PassQueen
a.k.a dumpsqueen.com

SUSE Certified Administrator in Enterprise Linux 15

SUSE sca_sles15

Version Demo

Total Demo Questions: 10

Total Premium Questions: 107

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Installation	4
Topic 2, Software Management	9
Topic 3, File System Administration	20
Topic 4, System Administration	50
Topic 5, Network Administration	9
Topic 6, Security Administration	15
Total	107

QUESTION NO: 1

Which command will take a foreground process and suspend it. and then place it in the background as a stopped job?

A. bg

B. ctri+&

C. Ctrl+Z

D. fg--bg--stop

E. &&

ANSWER: C

Explanation:

`Ctrl+Z` is the terminal key sequence used by shell job control to suspend the currently running foreground process. It sends a terminal stop signal to the foreground process group, ordinarily `SIGTSTP`. The shell then regains control of the terminal, returns a prompt, and records the process in its job table with a stopped status. This is why a job suspended with `Ctrl+Z` is displayed as stopped when using the `jobs` command. At that point, the process is no longer the active foreground process and can subsequently be managed as a shell job. If continued execution in the background is required, the user can issue `bg`, which resumes the stopped job without assigning it the terminal foreground. The distinction is important: `Ctrl+Z` performs the suspension and creates the stopped job state described in the question. SUSE documents shell job control, including stopping a foreground job with `Ctrl+Z` and continuing it with `bg`, in its administration documentation. See SUSE Linux Enterprise Server: Bash and job control and the GNU Bash manual on job-control basics.

QUESTION NO: 2

Which statements regarding subvolumes in a BtrFS filesystem are true (Choose two)

A. Every subvolume on a BtrFS filesystem must have the same quota if quotas are implemented.

B. Even though subvolumes in BtrFS appear to be subdirectories, subvolumes can be mounted separately using `-o subvol=`.

C. Every subvolume on a BtrFS fiLesystem must have a name that begins with the "@" character.

D. Multiple virtual filesystems can be created inside the BtrFS filesystem.

ANSWER: B D

Explanation:

Btrfs subvolumes are independently addressable filesystem trees within one Btrfs filesystem. Although a subvolume is visible from the top-level Btrfs tree much like a directory, it is not merely an ordinary directory: it has its own subvolume ID and can be selected as the root of a separate mount. To mount one explicitly, the Btrfs mount option is `-o subvol=<path>` (or, alternatively, `subvolid=<id>`). This capability lets an administrator mount separate trees such as a system root, home-directory tree, or application-data tree at different mount points while they remain on the same underlying Btrfs filesystem.

Accordingly, Btrfs supports multiple virtual filesystems inside the same filesystem through subvolumes. Each subvolume provides an independent tree and is a useful boundary for administration, including selective mounting and snapshot operations. SUSE uses this Btrfs design extensively; for example, a standard SLES installation can organize system data into subvolumes to support Snapper snapshots and rollback workflows. The subvolume mechanism is therefore central to Btrfs layout management rather than being a naming convention or a requirement that all trees use one common quota value. See the SUSE documentation on Snapper and Btrfs snapshots and the upstream Btrfs subvolume documentation.

QUESTION NO: 3

Which statements are true regarding a daemon? (Choose three)

A. A daemon process is associated to the system terminal.

B. A daemon process is associated to the user that launched the process.

C. A daemon will wait for some event to occur to trigger action on the part of the daemon.

D. A daemon process is launched by the system.

E. A daemon process is not associated with a terminal or a graphical environment.

F. The user that launched the daemon will control what the daemon does based on the user's input.

ANSWER: C D E

Explanation:

A daemon is a background service process that provides functionality without interactive user control. “A daemon will wait for some event to occur to trigger action on the part of the daemon” is true because daemons commonly remain running while waiting for requests, timers, network connections, device activity, or other events. For example, a network service daemon waits for incoming client connections before handling them.

“A daemon process is launched by the system” is also true in the normal SLES service-management model. On SUSE Linux Enterprise Server, `systemd` is the system and service manager; it starts, supervises, stops, and can restart daemon processes according to unit definitions and configured dependencies. A daemon may also be started on demand by systemd, rather than necessarily at boot.

“A daemon process is not associated with a terminal or a graphical environment” describes a defining daemon characteristic. A properly detached daemon runs independently of a controlling terminal and is intended to continue providing its service in the background. This permits the service to operate regardless of whether a user is logged in or a graphical session exists. See the SUSE systemd basics documentation and the systemd.service manual.

QUESTION NO: 4

What command would be used to add new physical volumes to an existing volume group?

A. vgextend

B. vSadd

C. pvmodify

D. pvadd

E. pvextend

F. vgmodify

ANSWER: A

Explanation:

`vgextend` is the LVM command for adding one or more initialized physical volumes to an existing volume group. Its normal syntax is `vgextend VolumeGroupName PhysicalVolumePath`, for example `vgextend vgdata /dev/sdb1`. Before this step, the new disk or partition must have been initialized as an LVM physical volume, typically with `pvcreate /dev/sdb1`. Extending the volume group incorporates the physical volume’s extents into the group’s available storage pool. That newly available capacity can then be assigned to an existing logical volume or used when creating a new logical volume.

