. It allows the use of *chfn* only after the user's password has been entered or the calling user is 'root'.

```

#%PAM-1.0
auth      sufficient    pam_rootok.so
auth      required      pam_unix2.so

```



```

account    required    pam_unix2.so
password   required    pam_deny.so
session    required    pam_deny.so

```

3.14.3 /etc/pam.d/chsh

This file configures the access control for the *chsh* command. It allows the use of *chsh* only after the user's password has been entered or the calling user is 'root'.

```

#%PAM-1.0
auth      sufficient    pam_rootok.so
auth      required      pam_unix2.so
account    required      pam_unix2.so
password   required      pam_deny.so
session    required      pam_deny.so

```

3.14.4 /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 optional *pam.env* module MAY be used to set environment variables from */etc/security/pam.env.conf*. The optional *pam.mail* module MAY be used to notify the user that there is new mail. The *pam.tally* module MUST be used to block the user after 5 failed login attempts. The optional *pam.limits* module MAY be used to enforce resource limits via */etc/security/limits.conf*.

The *pam.laus.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.laus.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
# If serial terminals are in use, pam_laus.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_tally.so onerr=fail no_magic_root
auth      requisite      pam_unix2.so
auth      required      pam_securetty.so
auth      required      pam_nologin.so
auth      required      pam_env.so          # optional
auth      required      pam_mail.so         # optional
account    required      pam_unix2.so
account    required      pam_tally.so deny=5 reset no_magic_root
password   requisite      pam_passwdqc.so ask_oldauthtok=update check_oldauthtok
password   requisite      pam_pwcheck.so   use_first_pass use_authtok
password   required      pam_unix2.so     use_first_pass use_authtok
session    required      pam_unix2.so
session    required      pam_limits.so    # optional
session    optional      pam_laus.so      # no lockout on failure

```

3.14.5 /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
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

```

3.14.6 /etc/pam.d/passwd

This service configuration applies to password changes. Please see also */etc/security/pam.pwcheck.conf*.

```

#%PAM-1.0
auth      required      pam_unix2.so
account   required      pam_unix2.so
password  requisite      pam_passwdqc.so ask_oldauthtok=update check_oldauthtok
password  requisite      pam_pwcheck.so use_first_pass use_authtok
password  required      pam_unix2.so use_first_pass use_authtok
session   required      pam_unix2.so

```

3.14.7 /etc/pam.d/sshd

This file configures the PAM usage for SSH.

```

#%PAM-1.0
auth      required      pam_securetty.so # deny root login in evaluated config
auth      required      pam_tally.so onerr=fail no_magic_root
auth      required      pam_unix2.so
auth      required      pam_nologin.so
auth      required      pam_env.so      # optional
account   required      pam_unix2.so
account   required      pam_nologin.so
account   required      pam_tally.so deny=5 reset no_magic_root
password  requisite      pam_passwdqc.so ask_oldauthtok=update check_oldauthtok
password  requisite      pam_pwcheck.so use_first_pass use_authtok
password  required      pam_unix2.so use_first_pass use_authtok
session   required      pam_unix2.so
session   required      pam_limits.so  # optional
session   required      pam_laui.so detach

```

3.14.8 /etc/pam.d/su

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

```

#%PAM-1.0
auth      sufficient      pam_rootok.so
auth      required        pam_wheel.so use_uid group=trusted
auth      required        pam_unix2.so
auth      required        pam_tally.so onerr=fail no_magic_root
account   required        pam_unix2.so
account   required        pam_tally.so no_magic_root deny=5 reset
password  required        pam_deny.so
session   required        pam_unix2.so

```

Forcing the root user to change the root password is not desired here, therefore the *pam_unix2.so* module is absent in the *password* branch and *pam_deny.so* is used instead.

3.14.9 /etc/pam.d/useradd

This file allows the root user to add accounts without entering the root password.

```

#%PAM-1.0
auth      sufficient      pam_rootok.so
auth      required        pam_deny.so
account   required        pam_permit.so
password  required        pam_permit.so
session   required        pam_deny.so

```

3.14.10 /etc/pam.d/vsftpd

This file configures the authentication for the FTP daemon. With the *listfile* module, users listed in */etc/ftpusers* are denied FTP access to the system.

```

#%PAM-1.0
auth      required        pam_tally.so onerr=fail no_magic_root
auth      required        pam_listfile.so item=user sense=deny \
                        file=/etc/ftpusers onerr=fail
auth      required        pam_unix2.so
account   required        pam_unix2.so
account   required        pam_tally.so deny=5 reset no_magic_root
account   required        pam_laas.so detach
password  required        pam_deny.so
session   required        pam_unix2.so

```

Note that the FTP protocol has no provisions for changing passwords, therefore the *pam_unix2.so* module is absent in the *password* branch and *pam_deny.so* is used instead.

