

IPWorks OS Hardening Guide

USER GUIDE

Copyright

© Ericsson AB 2017, 2018. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

1	Introduction	1
1.1	Prerequisite	1
1.1.1	Software	1
1.1.2	Personnel	1
1.2	Related Information	1
2	System Configuration Process Overview	3
3	System Configuration	5
3.1	Hardening Configuration of etc-overlay	5
3.1.1	Installing the LDE etc-overlay	6
3.1.2	Changing System configuration	7
3.2	Access Control	11
3.2.1	File permissions	12
3.3	Hardening Configuration in cluster.conf	12
3.3.1	SSH Restriction to a Specific Network	12
3.3.2	Disabling Root Access over SSH Connection	14
3.3.3	Disabling Core Files (Optional)	14
3.3.4	Reloading the configuration	15
4	LDE Firewall	17
	Reference List	19





1 Introduction

This document provides instructions to harden the Linux Distribution Extensions (LDE) product used for the IPWorks application. This means configuring the node to a generally agreeable security level.

The issues dealt with here are configuring the operating system services, managing user accounts and permissions, and installing additional third-party security software.

The fundamental principle to be followed is “the more minimal, the more secure”.

For Additional node and application-specific hardening information, refer to [IPWorks Application Components Hardening Guide](#).

1.1 Prerequisite

1.1.1 Software

The IPWorks software must be installed before OS hardening. For more information about how to install IPWorks, refer to [IPWorks Deployment Guide](#).

1.1.2 Personnel

The personnel following these instructions must be familiar with:

- Working in UNIX-like environment without X-Windowing system.
- UNIX shell commands.

Note: It is recommended to create the backup file whenever to modify a configuration file.

1.2 Related Information

Trademark information, typographic conventions, and definition and explanation of abbreviations and terminology can be found in the following documents:

- [Trademark Information](#)
- [Typographic Conventions](#)
- [Glossary of Terms and Acronyms](#)





2 System Configuration Process Overview

Figure 1 describes the configuration process of Node Hardening:

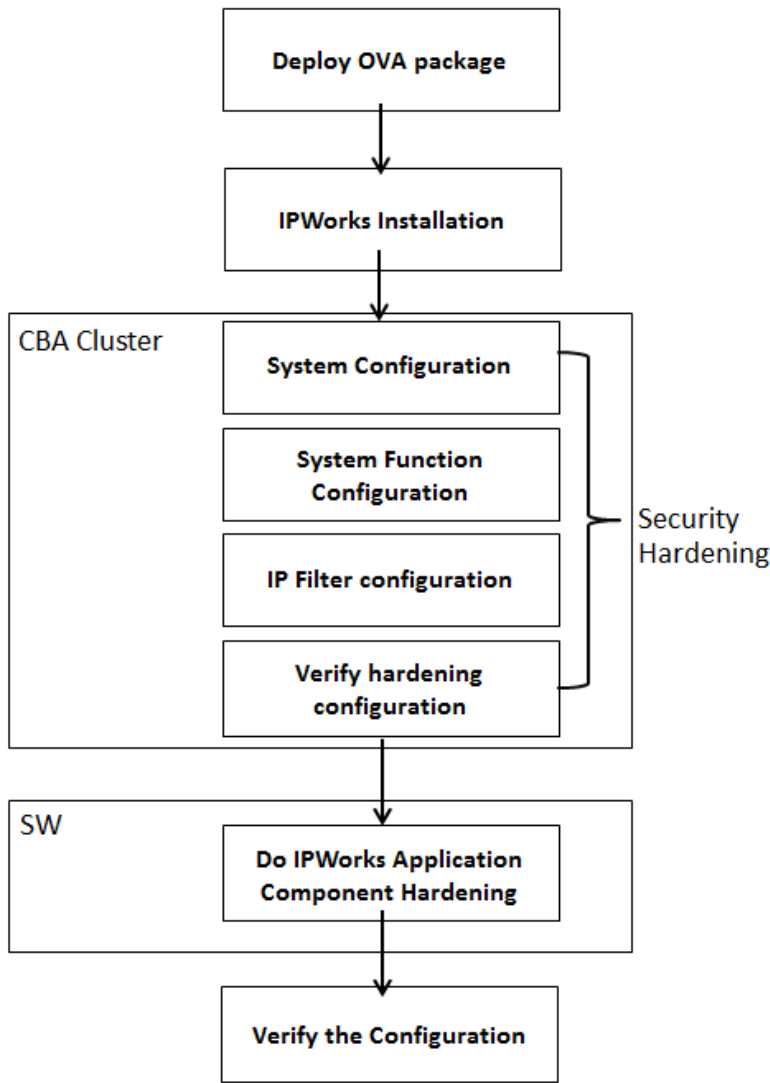


Figure 1 Configuration Process of Node Hardening





3 System Configuration

The system is configured either using a plain text file, stored as `/cluster/etc/cluster.conf` or for static changes to configuration files using the `etc-overlay-framework`. The `cluster.conf` file contains blade and site-specific parameters, such as number of included blades, IP addresses, external DNS, or NTP servers while the `etc-overlay-framework` is typically used for very application-specific configuration such as hardening.

