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

Klaus Weidner <klaus@atsec.com>

July 2, 2007; v2.4

Contents

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:

```
System x: x3550 (rack mount), HS20 and HS21 (blades)
Opteron (AMD): x3455 (rack mount), LS21 (blade)
System p: any POWER5 or POWER5+ system
System z: any z/Architecture compliant system or software
```

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 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 an Internet-connected server, or the case where services are to be provided to potentially hostile users.

It is assumed that the value of the stored assets merits moderately intensive penetration or masquerading attacks.

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.

1.3.3 Requirements for connectivity

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.

Any other systems with which the system communicates MUST be under the same management control and operate under the same security policy constraints.

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).

1.3.4 Requirements for administrators

There MUST be one or more competent individuals who are assigned to manage the system and the security of the information it contains. These individuals will have sole responsibility for the following functions: (a) create and maintain roles (b) establish and maintain relationships among roles (c) Assignment and Revocation of users to roles. In addition these individuals (as owners of the entire corporate data), along with object owners will have the ability to assign and revoke object access rights to roles.

The system administrative personnel MUST NOT be careless, willfully negligent, or hostile, and MUST follow and abide by the instructions provided by the administrator documentation.

In CAPP mode, 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 security features 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 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.

1.3.5 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.

Authorized users possess the necessary authorization to access at least some of the information managed by the system and are expected to act in a cooperating manner in a benign environment.

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.
- Rights for users to gain access and perform operations on information are based on their membership in one or more roles. These roles are granted to the users by the administrator. These roles MUST accurately reflect the users job function, responsibilities, qualifications, and/or competencies within the enterprise.
- A limited set of users is given the rights to create new data objects and they become owners for those data objects. The organization is the owner of the rest of the information under the control of system.
- Users are trusted to accomplish some task or group of tasks within a secure IT environment by exercising complete control over their data.

- 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.

2 Installation

The evaluation covers a fresh installation of SLES, 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.
- PCL 4 or PostScript level 1 compatible printers attached to the system using a parallel port or USB connection. You MAY also use a network printer.
- 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 or DVD is an authentic SUSE distribution CD/DVD for SLES 10 SP1. The original media are shipped in a sealed sleeve.

   If using downloaded ISO images, you MUST verify that the MD5 checksums of the image files are correct. The checksums are shown on the SUSE/Novell download web page and signed with the SUSE package signing GPG key. You MUST obtain the SUSE package signing key and ensure that the key is authentic, for example by getting the key from older SUSE distribution media that were previously authenticated, or from a trusted key server or other distribution method.

   The ISO images and signed MD5sums are available at these URLs:

   ```
   System x/Opteron, System p:
    http://download.novell.com/protected/Summary.jsp?buildid=2FNtOnmkx-w~

   System z:
    http://download.novell.com/protected/Summary.jsp?buildid=HfBRh4TspiE~
   ```

   After verifying and importing the SUSE package signing key (using `gpg --import`), use the following command to check if the signature is authentic:

   ```
   gpg --verify FILENAME
   ```

   Run `md5sum *.iso` to view the checksums for the downloaded image files, and compare them with those shown on the web page.

   You MUST use **SUSE LINUX Enterprise Server 10 SP1**. 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:

   - System x, System p, Opteron-based eServer: Insert the first CD and boot from CD-ROM.
   - System z: 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. **Select language:** choose **English (US)** to ensure that the messages shown during the installation match those described in this guide.

6. Accept the **License agreement**.

7. If prompted (due to having Linux installed already), choose **New installation**.

8. Configure **Clock and Time Zone**:

- RECOMMENDED: keep hardware clock time as **UTC**.
- RECOMMENDED: set the time zone as appropriate for the server location.
- REQUIRED: verify that the time displayed is accurate.