3.14.11 /etc/security/pam_pwcheck.conf

This file contains the default option for the *pam_pwcheck* module. This makes it easier to set a global policy. The *md5* option enables long passwords (up to 127 characters, see also the limit in */etc/login.defs*, and the *use_cracklib* option activates password quality checks against standard dictionary and permutation attacks. The *remember* option ensures that the user does not reuse passwords by keeping track of the specified number of previously used passwords in the file */etc/security/opasswd*.

```
password:    md5 use_cracklib remember=7
```

3.14.12 /etc/security/pam_unix2.conf

This file contains the default option for the *pam_unix2* module. This makes it easier to set a global policy. The *md5* option enables long passwords (up to 127 characters, see also the limit in */etc/login.defs*). The *trace* option activates session tracing (start/stop) via *syslog*.

```
auth:
account:
password:    md5
session:     trace
```

3.15 Setting up login controls

The system supports various options to control log ins in */etc/login.defs*. The following table explains the options and values that **MUST** be set for the EAL4 evaluated configuration. Note that these are not the complete file contents, other settings not listed here **MAY** be changed.

```
# Required settings for the Common Criteria CAPP/EAL4+ evaluated configuration
# are indicated with a '# CC: ' comment.

# Should login be allowed if we can't cd to the home directory?
# Default is yes.
#
# CC: MUST be 'no'
DEFAULT_HOME                no

# Delay in seconds before being allowed another attempt
# after a login failure
#
# CC: MUST be at least 3
FAIL_DELAY                  3

# Enable logging and display of /var/log/faillog login
# failure info.
#
# CC: MUST be set to 'no'
# CC: (this is handled by the PAM config (pam_tally) in /etc/pam.d/login)
FAILLOG_ENAB                no

# Enable logging and display of /var/log/lastlog login time info.
#
# CC: MUST be 'yes'
LASTLOG_ENAB                yes

# Enable display of unknown usernames when login failures
# are recorded.
#
# CC: MUST be 'no'
LOG_UNKFAIL_ENAB            no
```

```

# Max number of login retries if password is bad
#
# CC: MUST be <= 3
LOGIN_RETRIES          3

# Max time in seconds for login
#
# CC: MUST be <= 60
LOGIN_TIMEOUT          60

# Require password before chfn/chsh can make any changes.
#
# CC: MUST be 'yes'
CHFN_AUTH              yes

# Which fields may be changed by regular users using chfn - use
# any combination of letters "frwh" (full name, room number, work
# phone, home phone).  If not defined, no changes are allowed.
# For backward compatibility, "yes" = "rwh" and "no" = "frwh".
#
# CC: MUST NOT contain 'f'
# CC: MUST NOT be set to 'no'
CHFN_RESTRICT          rwh

# Password aging controls (used by useradd):
#
#      PASS_MAX_DAYS    Maximum number of days a password may be used.
#      PASS_MIN_DAYS    Minimum number of days allowed between password changes.
#      PASS_WARN_AGE    Number of days warning given before a password expires.
#
# CC: The settings MAY be modified but the following conditions MUST be met:
# CC:  PASS_MAX_DAYS <= 60
# CC:  PASS_MIN_DAYS >= 1
# CC:  PASS_MIN_DAYS < PASS_MAX_DAYS
# CC:  PASS_WARN_AGE <= 7
# CC:  PASS_MIN_LEN >= 8
PASS_MAX_DAYS  60
PASS_MIN_DAYS  1
PASS_WARN_AGE  7
PASS_MIN_LEN   8

# Umask which is used by useradd and newusers for creating
# new home directories.
#
# CC: MUST be '077'
UMASK          077

```

The UMASK entry sets the *default* permissions for new home directories to the most restrictive setting. Users MAY assign different permissions as described in section §6.4 "Access control for files and directories" of this guide. Note that the default umask for logged-in users is set in the */etc/profile* file, not here.

