Oracle Linux v7.6 Common Criteria Guidance Document

May 27, 2021

1.7

Prepared By:
Acumen Security
2400 Research Blvd Suite 395
Rockville, MD, 20850
www.acumensecurity.net

Prepared for:
Oracle Corporation
500 Oracle Parkway
Redwood Shores, CA 94065
USA
Tel.: +1.650.506.7000
www.oracle.com

Trademarks

Oracle Linux and the Oracle logo are trademarks or registered trademarks of Oracle Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds.

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

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

Legal Notice

Table of Contents

Revision History

Version	Date	Description
0.1	March 8, 2019	Initial draft
0.2	March 22, 2019	Updates to Section 1
0.3	April 11, 2019	Updates to Section 1, 2
0.4	April 15, 2019	Addition of Sections 9 - 11
0.5	May 1, 2019	Updates to Sections 11-14
0.6	May 9, 2019	Updates to Sections 14-18
0.7	June 25, 2019	Minor updates based on Oracle feedback
0.8	August 21, 2020	Updates on sections 5-6
0.9	September 1, 2020	Minor updates to Sections 5-6
1.0	October 22, 2020	Minor updates to Sections based on AGD requirements.
1.1	November 12, 2020	Addition of sample audit events
1.2	November 23, 2020	Addressing certifier's comments
1.3	December 7, 2020	Addressing OR comment
1.4	January 5, 2021	Minor updates to address AGD assurance activities
1.5	January 22, 2021	Minor updates to Section 7
1.6	March 24, 2021	Addressing certifier comments
1.7	May 27, 2021	Addressing certifier comments

1 Purpose of this document

This document is intended to be a supplement to the Oracle public user documentation. This Common Criteria guidance document contains configuration information needed to configure and administer the Oracle Linux v7.6. The Oracle Linux conforms to the Protection Profile for General Purpose Operating Systems Version 4.2.1 (OS PP v4.2.1). The information contained in this document is intended for Administrators who would be responsible for the configuration and management of the Oracle Linux 7.6.

1.1 TOE Overview

The Oracle Linux v7.6 (herein referred to as the TOE) is a Linux-based operating system. Oracle Linux is a general purpose, multi-user, multi-tasking Linux based operating system. It provides a platform for a variety of applications. In addition, virtual machines provide an execution environment for a large number of different operating systems. It satisfies all of the criterion to meet the Protection Profile for General Purpose Operating Systems Version 4.2.1.

1.2 TOE Description

The TOE in the evaluated configuration consists of the following platforms:

- X86 64-bit Intel Platform with Intel(R) Xeon(R) Silver 4114 processor
- AMD hardware-based platform with AMD EPYC 7XXX series processor
- KVM (kernel based virtual machine) platform

1.3 Assumptions

The following Assumptions are for the Operational Environment:

Assumptions	Operational Environment
A.PLATFORM	The OS relies upon a trustworthy computing platform for its execution. This underlying platform is out of scope of this PP.
A.PROPER_USER	The user of the OS is not willfully negligent or hostile, and uses the software in compliance with the applied enterprise security policy. At the same time, malicious software could act as the user, so requirements which confine malicious subjects are still in scope.
A.PROPER_ADMIN	The administrator of the OS is not careless, willfully negligent or hostile, and administers the OS within compliance of the applied enterprise security policy.

Table 1 Assumptions on the Operational Environment

1.4 TOE Delivery

The TOE software can be download from the Oracle website. When software updates are available via the http://www.oracle.com website, they can obtain, verify the integrity and install the updates.

2 Prerequisites for Installation

Before beginning this procedure, you must ensure that:

- Oracle Linux 7.6 ISO image required for Common Criteria certification as applicable for the list of platforms listed in Section 1.2 of this guidance document.
- Ensure that the ISO image is bootable from DVD-R using TSFT and NFS.
- Ensure that the ISO image's integrity verification has been performed for all files downloaded from the Oracle website. The integrity is verified using SHA-256 hash sum provided by the Oracle web site.
- The minimum system requirements are the following:
    - 1 GB of memory
    - 1 CPU or vCPU
    - 1 NIC
    - 5 GB minimum disk space, But 20 GB is the recommended size.

NOTE: A graphical user interface (GUI) or system administration or any other operation is not included in the evaluated configuration. From the initial setup menu, please select Minimal Installation.

3 Installation of the Oracle Linux v7.6

The steps below are applicable whether Oracle Linux v7.6 is installed on a server, or virtual device.