9. 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 :

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 System p 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:

```
/boot       75 MB
/         1200 MB
/tmp       200 MB
/home      100 MB
/var       384 MB
/var/log   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 "Extended User Attributes", "No access time", and "Mount read-only" as required.

Software

The following patterns are REQUIRED and MUST be marked with a check mark:

```
Server Base System
```

The following patterns are OPTIONAL:

```
Novell AppArmor
C/C++ Compiler and Tools
Print Server
```

Other patterns MUST NOT be checked.

On System p, you MUST select the **64bit Runtime Environment**, under "System Selection and System Tasks / Base Technologies".

In addition, you MUST select additional packages. Switch to the **Details...** view, and set **Filter** 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. For example, search for `audit` to quickly find all matching packages.

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.

On System p, after selecting the packages, Select [I]nformation -> [V]ersions to show the selected package versions. Use [S]earch to find the ***openssh*** and ***pwdutils*** package. Make sure you have the **"ppc"** variants of these two packages selected, instead of the "ppc64" ones.

```
### REQUIRED packages
amtu                        # Abstract Machine Test Utility
audit                       # User space tools for 2.6 kernel auditing
audit-libs                  # Dynamic library for libaudit
pwdutils-plugin-audit       # Plugin for the Linux Audit-Subsystem

### OPTIONAL packages
audit-devel                 # Development files for libaudit
vsftpd                      # FTP daemon (needs xinetd)
stunnel                     # set up encrypted SSL tunnels
```

Language

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

Booting (Expert)

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

Default Runlevel (Expert)

MUST be set to **3**.

10. To start the installation: press the **Accept** and **Install** buttons.
11. Installation will proceed. Insert the CDs as prompted by the installer.
12. 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.
13. The installer will continue in text mode, confirm the explanatory text about this.
14. **Password for "root", the administrator**

    - choose according to the password policy (§6.3)
    - in **Expert Options**, you MUST set Password Encryption: **MD5**
15. Configure an appropriate **Hostname and Domain Name**. You MUST NOT use the *Change Hostname via DHCP* option.
16. **Network Configuration**: Configure all installed network cards (zero or more) as appropriate for the platform. In the case of virtual network cards on System z, 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 DSL, modems, or ISDN adapters. You MUST NOT allow remote administration via VNC.

You MAY enable proxy settings if necessary.

Use the **Change...** menu to configure the **Network interfaces**.

You MUST use the *Traditional Method*, NOT the *NetworkManager* applet.

You MAY choose to enable or disable the firewall. It is RECOMMENDED to permit the SSH port.

You MAY enable IPv6.

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.

17. In the **Test Internet Connection** dialog, select **No, Skip This Test**. Use the **Next** button to continue.
18. In the **Installation Settings** dialog, you MAY change the **CA Management** settings if needed. The **OpenLDAP Server** MUST be disabled.
19. In the **User Authentication Method** dialog, select the **Local** authentication method.
20. In the **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 **User Management** dialog, and choose **Edit** to activate the **Existing Local User** dialog. Edit the **Password Settings** for the user you have just created according to the parameters described in section §3.9 "Setting up login controls" of this guide:

      ```
      Days before Password Expiration to Issue Warning          5
      Days after Password Expires with Usable Login            -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.
    - While still in the **Existing Local User** dialog, choose **Details**, select the **Additional group** pane and add membership in the group **trusted** for this administrator. **WARNING:** If you forget this step, the user will not be permitted to use the *su* tool. Use the **Accept** button to close the dialog.
    - You MAY use the **User Management** tool 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.

21. If prompted to install additional packages, do so by choosing **Continue**.
22. Confirm the **Release Notes** dialog by selecting **Next**.
23. In the **Hardware Configuration** dialog, you MUST NOT enable 3D Acceleration.
24. Confirm the **Installation Completed** dialog to start the system.
25. Wait for the freshly installed system to start, and verify that the issue message printed above the login prompt matches the installed system type and version. Then log in as ”root” and proceed with the next section.

3 Secure initial system configuration

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

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

Note that the system does not define audit rules as there is no universally applicable default for this. Please refer to section §5.3 ”Configuring the audit subsystem” of this guide for more information.

Log in as user ’root’ on the system console for the following steps.

3.1 Prerequisites

3.1.1 Filesystem configuration

You MUST remove the ”debugfs” line from the */etc/fstab* file. This filesystem is not supported in the evaluated configuration.

3.1.2 Replace pwdutils package for ppc64 systems

This section applies only to System p.

If the currently installed version of the ”pwdutils” package is the 64bit ”ppc64” version, you MUST replace it with the 32bit version. The automated configuration described in the next section will check for this and will refuse to continue if it detects the wrong version.

Use the following command to check which version you have installed:

```
rpm -q --queryformat='%{ARCH}\n' pwdutils
```

Locate the replacement file ”pwdutils-*.ppc.rpm” on the installation media, and use the ”rpm” command to replace the current package, for example:

```
rpm -Uvh --replacepkgs pwdutils-3.0.7.1-17.24.ppc.rpm
```

3.1.3 Ensure pam_tally and trusted program binaries match

This section applies only to System p.

The binaries for *sshd* and the *pam_tally* program MUST be installed with the same word size. The following command MUST indicate "32-bit" for both files:

```
file /usr/sbin/sshd /sbin/pam_tally
```

(The *sles-eal4* script will perform this check automatically.)

If the word sizes do not match, you MUST reinstall the "ppc" version of "openssh" and/or "pam" to ensure that both are 32bit.

3.2 Automated configuration of the system

The *certification-sles-ibm-eal4* package MUST be installed and used to achieve the evaluated configuration. This RPM package contains EAL4 specific configuration files, updates to the online documentation, 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.

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 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.

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.

WARNING: The `sles-eal4` script will reboot the system as the final step in the process, as described in the manual instructions in section §3.11 "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.

After software upgrades or installation of additional packages, you MUST ensure that the configuration remains secure. Please refer to sections §1.2 "How to use this document" and §4.4 "Installation of additional software" of this guide for additional information. You MAY re-run the *sles-eal4* script, but this does not guarantee that you will be in the evaluated configuration if you have added, deleted, modified, or replaced system components.

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.

The information in the remainder of this section provides background information about how this configuration was achieved, and mentions some changes you MAY make to the installed system while still remaining within the evaluated configuration. It is not intended to be a complete listing of the changes made to the system. Following the instructions in section §2 "Installation" of this guide followed by the automated reconfiguration is the only supported method to set up the evaluated configuration.

3.3 Add and remove packages

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.

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-eal4` script does not remove these packages if they are installed on the system:

```
Mesa-32bit
atk-32bit
audit-devel
audit-libs-32bit
audit-libs-64bit
autoconf
automake
bind-libs-32bit
bind-libs-64bit
binutils
binutils-32bit
binutils-64bit
bison
bison-32bit
blocxx-64bit
cairo-32bit
compat-32bit
compat-libstdc++-64bit
compat-openssl097g-32bit
compat-openssl097g-64bit
cpp
cpufrequtils-32bit
cpufrequtils-64bit
cups-libs-32bit
cups-libs-64bit
curl-32bit
curl-64bit
cvs
db-devel
dbus-1-32bit
dbus-1-64bit
dbus-1-glib-32bit
dbus-1-glib-64bit
device-mapper-32bit
device-mapper-64bit
expat-32bit
expat-64bit
expect
flex
flex-32bit
flex-64bit
krb5-32bit
krb5-64bit
libaio-32bit
libaio-64bit
libaio-devel
libaio-devel-32bit
libapparmor-32bit
libapparmor-64bit
libart_lgpl-32bit
libattr-devel
libcap-32bit
libcap-64bit
libcom_err-32bit
libcom_err-64bit
libdrm-32bit
libelf-32bit
libgcj-32bit
libgcrypt-32bit
libgcrypt-64bit
libgpg-error-32bit
libgpg-error-64bit
libgssapi-32bit
libgssapi-64bit
libidn-32bit
libidn-64bit
libjpeg-32bit
libjpeg-64bit
liblcms-32bit
liblcms-64bit
libmng-32bit
libmng-64bit
libnscd-32bit
libnscd-64bit
libpcap-32bit
libpcap-64bit
libpng-32bit
libpng-64bit
librtas
libstdc++-devel
libtiff-32bit
libtiff-64bit
parted-64bit
patch
perl-Digest-HMAC
perl-Digest-SHA1
perl-HTML-Parser
perl-HTML-Tagset
perl-Net_SSLeay
perl-TimeDate
perl-URI
perl-gettext
perl-libwww-perl
powerpc-utils
powerpc32
powersave-libs-32bit
powersave-libs-64bit
recode-64bit
s390-32
samba-32bit
samba-64bit
sles-preparation-power_en
sles-preparation-zseries_en
sqlite-32bit
sqlite-64bit
strace
strace-32bit
strace-64bit
sysfsutils-32bit
sysfsutils-64bit
sysvinit
tar
tcl
tcl-32bit
tcl-64bit
tcpd
tcpd-32bit
tcpd-64bit
tcsh
telnet
terminfo
texinfo
timezone
```

```
fontconfig-32bit
freetype2-32bit
freetype2-64bit
gcc
gcc-c++
gdb-32bit
gdb-64bit
gdbm-devel
gdbm-devel-32bit
gettext-32bit
gettext-64bit
gettext-devel
glib
glib2-32bit
glib2-64bit
glibc-devel
glibc-devel-32bit
glibc-devel-64bit
glitz-32bit
gmp-32bit
gmp-devel
gpg-pubkey
gpm-32bit
gpm-64bit
gtk2-32bit
hal-32bit
hal-64bit
howtoenh
hplip17-hpijs
kernel-ppc64
kernel-source
libtool-32bit
libtool-64bit
libusb-32bit
libusb-64bit
libxslt-32bit
libxslt-64bit
libzio-64bit
limal-nfs-server-perl
ltrace-32bit
make
mono-core-32bit
mpt-firmware
ncurses-devel
ncurses-devel-32bit
ncurses-devel-64bit
numactl
numactl-64bit
openct-32bit
openct-64bit
openmotif-libs-32bit
opensc-32bit
opensc-64bit
openslp-32bit
openslp-64bit
pam
pam-32bit
pam-64bit
pam-modules-32bit
pam-modules-64bit
pango-32bit
parted-32bit
tk
udev
update-desktop-files
usbutils
utempter
utempter-32bit
utempter-64bit
util-linux
vim
vsftpd
w3m
wget
xinetd
xorg-x11-libs-32bit
yast2
yast2-bootloader
yast2-core
yast2-country
yast2-inetd
yast2-installation
yast2-ldap
yast2-ldap-client
yast2-mail-aliases
yast2-mouse
yast2-ncurses
yast2-network
yast2-online-update
yast2-s390
zlib-devel
zlib-devel-32bit
```

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:

```
auditd
cron
network
random
syslog
```

The following services are OPTIONAL for runlevel 3:

```
boot.apparmor
cups
dbus
haldaemon
kbd
microcode
```

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

3.5 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 is the configuration entry in */etc/xinetd.d/vsftpd*:

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

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 §3.8 "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.

3.6 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.

3.7 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.

3.8 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 in the */etc/security/pam_pwdcheck.conf* file.

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:

3.8.1 /etc/pam.d/common-password

```
#
# /etc/pam.d/common-password - password-related modules common to all services
# CAPP configuration
#
# This file is included from other service-specific PAM config files,
# and should contain a list of modules that define  the services to be
# used to change user passwords.  The default is pam_unix2 in combination
# with pam_pwcheck.