The `cluster.conf` is applied on the system using the `lde-config` command while the `etc-overlay-framework` is applied by installing a rpm.

For the detail description and configuration of LDE hardening, refer to LDE Management Guide.

3.1 Hardening Configuration of etc-overlay

This section describes what can be hardened through the `etc-overlay-framework`, how to install the `lde-etc-overlay`, and how to create and install a customized `etc-overlay`.

The `etc-overlay-framework` is intended for use when complete files under the `/etc` directory need to be overwritten and this needs to be done before normal system startup.

A good example of a use of the `etc-overlay-framework` is for hardening changes where the configuration must be applied before networking is configured and the changes should be made in a non-dynamic fashion to prevent damaged configuration files.

Two `etc-overlays` are provided, one default and one hardened. This section contains a summary of the different modifications made to the default SLES configuration files in the two `etc-overlays`. Modifications that will be present in the default and hardened overlay respectively is described in Table 1.

Table 1 Overlays Provided By LDE

Modification	default-overlay	hardened-overlay
Default Umask 027		yes
Legal Warning at Login		yes
Inactivity Timer for User Account		No ⁽¹⁾



Modification	default-overlay	hardened-overlay
Strong Password Enforcement		No ⁽¹⁾
Auditing - Full Personal Accountability		yes

(1) This function is not supported yet.

For more information about the LDE hardened-overlay, refer to the Section The provided etc-overlays in LDE Management Guide.

For the hardening through etc-overlay, you can install the LDE hardened etc-overlay as a default hardened etc-overlay (see Section 3.1.1 on page 6). In addition, you can create a customized etc-overlay that contains the specify configurations for IPWorks (see Section 3.1.2 on page 7).

3.1.1 Installing the LDE etc-overlay

You can choose to install one of the etc-overlays (default or hardened etc-overlay), it is recommended to install the hardened etc-overlay.

To install the etc-overlay, do the following on SC-1 or SC-2:

1. Copy the following etc-overlay rpm packages to `/cluster/rpms` in SC-1 or SC-2.

```
lde-default-etc-overlay-cxp9024735-1.5.0-0.sle12.noarch.rpm
```

```
lde-hardened-etc-overlay-cxp9024737-1.5.0-0.sle12.noarch.rpm
```

These two rpm packages are located in the folder `/cee_deployment/hardening` in IPWorks Utilities Package. The IPWorks software packages can be retrieved from SW Gateway. For specific information, see the product release notes.

2. Install the etc-overlay to all the nodes by executing the following command:

```
cluster rpm --add <etc-overlay-rpm-filename> -n <node id>
```

Note: When the `<node id>` corresponds to a SC node (SC-1, SC-2), the command must be executed locally; when it corresponds to a PL node, the command can be executed either on a SC node or on the PL node. You can check the `<node id>` in `/cluster/etc/cluster.conf`, for example, "node 1 control SC-1", which means the node id of SC-1 node is 1.

3. Check whether the rpm `<etc-overlay-rpm-filename>` is installed.

```
cluster rpm -l -n <node id> | grep overlay
```

```
SC-1:/cluster/rpms # cluster rpm -l -n 1 | grep overlay
```



```
lde-hardened-etc-overlay-cxp9024737 1.5.0-0.sle12
lde-hardened-etc-overlay-cxp9024737-1.5.0-0.sle12.noarch.rpm
```

4. Reboot the cluster to make the etc-overlay work.

```
cluster reboot -a
```

3.1.2 Changing System configuration

Certain system configurations are necessary for IPWorks hardening, such as, Network Parameters. To realize the system configuration, you must create a customized etc-overlay that contains the configuration, and install the customized etc-overlay rpm to the node you want to configure.

Before you create the customized etc-overlay, you need to follow Section 3.1.2.1 on page 7 to create temp folders and back up default configuration. Then follow Section 3.1.2.2 on page 8 to add some new configuration to the temp folders. And Section 3.1.2.3 on page 10 shows how to create the customized etc-overlay rpm according to your temp configuration folder and how to install the customized etc-overlay rpm.