The ISO can be downloaded from the Oracle Cloud download center accessible from the following link:
https://edelivery.oracle.com/

NOTE: Users must register and configure an account to gain access to download software.

4 Enabling FIPS mode of operation

In order to be complaint for Common Criteria, the user must ensure FIPS mode is enabled.
Follow the steps below to put the OS into FIPS mode of operation:

1. Install the dracut-fips package:
yum install dracut-fips

The dracut-fips package provides the modules to build a dracut initramfs file system that performs an integrity check.

2. If the system CPU supports AES New Instructions (AES-NI), install the package.

Run the following command to check whether the system supports AES-NI:
grep aes /proc/cpuinfo

To install the package:
yum install dracut-fips-aesni

3. Recreate the initramfs file system.

dracut -f

4. Perform the following steps to configure the boot loader so that the system boots into FIPS mode:

Identify the boot partition and the UUID of the partition, for example:

```
# df /boot
Filesystem     1K-blocks   Used Available Use% Mounted on
/dev/sda1       508588 294476    214112  58% /boot

# blkid /dev/sda1
```

/dev/sda1: UUID="6046308a-75fc-418e-b284-72d8bfad34ba" TYPE="xfs"

As the root user, edit the /etc/default/grub file as follows:

- Add the fips=1 option to the boot loader configuration.

```
GRUB_CMDLINE_LINUX="vconsole.font=latarcyrheb-sun16
 rd.lvm.lv=ol/swap rd.lvm.lv=ol/root crashkernel=auto
 vconsole.keymap=uk rhgb quiet fips=1"
```

- If the contents of /boot reside on a partition other than the root partition, you must use the boot=UUID=***boot_UUID*** line to the boot loader configuration to specify the device that should be mounted onto /boot when the kernel loads.

```
GRUB_CMDLINE_LINUX="vconsole.font=latarcyrheb-sun16
 rd.lvm.lv=ol/swap rd.lvm.lv=ol/root crashkernel=auto
 vconsole.keymap=uk rhgb quiet
 boot=UUID=6046308a-75fc-418e-b284-72d8bfad34ba fips=1"
```

This is required for FIPS to perform kernel validation checks, where it verifies the kernel against the provided HMAC file in the /boot directory.

- Save your changes.

5. Rebuild the GRUB configuration as follows:

On BIOS-based systems, run the following command:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

On UEFI-based systems, run the following command:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

6. To ensure proper operation of the in-module integrity verification, prelinking must be disabled on all system files.

By default, the prelink package is not installed on the system. However, if it is installed, disable prelinking on all libraries and binaries as follows:

Set PRELINKING=no in the /etc/sysconfig/prelink configuration file.

If the libraries were already prelinked, undo the prelink on all of the system files as follows:

```
# prelink –u -a
```

7. Reboot to apply these settings.

8. Verify that FIPS mode is enabled:

cat /proc/sys/crypto/fips_enabled

The response should be 1 which shows that FIPS mode is enabled.

5 Configuring SSH

The OS implements SSH protocol that complies with RFCs 4251, 4252, 4253, 4254, and 6668 as a client and server. The TOE supports password-based authentication and public key based authentication.

5.1 Configuring SSH Server

The SSH server is allowed to use only approved ciphers. This can be configured with the configuration file /etc/ssh/sshd_config.

To configure an OpenSSH server

1. Install or update the openssh and openssh-server packages:

yum install openssh openssh-server

2. Start the sshd service and configure it to start following a system reboot:

systemctl start sshd

systemctl enable sshd

You can set sshd configuration options for features such as Kerberos authentication, X11 forwarding, and port forwarding in the /etc/ssh/sshd_config file.

NOTE: The default Oracle Linux installation includes the openssh and openssh-client packages.

- The following public key algorithms: ssh-rsa.
- The SSH client shall ensure that, as described in RFC 4253, packets greater than 262144 bytes in an SSH transport connection are dropped.
- The TOE supports the following encryption algorithms: aes128-ctr, aes256-ctr, aes128-cbc, and aes256-cbc.
- The TOE supports the following data integrity MAC algorithms: hmac-sha1, hmac-sha2-256, and hmac-sha2-512.
- The TOE supports the following key exchange algorithms: diffie-hellman-group14-sha1.