# The "nullok" option allows users to change an empty password, else
# empty passwords are treated as locked accounts.
#
# To enable Blowfish or MD5 passwords, you should edit
# /etc/default/passwd.
#
# Alternate strength checking for passwords should be configured
# in /etc/security/pam_pwcheck.conf.
#
# pam_make can be used to rebuild NIS maps after password change.
#
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
```

3.8.2 /etc/pam.d/login

This file configures the behavior of the *login* program. It allows root login only for terminals configured in */etc/securetty*. If the file */etc/nologin* is present, then only root can log in. The 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_loginuid.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_loginuid.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
auth     required       pam_securetty.so
auth     required       pam_tally.so deny=5 onerr=fail
auth     include        common-auth
auth     required       pam_nologin.so
account  include        common-account
account  required       pam_tally.so
password include        common-password
session  include        common-session
session  required       pam_lastlog.so nowtmp
session  required       pam_resmgr.so
session  optional       pam_mail.so standard
session  optional       pam_loginuid.so     # no lockout on failure
```

3.8.3 /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 deny=5 onerr=fail
auth     include        common-auth
auth     required       pam_nologin.so
account  include        common-account
account  required       pam_tally.so
password include        common-password
session  include        common-session
session  required       pam_loginuid.so require_auditd
```

3.8.4 /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     include        common-auth
account  include        common-account
password required       pam_deny.so
session  include        common-session
session  optional       pam_xauth.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.

3.8.5 /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_listfile.so item=user sense=deny \
                        file=/etc/ftpusers onerr=succeed
auth     required       pam_tally.so deny=5 onerr=fail
auth     include        common-auth
account  include        common-account
account  required       pam_tally.so
account  required       pam_loginuid.so require_auditd
password required       pam_deny.so
session  include        common-session
```

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.

3.8.6 /etc/pam.d/crond

This file configures the cron service to ensure that tasks run on a user's behalf have the correct associated audit uid.

```
#
# The PAM configuration file for the cron daemon
#
#
auth     sufficient     pam_rootok.so
auth     include        common-auth
account  include        common-account
password include        common-password
session  include        common-session
session  required       pam_loginuid.so require_auditd
```

3.8.7 /etc/security/pam_pwcheck.conf

This file contains the default options 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 cracklib remember=7
```

3.8.8 /etc/security/pam_unix2.conf

This file contains the default options 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
```

3.9 Setting up login controls

The system supports various options to control log ins in */etc/login.defs*. Comments in the file explain the options and values that MUST be set for the EAL4 evaluated configuration.

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.

3.9.1 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.

3.10 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.

3.10.1 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 $1$O471l/$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.

3.10.2 Yaboot boot loader configuration

Yaboot is used on the System p 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 System p, 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
```

3.10.3 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 `zipl` command must be re-run after any kernel or boot argument modifications. Please refer to the *zipl*(8) man page for more information.

The following example shows a typical */etc/zipl.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"
```

3.11 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/
```

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 create 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

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.

Errors in the non interactive jobs executed by `cron` 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` is controlled through the following *allow* and *deny* files:

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

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

In the SLES distribution, the *allow* file does not exist, and the *deny* file is used to prevent system-internal IDs and/or guest users from using the service. By default, the evaluated configuration permits all non-system users to use *cron*.

It is RECOMMENDED to restrict the use of *cron* 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/cron.deny
chmod 600 /etc/cron.deny
```

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

You MAY create a */etc/cron.allow* file to explicitly list users who are permitted to use this service. If you do create the file, it 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.

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. Alteratively, 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 §3.8 ”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 jobs
crontab -u $U -r

# 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_loginuid.so* module in the `session` stack. Please refer to section §3.8.2 "/etc/pam.d/login" of this guide for more information about the needed PAM configuration.