Note:

- Do not reboot the cluster during the configuration. Otherwise, certain things created in the /tmp folder will be deleted after cluster reboot.
- Log on to one SC node to perform the system configuration.

3.1.2.1 Prepare for the Customized etc-overlay

- If you did not install lde-etc-overlay before you start this section, create your temp configuration folder on SC-1:

```
mkdir -p /tmp/create-etc-overlay/all
```

```
mkdir -p /tmp/create-etc-overlay/payload
```

```
mkdir -p /tmp/create-etc-overlay/control
```

Note: Those folders are empty and no configuration files are backed up.

- If you installed the lde-etc-overlay before you start this section, you need to get the lde-etc-overlay tarball, add this configuration file to your configuration and create your temp configuration folder on SC-1:

```
mkdir -p /tmp/create-etc-overlay
```

```
tar -zxvf /usr/lib/lde/etc-overlay/etc-overlay.tgz -C
/tmp/create-etc-overlay/
```

Then you can find the lde-etc-overlay configuration in the folder you created:



```
ls /tmp/create-etc-overlay/all/
```

Creates some folders to store your configuration files on SC or PL nodes:

```
mkdir -p /tmp/create-etc-overlay/payload
```

```
mkdir -p /tmp/create-etc-overlay/control
```

Note: The files in "all" folder will be overwritten to all the nodes, the files in "control" folder will be overwritten to control nodes, the files in "payload" folder will be overwritten to payload nodes. And the precedence-order is "all" < "control" < "payload".

Configuration files can be modified or created under the /etc path. Back up those configuration files into the temp configuration folders. Otherwise the configuration files will lose after removing the installed etc-overlay-rpm.

Example

If the file /etc/fstab is created to set mount automatically, cp the /etc/fstab file to /tmp/create-etc-overlay/control/.

```
cp -p /etc/fstab /tmp/create-etc-overlay/control/
```

3.1.2.2 Modify Configuration Files

This section takes setting network parameters for example to show how to modify configuration files.

IP traffic parameters indirectly prevent the operating system from the intrusion in the Internet. The IP parameters bring extra checking to the system process. The intruder can send for example lots of ICMP messages in short period that match to the IP traffic parameters. This traffic flow can trigger the pre-configured system policy to reduce, reject, or block specific IP traffic in the system. The IP traffic parameters can affect the IP traffic performance.

To set the network parameters, cp the /etc/sysctl.conf file to /tmp/create-etc-overlay/all/.

```
cp -p /etc/sysctl.conf /tmp/create-etc-overlay/all/
```

Edit /tmp/create-etc-overlay/all/sysctl.conf and insert the parameter values below:

Disable ICMP broadcast echo activity. Otherwise, your system can be used as part of a Smurf attack:

```
net.ipv4.icmp_echo_ignore_broadcasts=1
```

Disable bogus ICMP error logging. Otherwise, your system logs bogus ICMP errors which can cause the log overflow. This is often used by hackers, who attempt to DoS your syslog.



```
net.ipv4.icmp_ignore_bogus_error_responses=1
```

Disable ICMP routing redirects. Otherwise, your system can have its routing table misadjusted by an attacker.

```
net.ipv4.conf.all.accept_redirects=0
```

```
net.ipv6.conf.all.accept_redirects=0
```

```
net.ipv4.conf.all.send_redirects=0
```

Disable IP Source Routing.

Now the IP source routing is only used by attackers trying to spoof IP addresses that you trust as internal hosts.

```
net.ipv4.conf.all.accept_source_route=0
```

```
net.ipv4.conf.all.forwarding=0
```

```
net.ipv6.conf.all.forwarding=0
```

```
net.ipv6.conf.all.accept_source_route=0
```

```
net.ipv6.conf.all.accept_ra_defrtr=0
```

```
net.ipv6.conf.all.accept_ra_pinfo=0
```

```
net.ipv6.conf.all.accept_ra_rtr_pref=0
```

```
net.ipv6.conf.all.accept_dad=1
```

```
net.ipv6.conf.all.accept_ra=1
```

```
net.ipv6.conf.default.autoconf=0
```

```
net.ipv6.conf.all.max_addresses=1
```

```
net.ipv6.conf.all.dad_transmits=0
```

```
net.ipv6.conf.all.router_solicitations=0
```

Enforce sanity checking, also called ingress filtering, or egress filtering. The point is to drop a packet if the source and destination IP addresses in the IP header do not make sense when considered in light of the physical interface on which it arrived.

```
net.ipv4.conf.all.rp_filter=1
```