5.2 Configuring SSH Client

To configure an OpenSSH client, install or update the openssh and openssh-client packages:

yum install openssh openssh-client

The SSH client is allowed to use only approved ciphers. This can be configured with the configuration file of /etc/ssh/ssh_config

- The following public key algorithms: ssh-rsa.
- The SSH client shall ensure that, as described in RFC 4253, packets greater than 262144 bytes in an SSH transport connection are dropped.
- The TOE supports the following encryption algorithms: aes128-ctr, aes256-ctr, aes128-cbc, and aes256-cbc.
- The TOE supports the following data integrity MAC algorithms: hmac-sha1, hmac-sha2-256, and hmac-sha2-512.
- The TOE supports the following key exchange algorithms: diffie-hellman-group14-sha1.

5.3 Using Public Key authentication

The OS supports public key authentication. In order to generate a public and public key pair, the following command should be used:

ssh-keygen

Press Enter each time that the command prompts you to enter a passphrase. Use the ssh-copy-id script to append the public key in the local ~/.ssh/id_rsa.pub file to the ~/.ssh/authorized_keys file on the remote system. You can now use the OpenSSH utilities to access the remote system without supplying a password. As the script suggests, you should use ssh to log into the remote system to verify that the ~/.ssh/ authorized_keys file contains only the keys for the systems from which you expect to connect.

6 Configuring TLS

Administrators are directed to not use ECC certificates in the evaluated configuration.

TOE supports RSA key sizes of 2048 bits, 3072 bits and 4096 bits for key generation The RSA keys are used in support of digital signature for TLS sessions.

The TOE supports FFC schemes using cryptographic key sizes of 2048-bits or greater. The FFC scheme is used as part of key generation.

Any Certificate Authority can be used to generate or sign the certificate. TLS v1.2 is supported.

The following ciphersuites are supported by the OS for TLS session establishments:

- TLS_RSA_WITH_AES_128_CBC_SHA as defined in RFC 5246
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256 as defined in RFC 5246
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256 as defined in RFC 5246

The following command is used to generate RSA private key:

openssl genrsa

The following command is used to generate a certificate signing request:

openssl req

The following command is used to sign certificate signing request:

openssl ca

7 Cryptographic Key Destruction

The TOE is capable of performing key destruction.

The TOE does not support delayed key destruction.

When using SSDs, the wear levelling mechanism prevents software to overwrite the exact physical location where the keys are stored.

Key data may still reside on the physical data store albeit it cannot be retrieved by the operating system anymore. Yet, forensic tools may recover that data. Thus, an SSD must be physically destroyed at the end of life to guarantee that no cryptographic keys remain.

The system uses many more keys than outlined in the preceding sections. Those keys are always ephemeral and maintained in RAM. These keys will be securely erased by the system without user intervention.

8 Configuring User Authentication

The Pluggable Authentication Modules (PAM) feature allows you to enforce strong user authentication and password policies, including rules for password complexity, length, age, expiration and the reuse of previous passwords. You can configure PAM to block user access after too many failed login attempts, after normal working hours, or if too many concurrent sessions are opened.

The PAM configuration file (/etc/pam.d/system-auth) contains the following default entries for testing a password's strength:

The line for pam_pwquality.so defines that a user gets three attempts to choose a good password. From the module's default settings, the password length must a minimum of six characters, of which three characters must be different from the previous password. The module only tests the quality of passwords for users who are defined in /etc/passwd.

```
password  requisite   pam_pwquality.so try_first_pass local_users_only retry=3 authtok_type=
password  sufficient  pam_unix.so sha512 shadow nullok try_first_pass use_authtok
password  required    pam_deny.so
```

The line for pam_unix.so specifies that the module tests the password previously specified in the stack before prompting for a password if necessary (pam_pwquality will already have performed such checks for users defined in /etc/passwd), uses SHA-512 password hashing and the /etc/shadow file, and allows access if the existing password is null. An alternate way of defining password requirements is available by selecting the Password Options tab in the Authentication Configuration GUI (system-config-authentication).

9 Mounting Filesystems

To access a file system's contents, you must attach its block device to a mount point in the directory hierarchy. You can use the mkdir command to create a directory for use as a mount point, for example:

mkdir /var/projects

You can use an existing directory as a mount point, but its contents are hidden until you unmount the overlying file system.