3.15.1 Maintaining *cracklib* dictionaries

The dictionary files used by *cracklib* are stored in */usr/lib/*:

```
/usr/lib/cracklib_dict.hwm  
/usr/lib/cracklib_dict.pwd  
/usr/lib/cracklib_dict.pwi
```

To create custom dictionary files instead of the supplied ones, the command */usr/sbin/create-cracklib-dict* MAY be used as follows:

```
/usr/sbin/create-cracklib-dict wordlist wordlist ...
```

This will generate a new set of dictionary files from the supplied word lists. Suggested word lists are included in the source RPM package of *cracklib*. We RECOMMEND adding dictionaries for your local language and other languages likely to be known by your user community.

3.16 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 SLES, 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.

3.16.1 GRUB boot loader configuration

The GRUB boot loader is used on the x86 and 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 SLES 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:

```
color white/blue black/light-gray
default 0
timeout 8
password --md5 $1$04711/$H/JW2MYeugX6Y1h3v.1Iz0

title linux
    kernel (hd0,1)/boot/vmlinuz root=/dev/sda2
    initrd (hd0,1)/boot/initrd
title failsafe
    lock
    kernel (hd0,1)/boot/vmlinuz.shipped root=/dev/sda2 ide=nodma apm=off \
        acpi=off vga=normal nosmp disableapic maxcpus=0 3
    initrd (hd0,1)/boot/initrd.shipped
```

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.

3.16.2 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.

The configuration is contained in the */etc/lilo.conf* file. Running the *lilo* tool creates the *yaboot.conf* file based on the information in the */etc/lilo.conf* file.

You need to re-run the *lilo(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 "SuSE Linux Enterprise Server Installation Guide" for iSeries and pSeries (pg. 37), the *yaboot.conf(5)* and *lilo(8)* manual pages, and the yaboot HOWTO for more information:

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

3.16.3 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:

```
# Generated by YaST2
[defaultboot]
default=ipl

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

3.16.4 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 "vmlintrd" 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/vmlintrd 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

and the "SuSE Linux Enterprise Server Installation Guide" for iSeries and pSeries.

3.17 Reboot and initial network connection

– *This concludes the sections covered by the automated configuration script –*

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.

Remember to remove any CD-ROM from the drive and/or configure the system to boot from hard disk only.

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 SLES 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 systems using the *grub* boot loader you can use the following 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 SLES 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/atjobs/
```

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/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
/var/spool/cron/allow
/var/spool/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 SLES 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 all non-system users 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 /var/spool/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.

The special-purpose *proc*, *sysfs*, and *tmpfs* filesystems are part of the evaluated configuration. These are virtual filesystems with no underlying physical storage, and represent data structures in kernel memory. Access to contents in these special filesystems is protected by the normal discretionary access control policy and additional permission checks.

Note that changing ownership or permissions of virtual files and directories is generally **NOT** supported for the *proc* and *sysfs* filesystems (corresponding to directories */proc/* and */sys/*), and attempts to do so will be ignored or result in error messages.

Note that use of the *usbfs* filesystem type is **NOT** permitted (and not needed) in the evaluated configuration, please refer to section §3.4.1 "Disable usbfs" of this guide for more information.

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 or DVDs.

CD/DVD devices **MUST** be accessed using the *iso9660* filesystem type. Using the *subfs* automounter is **NOT** permitted in the evaluated configuration. See also section §3.1.1 "Filesystem configuration" of this guide.

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 /media/cdrom -t iso9660 -o ro,nodev,nosuid,noexec
```

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

```
/dev/cdrom /media/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 **MUST** 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/bash.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 *trusted* 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'),trusted 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 §3.14 "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/atjobs -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 §3.14.4 “*/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 a 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.

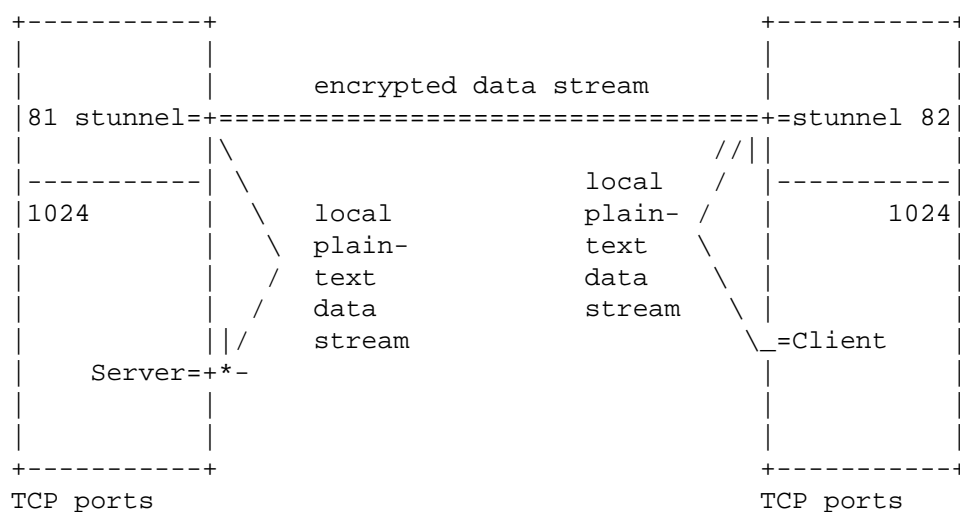
Stunnel has three main operating modes:

- 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 stunnel.cert`.

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 stunnel.cert >> /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+ezu0Pana7mmmvvg3Mi+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 (`stunnel.cert` 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 `CPath` 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 **MUST** 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:

```
# create secret key and self-signed certificate
openssl req -newkey rsa:1024 -nodes \
  -keyout /etc/stunnel/stunnel.pem \
  -new -x509 -sha1 -days 365 \
  -out /etc/stunnel/stunnel.cert

# set appropriate file permissions
chmod 400 /etc/stunnel/*.pem
chmod 444 /etc/stunnel/*.cert
```

```
# append copy of certificate to private key
echo >> /etc/stunnel/stunnel.pem
cat /etc/stunnel/stunnel.cert >> /etc/stunnel/stunnel.pem
```

The secret key contained in the `/etc/stunnel/stunnel.pem` file **MUST** be kept secret. The key files contain human-readable headers and footers along with the ASCII-encoded key, and the secret key is marked with the header "BEGIN RSA PRIVATE KEY".

You **MAY** distribute the public certificate stored in the `/etc/stunnel/stunnel.cert` file to clients, and is marked with the header "BEGIN CERTIFICATE".. Make sure you do not accidentally distribute the secret key instead.

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 `CAfile` 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 **MUST** 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.cert
```

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 one of the following cipher suites as defined in the SSL v3 protocol:

# Cipher	Proto	Key	Authen-	Encryption	Message
#		exchg	tication		auth code
#					
RC4-SHA	SSLv3	Kx=RSA	Au=RSA	Enc=RC4(128)	Mac=SHA1
DES-CBC3-SHA	SSLv3	Kx=RSA	Au=RSA	Enc=3DES(168)	Mac=SHA1
AES128-SHA	SSLv3	Kx=RSA	Au=RSA	Enc=AES(128)	Mac=SHA1
AES256-SHA	SSLv3	Kx=RSA	Au=RSA	Enc=AES(256)	Mac=SHA1

All other cipher suites and the other protocols supported by the OpenSSL library (TLSv1 and SSLv2) **MUST** be disabled.

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

```
ciphers = RC4-SHA:DES-CBC3-SHA:AES128-SHA:AES256-SHA
options = NO_TLSv1
options = NO_SSLv2
```

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:DES-CBC3-SHA:AES128-SHA:AES256-SHA
options = NO_TLSv1
options = NO_SSLv2
#
# 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 CApath instead of CAfile for externally signed certificates.
CAfile = /etc/stunnel/stunnel.cert
ciphers = RC4-SHA:DES-CBC3-SHA:AES128-SHA:AES256-SHA
options = NO_TLSv1
options = NO_SSLv2
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 **MUST** 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` in the client configuration file 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** specify

```
ciphers = RC4-SHA:DES-CBC3-SHA:AES128-SHA:AES256-SHA
options = NO_TLSv1
options = NO_SSLv2
```

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 3: 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 §3.3 "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 *mmdHHMMyyyy* 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.d/*
/etc/cron.daily/*
/etc/cron.hourly/*
/etc/cron.monthly/*
/etc/cron.weekly/*
/etc/crontab
/etc/ftpusers
/etc/group
/etc/gshadow
/etc/hosts
/etc/init.d/*
/etc/inittab
/etc/ld.so.conf
/etc/login.defs
/etc/modules.conf
/etc/pam.d/*
/etc/passwd
/etc/securetty
/etc/security/opasswd
/etc/security/pam_pwcheck.conf
/etc/security/pam_unix2.conf
/etc/shadow
/etc/ssh/ssh_config
/etc/ssh/sshd_config
```



```

/etc/stunnel/*
/etc/sysconfig/*
/etc/vsftpd.conf
/etc/xinetd.conf

/usr/lib/cracklib_dict.*

/var/log/audit.d/*
/var/log/faillog
/var/log/lastlog
/var/spool/atjobs/*
/var/spool/cron/*
/var/spool/cron/allow
/var/spool/cron/deny

```

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
YaST2	Directory for YaST2 log files
audit.d	Directory for LAuS logs
boot.msg	Messages from system startup
lastlog	Last successful log in (see <code>lastlog(8)</code>)
vsftpd.log	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>
xinetd.log	Written by xinetd, logging all connections

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 *AUDIT_ATTACH_ALL=1* in */etc/sysconfig/audit*).