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

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

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

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

```
init q
```

4.9 SYSV shared memory and IPC objects

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

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

4.10 Configuring secure network connections with *stunnel*

4.10.1 Introduction

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

Stunnel has three main operating modes:

- Accept incoming SSL-encrypted TCP connections, and run a specific program to handle the request.

  This is similar to how *xinetd* launches programs, and any program compatible with *xinetd* can also be used for this purpose. It must read and write the communication data on the *stdin* and *stdout* file descriptors and stay in the foreground. *stunnel* also supports switching user and group IDs before launching the program.

- Open a SSL connection to a remote SSL-capable TCP server, and copy data to and from *stdin* and *stdout*.

- Bind a TCP port to accept incoming unencrypted connections, and forward data using SSL to a prespecified remote server.

The following diagram shows a sample usage scenario:

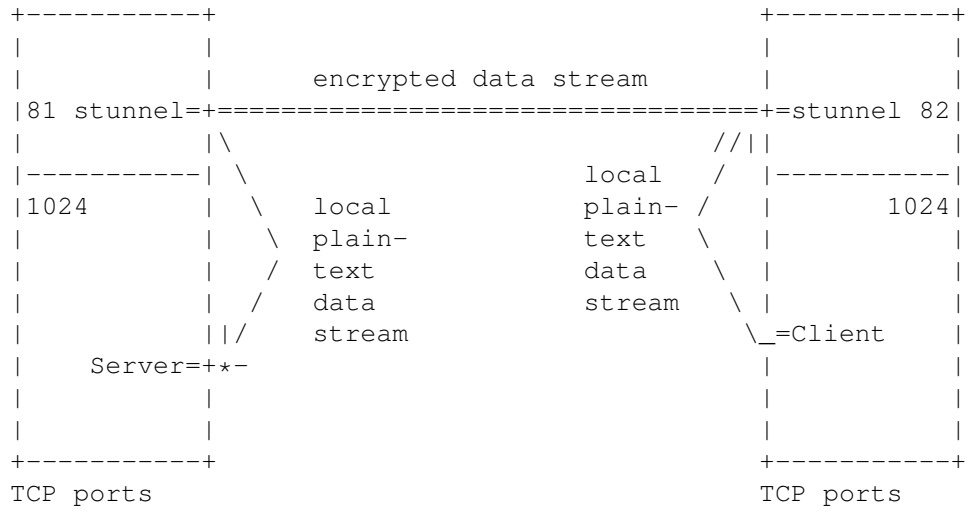

In this scenario, neither the client nor the server have administrator privileges, they are running as normal user processes. Also, the client and server do not support encryption directly.

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

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

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

stunnel MAY also be used by non-administratorive 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+ezu0Pana7mmmvg3Mi+BuqwlQ/iU+G/qrG6VGj
-----END RSA PRIVATE KEY-----

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

You MAY distribute the original signed certificate (*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 *CApath* 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 *mmddHHMMyyyy* to specify the numeric month, day, hour, minute and year respectively. For example, the following command sets the clock to May 1st 2004, 1pm in the local time zone:

```
date 050113002004
```

The *hwclock*(8) can query and modify the hardware clock on supported platforms, but is not available in virtual environments such as z/VM. 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.

4.13 AppArmor configuration

The evaluated configuration keeps the AppArmor system enabled in a static configuration, but does not depend on AppArmor for any security features. You MAY modify the AppArmor configuration, for example to add additional restrictions, but this is beyond the scope of the CAPP evaluation.

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/audit/*
/etc/cron.allow
/etc/cron.deny
/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/cron/*
```

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

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

```
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          Directory for audit logs
boot.msg       Messages from system startup
lastlog        Last successful log in  (see lastlog(8))
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 susbystem, see who(1)
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 refer to the *auditd*(8), *auditd.conf*(8), and *auditctl*(8) man pages for more information.