The mount command attaches the device containing the file system to the mount point:

mount [options] device mount_point

You can specify the device by its name, UUID, or label. For example, the following commands are equivalent ways of mounting the file system on the block device /dev/sdb1:

mount /dev/sdb1 /var/projects

mount UUID="ad8113d7-b279-4da8-b6e4-cfba045f66ff" /var/projects

mount LABEL="Projects" /var/projects

If you do not specify any arguments, mount displays all file systems that the system currently has mounted, for example:

mount

/dev/mapper/vg_host01-lv_root on / type ext4 (rw)

10 Creating User Accounts

To create a user account by using the useradd command:

Enter the following command to create a user account:

useradd [options] username

You can specify options to change the account's settings from the default ones.

By default, if you specify a user name argument but do not specify any options, useradd creates a locked user account using the next available UID and assigns a user private group (UPG) rather than the value defined for GROUP as the user's group.

Assign a password to the account to unlock it:

passwd username
The command prompts you to enter a password for the account.

If you want to change the password non-interactively (for example, from a script), use the chpasswd command instead:

echo "username:password" | chpasswd
Alternatively, you can use the newusers command to create a number of user accounts at the same time.

10.1 Locking an Account

To lock a user's account, enter:
passwd -l username
To unlock the account:
passwd -u username

10.2 Modifying or Deleting User Accounts

To modify a user account, use the usermod command:
Creating Groups
usermod [options] username
For example, to add a user to a supplementary group (other than his or her login group):
usermod -aG groupname username
You can use the groups command to display the groups to which a user belongs, for example:
groups root root : root bin daemon sys adm disk wheel
To delete a user's account, use the userdel command:
userdel username

10.3 Creating Groups

o create a group by using the groupadd command:
groupadd [options] groupname
Typically, you might want to use the -g option to specify the group ID (GID). For example:

groupadd -g 1000 devgrp

10.4 Modifying or Deleting Groups

To modify a group, use the groupmod command:

groupmod [options] username

To delete a user's account, use the groupdel command:

groupdel username

10.5 Changing User Passwords

One can change the user password by using the following command:

passwd <username>

10.6 Password Policy

It is recommended that users choose password that are strong. The initial password set by the administrator be changed when you first login to the system. Password length should be a minimum of 8 characters. Passwords can be comprised of special characters, upper case, lower case, and numeric characters.

11 Management Functions

The TOE maintains the following roles: Administrator and User.

The management functions are listed below:

Management Function	**Administrator**	**User**
Enable/disable [*session timeout*]	X	
Configure [*session]* inactivity timeout	X	
Configure local audit storage capacity	X	
Configure minimum password Length	X	
Configure minimum number of special characters in password	X	
Configure minimum number of numeric characters in password	X	
Configure minimum number of uppercase characters in password	X	
Configure minimum number of lowercase characters in password	X	

Management Function	Administrator	User
Configure lockout policy for unsuccessful authentication attempts through [*limiting number of attempts during a time period]*	X	
Configure host-based firewall	X	
Configure name/address of directory server with which to bind		
Configure name/address of remote management server from which to receive management settings		
Configure name/address of audit/logging server to which to send audit/logging records		
Configure audit rules	X	
Configure name/address of network time server	X	
Enable/disable automatic software update	X	
Configure WiFi interface		
Enable/disable Bluetooth interface		
Enable/disable [*no other devices*]	X	
No other management functions	X	

12 System Firewall

To implement a general-purpose firewall, you can use the Firewall Configuration GUI (firewall-config), provided by the firewall-config package. To create or modify a firewall configuration from the command line, use the firewall-cmd utility (or, if you prefer, the iptables, or ip6tables utilities) to configure the packet filtering rules. The packet filtering rules are recorded in the /etc/firewalld hierarchy for firewalld and in the /etc/sysconfig/iptables and /etc/sysconfig/ip6tables files for iptables and ip6tables.

13 Network Time Service

The ntpd daemon can synchronise the system clock with remote NTP servers, with local reference clocks.

To configure the ntpd service on a system:

1. Install the ntp package.

yum install ntp

2. Edit /etc/ntp.conf to set up the configuration for ntpd.

14 Session Timeout

The OS supports CLI and SSH session timeouts.

The session inactivity on the terminal is defined by a time-out in either /etc/screenrc or ~/.screenrc using the idle X lockscreen configuration value where X is an integer value specifying the idle time in seconds before the screen is locked.

For remote SSH, the session timeout can be set with the ClientAliveInterval. As an administrator user, open the sshd_config file:

```
# vi /etc/ssh/sshd_config
```

Locate the ClientAliveInterval option to 60 (in seconds) or add the value if it is not there.

```
ClientAliveInterval 60
```

Note : ClientAliveInterval: number of seconds that the server will wait before sending a null packet to the client (to keep the connection alive).