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/init.d/audit restart
```

You **MUST NOT** attempt to reload the configuration by sending *auditd* a *HUP* signal or by running */etc/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 the `SuSEconfig` command and used as input to re-write other configuration files on the system.

The following is a brief overview of the security relevant files, including the specification of permitted changes.

In the evaluated configuration, no changes are permitted that would require running the `SuSEconfig` command to re-write other configuration files. You MAY run `SuSEconfig`, but it will have no effect on the evaluated configuration.

5.4.1 `suseconfig`

This file specifies global configuration variables. Most notably `ENABLE_SUSECONFIG`, which specifies whether `SuSEconfig` is allowed to modify other configuration files based on the variables in `/etc/sysconfig`.

Security relevant entries that MUST NOT be changed are:

```
ENABLE_SUSECONFIG="yes" # Is SuSEconfig allowed to modify configuration files?
MAIL_REPORTS_TO="root"  # Where are system status mails sent to
CWD_IN_ROOT_PATH="no"   # There MUST NOT be an entry for the current directory
CWD_IN_USER_PATH="no"   # There MUST NOT be an entry for the current directory
```

5.4.2 *security*

Specifies the operation mode and the configuration file for the SUSE permission system. Read by the *chkstat(8)* program which is run automatically by *yast2* after installation of new software. The following settings MUST NOT be changed:

```
CHECK_PERMISSIONS=set  
PERMISSION_SECURITY="eal4"
```

5.4.3 *cron*

Configures standard system cron jobs, like deletion of old files in */tmp* or update of the *man* databases. The settings are read by the shell scripts */etc/cron.daily/**. Security relevant variables are the following settings which MUST NOT be changed:

```
MAX_DAYS_IN_TMP=0           # How many days can files stay in /tmp  
TMP_DIRS_TO_CLEAR="/tmp /var/tmp" # Which temporary directories are checked  
OWNER_TO_KEEP_IN_TMP="root"   # Ids for which files will not be erased  
CLEAR_TMP_DIRS_AT_BOOTUP="no" # No cleaning of temp directories at boot
```

5.4.4 *language*

Sets up the default locale. This MUST NOT be changed, non-root users MAY override these default settings in their shell profiles.

5.4.5 *backup*

Configures the backup of the RPM database. MAY be changed.

5.4.6 *boot*

Configures the verbosity and interaction level of the boot process for debugging. Read by bootup scripts in */etc/init.d/*. MAY be changed.

5.4.7 *displaymanager*

This would configure the display manager for a workstation. It is not used in the evaluated configuration.

5.4.8 *kernel*

Configures modules to be installed in the *initrd* for system boot. MUST NOT be changed.

5.4.9 *clock*

Configures time zone and system clock, read during system boot. MAY be changed.

5.4.10 proxy

Configures global variables for the use of proxies. Not used in the evaluated configuration.

5.4.11 windowmanager

Would select the window manager on a workstation. Not used in the evaluated configuration.

5.4.12 sysctl

Configures some system variables for the boot process. The following are security relevant and **MUST NOT** be changed:

```
IP_DYNIP=no           # The system only has a static address
IP_TCP_SYNCOOKIES=yes # Syn Flood protection
IP_FORWARD=no         # Has to be set to yes if the system acts as a router.
ENABLE_SYSRQ=no       # System request key MUST be disabled.
```

5.4.13 java

Would configure the Java run time environment if installed. Not used in the evaluated configuration.

5.4.14 mail

Configures the MTA.

Security relevant variables are:

```
SMTPD_LISTEN_REMOTE="no" # MAY be set to yes to accept remote connections.
```

5.4.15 hardware

Configures hardware parameters (DMA), read during system boot. **MAY** be changed.

5.4.16 printer

Sets the default printer. **MUST NOT** be changed, but non-root users may override the setting in their shell profiles.

5.4.17 news

Usenet news / NNTP settings. Not used in the evaluated configuration.

5.4.18 console

Sets up the console configuration (font, code page, frame buffer). **MUST NOT** be changed.

5.4.19 keyboard

Sets up the console keyboard (repeat rate, layout, number of virtual consoles). MAY be changed.

5.4.20 mouse

Sets up the mouse type. Not used in the evaluated configuration.

5.4.21 lvm

Sets up LVM. MAY be changed.

5.4.22 network

This directory contains the networking configuration and scripts for the interfaces and routes. MAY be modified as needed, but IP addresses MUST be static (no DHCP).

5.4.23 syslog

Configures the *syslog* daemon. MAY be changed.

5.4.24 SuSEfirewall2

Configures the firewall. Not used in the evaluated configuration.

5.4.25 hotplug

Configures dynamically attached devices (USB, Firewire). OPTIONAL in the evaluated configuration.

5.4.26 ssh

Configures command line options for the SSH daemon. MUST NOT be changed.

5.4.27 postfix

Configures the basic MTA setup. MUST NOT be changed.

5.4.28 bootloader

Configures the type of bootloader to use and where to store the boot record. MUST NOT be changed.

5.4.29 audit

Configures tunable parameters for the kernel part of the audit subsystem. MUST NOT be changed.

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:

```
less /usr/share/doc/packages/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 SLES documentation is also installed in electronic form. `/usr/share/doc/packages/sles-inst-*/` contains the installation guide in PDF format, and `/usr/share/doc/packages/sles-admin-*/` the administration manual.

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 will be prompted 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;citd' too weak; contained in this document"
=> P'9tw;citd

```

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. Be aware that manipulations to the environment a program is run in can also cause security flaws, such as leaking sensitive information. Do not use the shell variables LD_LIBRARY_PATH or LD_PRELOAD that modify the shared library configuration used by dynamically linked programs.

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.

SuSE Linux Enterprise Server Installation Guide, [/usr/share/doc/packages/sles-inst-x86+x86-64_en/](#)
[/usr/share/doc/packages/sles-inst-ipseries_en/](#) [/usr/share/doc/packages/sles-inst-zseries_en/](#)

SuSE Linux Enterprise Server Administrator Guide, [/usr/share/doc/packages/sles-admin-x86+x86-64_en/](#)
[/usr/share/doc/packages/sles-admin-ipseries_en/](#) [/usr/share/doc/packages/sles-admin-zseries_en/](#)

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

Aleen 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          = yes;
    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`

```

#

```

```

# This is a sample filter.conf file.
# Please take a look at filesets.conf first if you
# wish to customize what system calls will be logged.
#
# The syntax of this file is described in filter.conf(5).
#

#
# Various primitive predicates
predicate      is-null          = eq(0);
predicate      is-negative      = lt(0);
predicate      is-system-uid     = lt(100);
predicate      is-lower-1024    = lt(-1024);

#
# Predicate to check open(2) mode: true iff
# (mode & O_ACCMODE) == O_RDONLY
predicate      is-rdonly        = mask(O_ACCMODE, O_RDONLY);

#
# Predicates for testing file type, valid when applied
# to a file type argument
predicate      __isreg           = mask(S_IFMT, S_IFREG);
predicate      __isdir           = mask(S_IFMT, S_IFDIR);
predicate      __ischr           = mask(S_IFMT, S_IFCHR);
predicate      __isblk           = mask(S_IFMT, S_IFBLK);
predicate      __issock          = mask(S_IFMT, S_IFSOCK);
predicate      __islnk           = mask(S_IFMT, S_IFLNK);
predicate      s_isreg           = __isreg(file-mode);
predicate      s_isdir           = __isdir(file-mode);
predicate      s_ischr           = __ischr(file-mode);
predicate      s_isblk           = __isblk(file-mode);
predicate      s_issock          = __issock(file-mode);
predicate      s_islnk           = __islnk(file-mode);
predicate      is-tempdir        = mask(01777, 01777);
predicate      is-world-writable = mask(0666, 0666);

#
# Predicates dealing with process exit code
predicate      if-crash-signal =
                                !mask(__WSIGMASK, 0)
                                && (mask(__WSIGMASK, __WSIGILL) ||
                                    mask(__WSIGMASK, __WSIGABRT) ||
                                    mask(__WSIGMASK, __WSIGSEGV) ||
                                    mask(__WSIGMASK, __WSIGSTKFLT));

#
# 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,
    SIOCDEFADDR,
    SIOCDAEP,
    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",
};
```