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 audit package provides a suggested audit configuration in the */usr/share/doc/packages/audit/sample.rules* file. It contains a suggested setup for a typical multiuser system, all access to security relevant files is audited, along with other security relevant events such as system reconfiguration. You MAY copy the sample rules files to */etc/audit.rules* and modify the configuration according to your local requirements, including the option of using an empty audit rules file to disable auditing if not required.

You MAY selectively disable and enable auditing for specific events or users as required by modifying the *audit.rules* file. For example, you can include and exclude specific users from auditing by adding filters based on the loginuid, such as the following entry:

```
-a exit,always -F auid!=trusteduser -S chown
```

The audit system also supports filtering on success or failure of system call operations:

```
-F success=1   # for successful syscalls
-F success!=1  # for unsuccessful syscalls
```

You MAY configure filesystem watches using the `-w` option. Note that filesystem watches are order sensitive if you create multiple watches for the same inode, for example if creating separate watches for multiple hard links to a single file. You can filter filesystem watches, for example to exclude a user ID from being audited:

```
-w /etc/shadow -k Secret
-a watch,never -F auid=trusteduser
-a exit,possible -S open
```

It is RECOMMENDED that you always reconfigure the audit system by modifying the */etc/audit.rules* file and then running the following command to reload the audit rules:

```
# as role "auditadm_r"
auditctl -R /etc/audit.rules
```

This procedure ensures that the state of the audit system always matches the content of the */etc/audit.rules* file. You SHOULD NOT manually add and remove audit rules and watches on the command line as those changes are not persistent.

Note that reloading audit rules involves initially deleting all audit rules, and for a short time the system will be operating with no or only a partial set of audit rules. It is RECOMMENDED to make changes to the audit rules when no users are logged in on the system, for example by using single user mode or a reboot to activate the changes.

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

5.3.3 Reading and searching the audit records

Use the *ausearch*(8) tool 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 datestamped copy of the applicable configuration with the log files for future reference.

For example:

```
# search for events with a specific login UID
ausearch -ul jdoe

# search for events by process ID
ausearch -p 4690
```

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

Audit messages from the *pwdutils* user management tools, including *useradd* and *gpasswd*, have the following format:

```
type=USER_CHAUTHTOK msg=audit(11.059:63): user pid=33 uid=501 auid=501
msg='op=permission denied - account=bin, id=1, by=1000 id=2 exe=gpasswd
     (hostname=?, addr=?, terminal=pts/1 res=success)'
```

For these specific messages, you MUST refer to the text contained in the "op=" field to determine the success or failure of the operation. The "res=success" parenthetical value is always the same and MUST be ignored.

For some system calls on some platforms, the system call arguments in the audit record can be slightly different than you may expect from the program source code due to modifications to the arguments in the C library or in kernel wrapper functions. For example, the *mq_open*(3) glibc library function strips the leading '/' character from the path argument before passing it to the *mq_open*(2) system call, leading to a one character difference in the audit record data. Similarily, some system calls such as *semctl*(2), *getxattr*(2), and *mknodat*(2) can have additional internal flags automatically added to the flag argument. These minor modifications do not change the security relevant information in the audit record.

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 audit log file or output generated by the low-level *ausearch* tool.

5.3.4 Starting and stopping the audit subsystem

If the audit daemon is terminated, no audit events are saved until it is restarted. To avoid lost audit records when you have modified the filter configuration, you MUST use the command `/etc/init.d/audit reload` to re-load the filters.

You MUST NOT use the *KILL* signal (-9) to stop the audit daemon, doing so would prevent it from cleanly shutting down.

It is RECOMMENDED that you add the kernel parameter `audit=1` to your boot loader configuration file to ensure that all processes, including those launched before the *auditd* service, are properly attached to the audit subsystem. Please refer to the documentation of your boot loader and section §3.10 "Configuring the boot loader" of this document for more details.

5.3.5 Storage of audit records

The default audit configuration stores audit records in the */var/log/audit/audit.log* file. This is configured in the */etc/audit/auditd.conf* file. You MAY change the *auditd.conf* file to suit your local requirements.