Restart sshd daemon :

```
# sudo systemctl restart sshd.service
```

15 Storage of Sensitive Data

Keys and configuration files are stored in /etc directory. Privileges are controlled by permissions to invoke applications and to access data. Due to privileges being controlled by permissions, this prevents users from performing management functions that they do not have access to.

16 Application Developers

Application developers should use the following compiler options as best practice when developing applications invoking the gcc compiler and linker.

The stack-protector-strong flag has been developed to broaden the scope of the stack protection without extending it to every function in the program.

-fstack-protector-strong --param=ssp-buffer-size=4

ASLR improves executable security in terms of memory randomization and access protection.

-fpie -Wl,-pie

17 Setting System Time and Date

Date and time representation on a system can be set to match a specific timezone. To list all of the available timezones, run:

timedatectl list-timezones

To set the system timezone to match a value returned from the available timezones, you can run:

timedatectl set-timezone America/Los_Angeles

One can check your system's current date and time configuration by running the **timedatectl** command on its own

To set system time manually, you can use the timedatectl set-time command. For example. you can run:

timedatectl set-time "2018-10-28 01:59:59"

This command sets the current system time based on the time specified assuming the currently set system timezone. The command also updates the system Real Time Clock (RTC).

18 Applying Updates

Oracle provides regular updates to the Oracle Linux operating system. After initial installation, the update mechanism is fully configured to obtain updates.

To upgrade the system to the latest version of Oracle Linux, use the following command:

Use the following command to upgrade the system to the latest available update of Oracle Linux version.

```
# yum update
```

The command fetches the current list of available updates for the current installation base. If updates are available for one or more of the currently installed packages, the update command will list them and asks the administrator whether to install them. In case the update fails for any reason, yum will list the reason and prevent the update operation.

Use the following command to install or update a specific package:

```
# yum update package
```

The automated installation of updates can be done by using yum-cron which is installed using the following command:

```
# yum --enablerepo=ol7_optional_latest install yum-cron
```

NOTE: Please note that updated versions are not covered by the certification.

19 Auditing

Auditing collects data at the kernel level that you can analyze to identify unauthorized activity. Auditing collects more data in greater detail than system logging. The process of examining audit trails to locate events of interest can be a significant challenge that you will probably need to automate.

The TOE generates audit events for all start-up and shut-down function. The TOE leverages the Lightweight Audit Framework (LAF) audit system. Audit events are generated for the following audit functions:

- Start-up and shut-down of the audit functions
- Authentication events (Success/Failure)
- Use of privileged/special rights events (Successful and unsuccessful security, audit, and configuration changes)
- Privilege or role escalation events (Success/Failure)

Sample audit records for each of the above audit functions are listed below:

Start-up of the audit function

type=SERVICE_START msg=audit(1605110310.137:92): pid=1 uid=0 auid=4294967295 ses=4294967295 subj=system_u:system_r:init_t:s0 msg='unit=kdump comm="systemd" exe="/usr/lib/systemd/systemd" hostname=? addr=? terminal=? res=success'

Shut-down of the audit function

type=SERVICE_STOP msg=audit(1605110314.650:93): pid=1 uid=0 auid=4294967295 ses=4294967295 subj=system_u:system_r:init_t:s0 msg='unit=NetworkManager-dispatcher comm="systemd" exe="/usr/lib/systemd/systemd" hostname=? addr=? terminal=? res=success'

Authentication Events

type=USER_AUTH msg=audit(1605109361.516:629): pid=24756 uid=0 auid=4294967295 ses=4294967295 subj=system_u:system_r:sshd_t:s0-s0:c0.c1023 msg='op=PAM:authentication grantors=pam_unix acct="root" exe="/usr/sbin/sshd" hostname=? addr=? terminal=ssh res=success'

type=USER_AUTH msg=audit(1605109357.161:627): pid=24756 uid=0 auid=4294967295 ses=4294967295 subj=system_u:system_r:sshd_t:s0-s0:c0.c1023 msg='op=PAM:authentication grantors=? acct="root" exe="/usr/sbin/sshd" hostname=? addr=? terminal=ssh res=failed'