Log and drop "Martian" packets. A "Martian" packet is one for which the host does not have a route back to the source IP address (it apparently dropped in from Mars). These days most hosts have a default route, meaning that there would be no such thing as a Martian packet, but to be safe and complete.



```
net.ipv4.conf.all.log_martians=1
```

Increase resilience under heavy TCP load (which makes the system more resistant to SYN Flood attacks).

```
net.ipv4.tcp_max_syn_backlog=1280
```

```
net.ipv4.tcp_syncookies=1
```

If IPv6 addresses are used for traffic IPs, do the followings:

1. Remove etc-overlay.

```
cluster rpm --remove <etc-overlay-rpm-filename> -n <node id>
```

2. Remove rpm in /cluster/rpms.

```
rm /cluster/rpms/<etc-overlay-rpm-filename>
```

3. Add following content in the /tmp/create-etc-overlay/all/sysctl.conf:

```
net.ipv6.conf.all.autoconf=0
net.ipv6.conf.default.autoconf=0
net.ipv6.conf.eth1.autoconf=0
net.ipv6.conf.eth2.autoconf=0
net.ipv6.conf.all.use_tempaddr = 0
net.ipv6.conf.default.use_tempaddr=0
net.ipv6.conf.eth1.use_tempaddr=0
net.ipv6.conf.eth2.use_tempaddr=0
```

3.1.2.3 Install the Customized etc-overlay

After the configuration, you need to follow the steps below to deploy your new etc-overlay.

1. Create your etc-overlay

```
lde-etc-overlay.py create-rpm <toplevel-folder-name> \
-f test -V <version> -R <release>
```

Note: Here -f option only support “test”. It’s an LDE issue, and will be fixed later.

Naming convention of the built etc-overlay RPM file:

```
test-etc-overlay-<version>-<release>.noarch.rpm
```

Example:

```
lde-etc-overlay.py create-rpm /tmp/create-etc-overlay/ -f test
-V 1.0 -R R1A
```



You will get test-etc-overlay-1.0-R1A.noarch.rpm in your current folder.

2. Copy your new etc-overlay to /cluster/rpms/.

```
cp <etc-overlay-rpm-filename> /cluster/rpms/
```

3. Remove your installed lde-etc-overlay for all the nodes.

Note: If you did not install the lde-etc-overlay, skip this step.

```
cluster rpm --remove <etc-overlay-rpm-filename> -n <node id>
```

Note: When the <node id> corresponds to a SC node (SC-1, SC-2), the command must be executed locally; when it corresponds to a PL node, the command can be executed either on a SC node or on the PL node. You can check the <node id> in /cluster/etc/cluster.conf, for example, "node 1 control SC-1", which means the node id of SC-1 node is 1.

4. Reboot the cluster.

```
cluster reboot -a
```

Note: After this cluster reboot, the MySQL will not be startup automatically, please execute the following command on one SC node to startup the MySQL NDB Cluster.

```
SC-1:~ # /etc/init.d/ipworks.mysql stop-ndbcluster
```

```
SC-1:~ # /etc/init.d/ipworks.mysql start-ndbcluster
```

5. Install the etc-overlay you created for all the nodes.

```
cluster rpm --add <etc-overlay-rpm-filename> -n <node id>
```

Note: When the <node id> corresponds to a SC node (SC-1, SC-2), the command must be executed locally; when it corresponds to a PL node, the command can be executed either on a SC node or on the PL node. You can check the <node id> in /cluster/etc/cluster.conf, for example, "node 1 control SC-1", which means the node id of SC-1 node is 1.

6. Reboot the cluster to make the etc-overlay work.

```
cluster reboot -a
```

To verify the configuration, you can access the /etc/ folder to check if the files you modified are there.

3.2 Access Control

This section shows the access control performed on the SC node directly, the configuration only works on SC nodes.



3.2.1 File permissions

Permissions on system files need to be set correctly so that a user or process cannot make unauthorized changes to these files. Most of the permissions setting are automatically properly set during operating system installation. However, SUSE systems have various levels of pre-set permissions, which are defined in the files `/etc/permissions`, `/etc/permissions.easy`, `/etc/permissions.local`, `/etc/permissions.secure` and `/etc/permissions.paranoid`.

There are also some permission settings defined for the operation of particular programs in files under `/etc/permissions.d/`.

Set permissions to secure settings by executing the following command:

```
# chkstat -set /etc/permissions.secure
```

```
# chkstat -n --system --set /usr/bin/fusermount
```

There are a few files that need to have SUID root permissions set: such programs are a possible attack vector: you can check which files on the system have SUID permissions with a command such as:

```
# find / -type f -perm -u=s -ls 2>/dev/null
```

Files that are both executable and writable by others can be found using

```
# find / -type f -perm -o=w,u=x -ls 2>/dev/null
```

It is critical to protect system log files from being modified by unauthorized individuals. Also, certain logs contain sensitive data that should only be available to the system administrator. From the entire system point of view, it is important to check the file permissions.

