DUMPSBOSS.

Oracle Linux 8 Advanced System Administration

Oracle 1z0-106

Version Demo

Total Demo Questions: 7

Total Premium Questions: 71

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Automate System Administration Tasks	26
Topic 2, Configure Advanced Networking	11
Topic 3, Configure Advanced Storage	10
Topic 4, Configure System Security	18
Topic 5, Configure Virtualization	4
Topic 6, Configure High Availability	1
Topic 7, Mix Questions	1
Total	71

QUESTION NO: 1

Which is true about the /etc/sysconfig directory in an Oracle Linux 8 system?

A. It is used to access device and device driver information.

B. Files in this directory hierarchy contain information about running processes.

C. Its contents depend on the packages installed on the system.

D. Files in this directory hierarchy contain information about system hardware.

ANSWER: C

Explanation:

Its contents depend on the packages installed on the system. In Oracle Linux 8, `/etc/sysconfig` is conventionally used for system-wide configuration files associated with particular services, daemons, and utilities. A package can install one or more files there to provide settings that affect how its associated software starts or operates. Consequently, the precise set of files is not fixed across all Oracle Linux installations: a minimal system and a system with additional networking, firewall, authentication, performance-tuning, or service packages can have different `/etc/sysconfig` contents. Administrators commonly inspect or modify applicable files in this directory when configuring installed software, subject to the configuration method documented by the specific package. This purpose follows the Filesystem Hierarchy Standard convention for `/etc/sysconfig`, which describes it as containing host-specific system configuration data. Oracle Linux is binary compatible with Red Hat Enterprise Linux, so the documented RHEL 8 filesystem layout and package configuration conventions apply to Oracle Linux 8 administration as well. See the Filesystem Hierarchy Standard: /etc and Oracle's Oracle Linux 8 documentation.

QUESTION NO: 2

Which two statements are true about systemd timer units?

A. A timer unit normally activates a service unit when its configured time elapses.

B. systemctl list-timers displays scheduled timer units and their elapse information.

C. A timer unit can execute shell commands only when they are written in the timer unit file.

D. Enabling a timer automatically enables every service unit on the host.

E. Persistent=true prevents a timer from running after a missed scheduled time.

ANSWER: A B

Explanation:

A timer unit activates another unit, normally a service unit with the same base name. For example, `cleanup.timer` commonly activates `cleanup.service` when the timer elapses. The timer defines when activation occurs; the service defines the work that is performed.

`systemctl list-timers` displays timer units known to systemd, including their next and previous elapse times where available. A timer with `Persistent=true` can cause a missed calendar-based activation to occur after the timer becomes active again, which is useful when the system was powered off at the scheduled time.

QUESTION NO: 3

Examine this command:

ssh -L 5011:127.0.0.1:80 bob@10.10.2.20 -f sleep 30

Which two are true upon execution?

A. A reverse tunnel is created back to the local host on port 80.

B. A local port forward is created between client and server.

C. A socket remains open for 30 minutes unless a connection is established.

D. A web server is listening on port 5011.

E. An SSH connection process is forked to the background.

ANSWER: B E

Explanation:

The command creates a local port forward between the SSH client and the SSH server. In the `-L 5011:127.0.0.1:80` specification, `5011` is the local listening port on the machine where the command is run. Connections made to that local port are carried through the encrypted SSH connection to `10.10.2.20`; from there, SSH connects to `127.0.0.1` on the remote host at TCP port `80`. Thus, the loopback address identifies a service as seen from the remote SSH server, rather than a destination on the local client. This is the standard behavior of SSH local forwarding, documented in the ssh(1) manual.

An SSH connection process is also forked to the background. The `-f` option directs SSH to go to the background immediately before command execution, after authentication and forwarding setup. The remote command is `sleep 30`, so it keeps the session available for approximately 30 seconds; during that interval, the local forwarding listener can accept connections. This use of background forwarding is described in the OpenSSH manual pages.

QUESTION NO: 4

A host has two active NetworkManager connection profiles. The profile named `backend` has a static IPv4 address and a gateway, but it must be used only for traffic to its directly connected subnet. The default route must continue to be supplied only by the `frontend` profile.

Which command changes the `backend` profile accordingly while retaining its static IPv4 address configuration?

A. `nmcli connection modify backend ipv4.never-default yes`

B. `nmcli connection modify backend ipv4.gateway ""`

C. `nmcli connection modify backend ipv4.method auto`

D. `nmcli connection modify backend ipv4.route-metric 500`

ANSWER: A

Explanation:

`nmcli connection modify backend ipv4.never-default yes` prevents NetworkManager from using the `backend` profile to provide an IPv4 default route. The profile can retain its configured IPv4 address and routes for its directly connected network, while the default route remains associated with `frontend`.

Removing the gateway also prevents a default route, but it unnecessarily changes the profile's gateway configuration. Setting `ipv4.method` to `auto` replaces the static-addressing method, and changing route metrics still permits the connection to install a default route.

QUESTION NO: 5

Examine this command:

cryptsetup luksOpen /dev/xvdd1 cryptfs

What happens upon execution?