Privilege Escalation

type=USER_AUTH msg=audit(1605117077.712:186): pid=7028 uid=1000 auid=1000 ses=4 subj=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 msg='op=PAM:authentication grantors=pam_unix acct="root" exe="/usr/bin/su" hostname=? addr=? terminal=pts/0 res=success

type=USER_AUTH msg=audit(1605117066.821:185): pid=7025 uid=1000 auid=1000 ses=4 subj=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 msg='op=PAM:authentication grantors=? acct="root" exe="/usr/bin/su" hostname=? addr=? terminal=pts/0 res=failed'

Use of privileged/special rights

type=USER_CHAUTHTOK msg=audit(1605117783.298:291): pid=7171 uid=1000 auid=0 ses=7 subj=unconfined_u:unconfined_r:passwd_t:s0-s0:c0.c1023 msg='op=change password id=1000 exe="/usr/bin/passwd" hostname=? addr=? terminal=pts/0 res=success'

type=USER_CHAUTHTOK msg=audit(1605117611.485:249): pid=7102 uid=1000 auid=1000 ses=6 subj=unconfined_u:unconfined_r:passwd_t:s0-s0:c0.c1023 msg='op=PAM:chauthtok grantors=? acct="tester" exe="/usr/bin/passwd" hostname=? addr=? terminal=pts/0 res=failed'

Each audit record contains the following information:

- Date and time is marked with “time”
- Type of event is referenced with “type”
- Subject identity is specified by “uid” containing the numeric user ID of the respective user
- Outcome of the event identified with the field “success”
- User identity is given with the “auid” field

The audit configuration file, /etc/audit/auditd.conf, defines the data retention policy, the maximum size of the audit volume, the action to take if the capacity of the audit volume is exceeded, and the locations of local and remote audit trail volumes. The default audit trail volume is /var/log/audit/audit.log.

By default, auditing captures specific events such as system logins, modifications to accounts, and **sudo** actions. You can also configure auditing to capture detailed system call activity or modifications to certain files. The kernel audit daemon (auditd) records the events that you configure, including the event type, a time stamp, the associated user ID, and success or failure of the system call.

The entries in the audit rules file, /etc/audit/audit.rules, determine which events are audited. Each rule is a command-line option that is passed to the **auditctl** command. You should typically configure this file to match your site's security policy.

The following are examples of rules that you might set in the /etc/audit/audit.rules file.

Record all unsuccessful exits from open and truncate system calls for files in the /etc directory hierarchy.

-a exit,always -S open -S truncate -F /etc -F success=0

Record all files opened by a user with UID 10.

-a exit,always -S open -F uid=10

Record all files that have been written to or that have their attributes changed by any user who originally logged in with a UID of 500 or greater.

-a exit,always -S open -F auid>=500 -F perm=wa

Record requests for write or file attribute change access to /etc/sudoers, and tag such record with the string sudoers-change.

-w /etc/sudoers -p wa -k sudoers-change

Record requests for write and file attribute change access to the /etc directory hierarchy.

-w /etc/ -p wa

Require a reboot after changing the audit configuration. If specified, this rule should appear at the end of the /etc/audit/audit.rules file.

-e 2

19.1 Starting and Stopping Audit

If the audit daemon is stopped, audit events are not saved until the system is restarted.

To avoid loss of audit records when you have modified the filter configuration, use the following command:

systemctl reload auditd

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 connected to the audit subsystem.

19.2 Storage of Audit Records

The audit configuration stores audit records in the /var/log/audit/ directory by default. This is configured in the /etc/audit/auditd.conf file. The auditd.conf file can be altered based on the users' local requirements.

You can 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 following settings are found in the /etc/audit/auditd.conf file if a fail-secure audit system is required:

admin_space_left_action = SINGLE

disk_full_action = HALT

disk_error_action = HALT

19.3 Retrieving Audit Records

The following command can be used to retrieve information from the audit events:

Searching for events by process ID:

```
# ausearch -p 4690
```

20 Access Control Lists

POSIX Access Control Lists (ACLs) provide a richer access control model than traditional UNIX Discretionary Access Control (DAC) that sets read, write, and execute permissions for the owner, group, and all other system users. You can configure ACLs that define access rights for more than just a single user or group, and specify rights for programs, processes, files, and directories. If you set a default ACL on a directory, its descendents inherit the same rights automatically. An ACL consists of a set of rules that specify how a specific user or group can access the file or directory with which the ACL is associated. A regular ACL entry specifies access information for a single file or directory. A default ACL entry is set on directories only, and specifies default access information for any file within the directory that does not have an access ACL.