3.3 Hardening Configuration in `cluster.conf`

This section describes what can be hardened through the configuration file `/cluster/etc/cluster.conf`.

You can add parameters (described in the following subsections) in the `/cluster/etc/cluster.conf`.

3.3.1 SSH Restriction to a Specific Network

The ssh parameters are shown in Table 2; the iptables parameters are shown in Table 4.

For details, refer to [IPWorks IPTables Service Configuration](#).



Table 2 Ssh

Syntax	ssh <target> <network>	
Description	Restrict SSH to listen to a specific network. This parameter can be defined multiple times if SSH shall listen to more than one network. ⁽¹⁾	
Options	<target> <network>	Target blades. Name of the network that SSH should listen to.
Examples	ssh payload internal	
Exceptions	If the ssh keyword is defined for a network, SSH cannot be used towards movable IPs on that network. If SSH access to movable IPs is required, all ssh parameters should be removed and iptables used to restrict SSH traffic.	

(1) If ssh is not defined for a blade, no restriction on SSH is made, meaning it listens to all available interfaces.

Table 3 Ssh.port

Syntax	ssh.port <target> <port number>	
Description	The port that the ssh service should listen for connections on, instead of the default 22. ⁽¹⁾	
Options	<target> <port number>	Target blades. The network port that ssh should listen on for connections. Value range is 1 - 65535.
Examples	ssh.port all 1024	

(1) If you need to change the ssh port, only change the ssh port for your OAM network on control node. Port 22 is still listened on the active COM node, it is used by the COM/COMSA service.

Table 4 Iptables

Syntax	iptables <target> <command>	
Description	Defines a rule in iptables. ⁽¹⁾	Rules are run in the order specified in this configuration.



Options	<code><target></code> <code><command></code>	Target blades. Specifies the parameters that should be passed to iptables. This can be any parameter accepted by iptables.
Examples	On all nodes, drop packets destined from source address 10.0.0.1: <pre>iptables all -A INPUT -s 10.0.0.1 -j DROP</pre> On all nodes, accept SSH traffic destined for the 192.168.0.0/24 network and drop all other SSH traffic: <pre>iptables all -A INPUT -p tcp --dport 22 -d 192.168.0.0/24 -j ACCEPT</pre> <pre>iptables all -A INPUT -p tcp --dport 22 -j DROP</pre>	

(1) To make the rule take effect, cluster reboot is also required after cluster configuration reloading.

Note: After rebooting the cluster, execute the following command to check if the rule takes effect.

```
iptables -L
```

3.3.2 Disabling Root Access over SSH Connection

The ssh.rootlogin parameters are shown in Table 5.

Table 5 Ssh.rootlogin

Syntax	<code>ssh.rootlogin <target> <value></code>	
Description	Disable root login over the ssh connection. ⁽¹⁾	
Options	<code><target></code> <code><value></code>	Target blades. Either on to permit root login over ssh or off to not permit root login over ssh.
Examples	<code>ssh.rootlogin all off</code>	

(1) If `ssh.rootlogin` is not defined for a blade, the default is to permit root login over SSH

3.3.3 Disabling Core Files (Optional)

The coredump parameters are shown in Table 6.



Table 6 Coredump

Syntax	coredump <target> <size> coredump <target> unlimited	
Description	Specifies the default core dump size for processes.	
Options	<target> <size> unlimited	Target node(s) The maximum size of the core dump in Kilobytes. The theoretical unlimited size of a core dump.
Examples	coredump all 0 coredump 1 0	

3.3.4

Reloading the configuration

After modifying the `/cluster/etc/cluster.conf`, you need to execute the following command to make it work.

```
lde-config --reload --all
```

For the iptables parameters modification in Table 4 and the core files configuration in Table 6, you need to reboot the cluster to take effect for the change.





4 LDE Firewall

The cluster configuration provides the possibility to control network access towards the cluster through the use of standard packet filtering features found in Linux, that is iptables.

For more information on how to configure iptables, refer to [IPWorks IPTables Service Configuration](#).





Reference List

- [1] Trademark Information
- [2] Typographic Conventions
- [3] Glossary of Terms and Acronyms
- [4] IPWorks Application Components Hardening Guide
- [5] IPWorks Deployment Guide, 21/1553-AVA 901 33/2 Uen
- [6] IPWorks IPTables Service Configuration
- [7] LDE Management Guide, 1/1553-CAA 901 2978/4 Uen