A. It creates the /dev/mapper/cryptfs device mapping file.

B. It creates the LUKS partition on /dev/xvdd1.

C. It creates the /dev/mapper/xvdd1 device mapping file.

D. It creates the /dev/mapper/xvdd1/cryptfs device mapping file.

E. It creates the /dev/mapper/xvdd1-cryptfs device mapping file.

ANSWER: A

Explanation:

The command `cryptsetup luksOpen /dev/xvdd1 cryptfs` activates an existing LUKS-encrypted block device. `/dev/xvdd1` identifies the device containing the LUKS header and encrypted data, while `cryptfs` is the name assigned to the newly activated device-mapper mapping. During activation, cryptsetup requests the required authentication material, such as a passphrase or key file, and uses it to unlock the encryption key stored in the LUKS metadata. Once successfully unlocked, device mapper exposes the decrypted view through `/dev/mapper/cryptfs`. This mapped device can then be treated as an ordinary block device: for example, it can contain a file system that may be mounted, or it can be used by LVM. The operation does not format or initialize the source device; it makes the already encrypted contents accessible through the specified mapping name. The mapping remains available until it is explicitly closed, typically with `cryptsetup luksClose cryptfs`, or until shutdown. Oracle Linux storage documentation describes device-mapper-based encrypted storage administration, and the cryptsetup manual documents the `luksOpen` action and its mapping-name argument. See Oracle Linux 8: Managing Storage Devices and the cryptsetup manual page.

QUESTION NO: 6

Examine this command and output:

mdadm --detail /dev/md0

/dev/md0:

Version: 1.2

Creation Time: Tue Oct 27 16:53:38 2020

Raid Level: raid5

Array Size: 207872 (203.03 MiB 212.86 MB)

Used Dev Size: 103936 (101.52 MiB 106.43 MB)

Raid Devices: 3

Total Devices: 3

Persistence : Superblock is persistent

Update Time: Tue Oct 27 16:53:38 2020

State: clean, degraded, recovering

Active Devices: 2

Working Devices: 3

Failed Devices: 0

Spare Devices: 1

Layout: left-symmetric

Chunk Size: 512K

Rebuild Status: 60% complete

Name: ol8.example.com:0 (local to host ol8.example.com)

UUID: 70f8bd2f:0505d92d:750a781e:c224508d

Events: 66

Number Major Minor RaidDevice State

0 8 49 0 active sync /dev/sdd1

1 8 65 1 active sync /dev/sde1

3 8 81 2 spare rebuilding /dev/sdf1

Which two are true?

A. A RAID device failed and has returned to normal operating status.

B. The RAID set read and write performance is currently sub-optimal.

C. A new RAID device was just added to replace a failed one.

D. An extra device was added to this RAID set to increase its size.

E. Only write performance is currently sub-optimal on this RAID set.

ANSWER: B C

Explanation:

The RAID set read and write performance is currently sub-optimal because the RAID-5 array reports `degraded, recovering` and a rebuild is only 60 percent complete. Only two members are presently active, so normal RAID-5 redundancy has not yet been restored. During recovery, md processes data from the active members and writes reconstructed stripe contents to the replacement member. This reconstruction activity competes with ordinary I/O for disk bandwidth and controller resources. In addition, RAID-5 writes continue to require parity handling while the array operates with a missing original member, so workloads should be expected to experience reduced performance until synchronization finishes.

A new RAID device was just added to replace a failed one is supported by the member listing: `/dev/sdf1` has the state `spare rebuilding`. The array has three configured RAID devices but only two active devices, along with one spare that is being promoted and synchronized into RAID-device slot 2. The previous member is no longer an active array member; the replacement is rebuilding its data and parity contents. Once rebuilding completes, the replacement becomes active and the array returns to its normal protected state. Oracle Linux documents software RAID administration with `mdadm` in its Managing Storage Devices guide. The mdadm manual page also describes rebuilding and spare-device behavior.

QUESTION NO: 7

Examine this segment of /etc/rsyslog.conf:

Log all kernel messages to the console.

Logging much else clutters up the screen.

#kern.* /dev/console

Log anything (except mail) of level info or higher.

Don’t log private authentication messages!

*.info;mail.none;authpriv.none;cron.none /var/log/messages

The authpriv file has restricted access.

authpriv.* /var/log/secure

Log all the mail messages in one place.

mail.* -/var/log/maillog

Log cron stuff

cron.* /var/log/cron

Everybody gets emergency messages

.emerg :omusrmsg:

Now examine this log output:

less

Nov 9 20:32:16 server02 sudo[4570]: pam_unix(sudo:session): session opened for user opc (uid=0)

Nov 9 20:32:17 server02 sudo[4570]: pam_unix(sudo:session): session closed

Nov 9 20:32:24 server02 unix_chkpwd[4661]: password check failed for user

Nov 9 20:32:24 server02 su[4581]: pam_unix(su:auth): authentication failed; logname= uid=1000 euid=0 tty=pts/0 ruser=opc rhost= user=root

Which setting enabled the reporting of this log file output?

A. authpriv.* /var/log/secure

B. *.emerg *

C. *.info;mail.none;authpriv.none;cron.none /var/log/messages

D. #kern.* /dev/sssd/sssd.log

E. cron.* /var/log/cron

ANSWER: A

Explanation:

`authpriv.* /var/log/secure` enables this reporting because the displayed records are private authentication and authorization events. The `sudo` session-open and session-close messages, the failed password check from `unix_chkpwd`, and the PAM authentication failure generated during `su` activity are emitted through the `authpriv` syslog facility. In rsyslog selector syntax, `authpriv.*` selects every severity level from that facility, so both successful session activity and failed authentication attempts are eligible for logging. The action portion sends the selected records to `/var/log/secure`, the conventional protected authentication log location on Oracle Linux and other Red Hat-compatible systems.

The configuration also deliberately keeps these records out of the general messages log by specifying `authpriv.none` on that rule. This separation is important because authentication data can include sensitive operational details, such as account names, privileged command activity, terminal information, and failed-access events. Oracle Linux documents rsyslog as the system logging service and describes use of facility-and-priority selectors to route messages to specific destinations. See the Oracle Linux 8 logging documentation and the Oracle Linux rsyslog configuration documentation.