20.1 Configuring Access Control Lists

Install the acl package:

```
# yum install acl
```

Edit /etc/fstab and change the entries for the file systems with which you want to use ACLs so that they include the appropriate option that supports ACLs, for example:

```
LABEL=/work     /work      ext4     acl     0 0
```

20.2 Setting and Displaying Access Control Lists

To add or modify the ACL rules for file, use the setfacl command:

```
# setfacl -m rules file ...
```

The rules take the following forms:

```
[d:]u:user[:permissions]
```

Sets the access ACL for the user specified by name or user ID. The permissions apply to the owner if a user is not specified.

```
[d:]g:group[:permissions]
```

Sets the access ACL for a group specified by name or group ID. The permissions apply to the owning group if a group is not specified.
[d:]m[:][:permissions]

Sets the effective rights mask, which is the union of all permissions of the owning group and all of the user and group entries.

[d:]o[:][:permissions]

Sets the access ACL for other (everyone else to whom no other rule applies).

The permissions are r, w, and x for read, write, and execute as used with chmod. The d: prefix is used to apply the rule to the default ACL for a directory.

To display a file's ACL, use the getfacl command, for example:

getfacl foofile

If extended ACLs are active on a file, the -l option to ls displays a plus sign (+) after the permissions, for example:

ls -l foofile
-rw-r--r--+ 1 bob bob 105322 Apr 11 11:02 foofile

The following are examples of how to set and display ACLs for directories and files.

Grant read access to a file or directory by a user.

setfacl -m u:*user*:r *file*

Display the name, owner, group, and ACL for a file or directory.

getfacl *file*

Remove write access to a file for all groups and users by modifying the effective rights mask rather than the ACL.

setfacl -m m::rx *file*

The -x option removes rules for a user or group.

Remove the rules for a user from the ACL of a file.

setfacl -x u:*user file*

Remove the rules for a group from the ACL of a file.

setfacl -x g:*group file*

The -b option removes all extended ACL entries from a file or directory.

setfacl -b *file*

Copy the ACL of file *f1* to file *f2*.

getfacl *f1* | setfacl --set-file=- *f2*

Set a default ACL of read and execute access for other on a directory:

setfacl -m d:o:rx *directory*

Promote the ACL settings of a directory to default ACL settings that can be inherited.

getfacl --access *directory* | setfacl -d -M- *directory*

The -k option removes the default ACL from a directory.

setfacl -k *directory*

21 Self-tests

When an Oracle Linux system boots, it performs the following operations:

The computer's BIOS performs a power-on self-test (POST), and then locates and initializes any peripheral devices including the hard disk.

The BIOS reads the Master Boot Record (MBR) into memory from the boot device. (For GUID Partition Table (GPT) disks, this MBR is the protective MBR on the first sector of the disk.) The MBR stores information about the organization of partitions on that device. On a computer with x86 architecture, the MBR occupies the first 512 bytes of the boot device. The first 446 bytes contain boot code that points to the boot loader program, which can be on the same device or on another device. The next 64 bytes contain the partition table. The final two bytes are the boot signature, which is used for error detection.

The default boot loader program used on Oracle Linux is GRUB 2, which stands for Grand Unified Bootloader version 2.

The boot loader loads the vmlinuz kernel image file into memory and extracts the contents of the initramfs image file into a temporary, memory-based file system (tmpfs).

The kernel loads the driver modules from the initramfs file system that are needed to access the root file system.

The kernel starts the systemd process with a process ID of 1 (PID 1). systemd is the ancestor of all processes on a system. systemd reads its configuration from files in the /etc/systemd directory. The /etc/systemd/system.conf file controls how systemd handles system initialization.

systemd reads the file linked by /etc/systemd/system/default.target, for example /usr/lib/systemd/system/multi-user.target, to determine the default system target.

GRUB 2 can load many operating systems in addition to Oracle Linux and it can chain-load proprietary operating systems. GRUB 2 understands the formats of file systems and kernel executables, which allows it to load an arbitrary operating system without needing to know the exact location of the kernel on the boot device. GRUB 2 requires only the file name and drive partitions to load a kernel. One can configure this information by using the GRUB 2 menu or by entering it on the command line.

The following command generates the configuration file using the template scripts in /etc/grub.d and menu-configuration settings taken from the configuration file, /etc/default/grub.

grub2-mkconfig

The default menu entry is determined by the value of the GRUB_DEFAULT parameter in /etc/default/grub.