It is RECOMMENDED that you configure the audit daemon settings appropriately for your local requirements, for example by changing the log file retention policy to never delete old audit logs with the following setting in the */etc/audit/auditd.conf* file:

```
max_log_file_action = KEEP_LOGS
```

The most important settings concern handling situations where the audit system is at risk of losing audit information, such as due to lack of disk space or other error conditions. You MAY choose actions appropriate for your environment, such as switching to single user mode (action `single`) or shutting down the system (action `halt`) to prevent auditable actions when the audit records cannot be stored.

Halting the system is RECOMMENDED and most certain way to ensure all user processes are stopped. The following settings are RECOMMENDED in the */etc/auditd.conf* file if a fail-secure audit system is required:

```
admin_space_left_action = SINGLE
disk_full_action = HALT
disk_error_action = HALT
```

It is RECOMMENDED that you configure appropriate disk space thresholds and notification methods to receive an advance warning when the space for audit records is running low.

It is RECOMMENDED that you use a dedicated partition for the */var/log/audit/* directory to ensure that *auditd* has full control over the disk space usage with no other processes interfering.

Please refer to the *auditd.conf*(5) man page for more information about the storage and handling of audit records.

5.3.6 Reliability of audit data

You MAY choose an appropriate balance between availability of the system and secure failure mode in case of audit system malfunctions based on your local requirements.

You MAY configure the system to cease all processing immediately in case of critical errors in the audit system. When such an error is detected, the system will then immediately enter "panic" mode and will need to be manually rebooted. To use this mode, add the following line to the */etc/audit/audit.rules* file:

```
-f 2
```

Please refer to the *auditctl*(8) man page for more information about the failure handling modes.

You MAY edit the */etc/libaudit.conf* file to configure the desired action for applications that cannot communicate with the audit system. Please refer to the *get_auditfail_action*(3) man page for more information.

auditd writes audit records using the normal Linux filesystem buffering, which means that information can be lost in a crash because it has not been written to the physical disk yet. Configuration options control how *auditd* handles disk writes and allow the administrator to choose an appropriate balance between performance and reliability.

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 the buffering settings do not affect normal operation.

The default setting is `flush = DATA`, ensuring that record data is written to disk, but metadata such as the last file time might be inconsistent.

The highest performance mode is `flush = none`, but be aware that this can cause loss of audit records in the event of a system crash.

If you want to ensure that auditd always forces a disk write for each record, you MAY set the `flush = SYNC` option in */etc/audit/auditd.conf*, but be aware that this will result in significantly reduced performance and high strain on the disk.

A compromise between crash reliability and performance is to ensure a disk sync after writing a specific number of records to provide an upper limit for the number of records lost in a crash. For this, use a combination of `flush = INCREMENTAL` and a numeric setting for the `freq` parameter, for example:

```
flush = INCREMENTAL
freq = 100
```

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

5.4 System configuration variables in */etc/sysconfig*

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

6 Security guidelines for users

6.1 Online Documentation

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

```
man ls
man man
```

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

```
apropos password
```

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

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

```
info diff
```

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

```
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 scrollback 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 scrollback 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/manual/sles-preparation-*/*

SUSE Linux Enterprise Server Administrator Guide, */usr/share/doc/manual/sles-admin_en/*

David A. Wheeler, "Secure Programming for Linux and Unix HOWTO", */usr/share/doc/howto/en/html/Secure-Programs-HOWTO.html*, *http://tldp.org/HOWTO/Secure-Programs-HOWTO/*

Kevin Fenzi, Dave Wreski, "Linux Security HOWTO", */usr/share/doc/howto/en/html/Security-HOWTO.html*, *http://www.linuxsecurity.com/docs/LDP/Security-HOWTO/*

7.2 Literature

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

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

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

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

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

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

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

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