The following command sets the default entry for all subsequent reboots:

grub2-set-default

The following command sets the default entry for the next reboot only:

grub2-reboot

22 Reference Identifiers

The TOE verifies that the presented identifier matches the reference identifier according to RFC 6125.
The reference identifiers that are supported are the following:

- DNS host name or IP address can be set in the Common Name field
- DNS host name and URI names can be set in the SAN field

The TOE will verify the above identifiers with the presented certificates to ensure that it matches.
Wild cards are supported and certificate pinning is unsupported.

23 Using KVM with Oracle Linux

The Kernel-based Virtual Machine (KVM) feature provides a set of modules that enable you to use the Oracle Linux kernel as a hypervisor. KVM supports x86 and aarch64 processor architectures.

KVM is built into the Oracle Linux Unbreakable Enterprise Kernel (UEK) release by default. On an Oracle Linux system, one can run the following command to verify which modules are loaded in the kernel:

```
# lsmod | grep kvm
kvm_intel  167936  0
kvm  516096  1 kvm_intel
```

The KVM modules loaded may vary depending on your processor family and architecture.

Oracle Linux provides several virtualization packages that enable you work with KVM. One can use the **yum** command to install whichever virtualization package as required.

The following packages are the minimum required for a virtualization host:

libvirt: This package provides an interface to KVM, as well as the libvirtd daemon for managing guest virtual machines.

qemu-kvm: This package installs the QEMU emulator that performs hardware virtualization so that guests can access host CPU and other resources.

As an alternative to installing virtualization packages individually, one can install virtualization package groups.

The Virtualization Host package group contains the minimum set of packages required for a virtualization host. If your Oracle Linux system includes a GUI environment, one can also choose to install the Virtualization Client package group.

Use the following command to determine which packages are included in a group:

```
# yum groupinfo "Virtualization Host"
```

23.1 Installation of virtualization host with Oracle Linux graphical installation program

1. Boot the Oracle Linux installation media and proceed to the Software Selection screen.
2. Select one of the following virtualization host types:

   Minimum Virtualization Host

   a. Select Virtualization Host in the Base Environment section.
   b. Select Virtualization Host in the Add-ons for Selected Environment section.

   Virtualization Host with GUI

c. Select Server with GUI in the Base Environment section.
d. Select the following package groups in the Add-ons for Selected Environment section:
    - Virtualization Client
    - Virtualization Hypervisor
    - Virtualization Tools
3. Follow the prompts to complete the installation.

23.2 Installation of virtualization hosts by individual packages or package groups

Specify virtualization packages individually, as in the following example:

```
%packages
libvirt
qemu-kvm
```

Specify the appropriate package groups for the installation type in the %packages section of the kickstart file by using the @GroupID format:

Minimum Virtualization Host

```
%packages
@virtualization-hypervisor
@virtualization-tools
# The following group is optional. Uncomment the line to include it:
#@virtualization-platform
```

Virtualization Host with GUI

```
%packages
@virtualization-hypervisor
@virtualization-tools
@virtualization-client
```

23.3 Installation of virtualization package on an existing Linux system

1. Log in as the root user on the target Oracle Linux system.
2. Ensure that your system has the appropriate Yum repository or ULN channel enabled for the virtualization package versions that you wish to install.
3. Run the yum update command to ensure the system is up to date.
4. Install virtualization packages by doing one of the following:

- Use the yum install command to manually install virtualization packages, for example:

  # yum install libvirt qemu-kvm

- Use the yum groupinstall command to install a virtualization package group, for example:

  # yum groupinstall "Virtualization Host"

23.4 Upgrading Virtualization packages

Virtualization packages can be updated by using the standard **yum update** command.

For example, one would update to the latest supported virtualization packages that are available in the ol7_kvm_utils repository as follows:

yum --disablerepo="*" --enablerepo="ol7_kvm_utils" update

If one wants to downgrade packages to a version in an alternate repository or channel, for example, to downgrade from the virtualization packages in the ol7_kvm_utils repository to the version of the same packages in the ol7_latest repository, you must first remove the existing packages before installing the packages from the alternate repository:

yum remove libvirt* qemu*
yum --disablerepo="*" --enablerepo="ol7_latest" install libvirt qemu-kvm

24 References

Document Name	Date
Oracle Linux 7 Administrator's Guide - E54669-78	October, 2020
Oracle Linux 7 Installation Guide - E54695-26	October 2020
Oracle Linux 7 Security Guide - E54670-27	December 2020
Oracle Linux Security Target v3.9	May 27, 2021