This distinction is important in LVM administration: physical volumes provide the underlying storage, volume groups aggregate that storage, and logical volumes consume capacity from the group. SUSE’s storage documentation describes volume groups as collections of physical volumes and documents the process of expanding available LVM storage. The

`vgextend` manual page specifically defines the command as adding physical volumes to a volume group. See the SUSE LVM administration documentation and the vgextend manual page.

QUESTION NO: 5

What are the basic components of RPM? (Choose two)

A. RPM Compiler

B. RPM Editor

C. RPM Packager

D. RPM Database

E. RPM Package Manager

ANSWER: D E

Explanation:

RPM Package Manager and RPM Database are the two fundamental parts of the RPM packaging system. The RPM Package Manager is the `rpm` utility and its associated mechanisms for working with RPM packages. It performs core local package operations, including installing, upgrading, removing, querying, and verifying packages. These operations allow an administrator to manage individual package files and to inspect package metadata directly on an Enterprise Linux system.

The RPM Database supplies the persistent inventory required for those operations. It records which packages are installed and maintains their package metadata, such as package names, versions, release information, file lists, dependencies, and installation state. When an administrator queries an installed package or verifies package-owned files, RPM uses this database rather than merely examining package files that may be available in repositories. Together, the package-management utility and its installed-package database provide the essential RPM functionality used by SUSE Linux Enterprise Server. SUSE documentation describes RPM as the low-level package-management tool and explains its package query and verification capabilities; the upstream RPM documentation also describes the database-backed package-management model. See the SUSE Linux Enterprise Server package-management documentation and the RPM documentation.

QUESTION NO: 6

You have just added a new directory to be exported via NFS. what command should you run next to make that directory available?

A. /usr/sbin/rpc.mountd -restart

B. /bin/nfsd - reload

C. /etcyinit.d/nfsd -reload

D. systemctl restart nfsserver.service

E. systmd -reload nfsserver.target

F. exportfs -r

ANSWER: F

Explanation:

`exportfs -r` is the appropriate command after adding an export definition to `/etc/exports`. The `exportfs` utility manages the list of directories made available by the NFS server. Its `-r` option re-exports all entries by synchronizing the active export table with the definitions currently present in `/etc/exports` and files included from it. This applies the newly

added directory without requiring a full restart of the NFS server, which is preferable because existing NFS service and client mounts can continue operating while the export configuration is refreshed. Administrators commonly use this command after editing export definitions, then can inspect the active exports with `exportfs -v` if verification is needed. SUSE documents managing exported directories through `/etc/exports` and the `exportfs` command in its NFS server guidance. See the SUSE Linux Enterprise Server NFS documentation and the exportfs manual page for the reload behavior and available options.

QUESTION NO: 7

firewalld maintains two separate configurations. What are they named? (Choose two)

A. Runtime

B. Temporary

C. Production

D. Permanent

E. Static

F. Testing

G. Fixed

ANSWER: A D

Explanation:

firewalld maintains a **runtime** configuration and a **permanent** configuration. The runtime configuration is the active ruleset currently applied by the running firewalld daemon. Commands issued with `firewall-cmd` normally modify this active configuration immediately, making it useful for applying and validating a change without first requiring a service reload or system restart.

The permanent configuration is the saved configuration from which firewalld obtains its rules when the service starts or reloads. To save a setting persistently, administrators use the `--permanent` option, then reload firewalld when they want the saved settings to replace the active runtime settings. This separation supports safer administration: a rule can be tested at runtime before it is committed as a persistent configuration. SUSE documents firewalld administration and the distinction between runtime and permanent settings in its Security and Hardening Guide. See SUSE Security and Hardening Guide: Masquerading and firewalld and the firewall-cmd manual page.

QUESTION NO: 8

By default, what does the VNC Client and the VNC server daemon use for secure communications?

A. On SLE 15 the VNC Client assumes the server is set up as a C

B. Both the VNC Client and the VNC server assume you have access to a valid 3rd party CA like Verisign.

C. The VNC server prompts the VNC Client for a password that was configured when the VNC server was installed.

D. VNC is based on the Telnet protocol so encrypted communication between the VNC Client and Server is not possible.

E. A self-signed SSL certificate

ANSWER: E

Explanation:

A self-signed SSL certificate is used by default to establish secure communication between the VNC client and the VNC server daemon. The certificate provides the TLS encryption material needed to protect the remote desktop session while it

traverses the network. Because it is self-signed rather than issued by a public certificate authority, a client may display a trust warning on the first connection. The administrator or user must then verify and accept the certificate fingerprint when appropriate. This default is practical for VNC deployments because it enables encrypted communication without requiring an organization to obtain, install, and maintain a third-party CA-issued certificate for every VNC server. In environments with stricter identity-validation requirements, administrators can deploy certificates signed by an internal or public CA instead, allowing clients to validate server identity through an established trust chain. SUSE documents VNC configuration and secure remote graphical access in its VNC documentation. TigerVNC also documents the TLS-related security types and certificate configuration used for encrypted VNC connections in its vncserver manual.

QUESTION NO: 9

The openSUSE Leap version of Linux is designed for which types of user environments? (Choose two)

A. Environments requiring High-Availability

B. Desktop power users

C. Users requiring only secure terminal access

D. Data Center environments

E. Developers

ANSWER: B E

Explanation:

Desktop power users and **Developers** are the intended environments for openSUSE Leap. Leap is a community-developed, stable general-purpose distribution that provides a polished desktop experience while also including extensive development tooling, current programming environments, containers, virtualization support, and administration utilities. This combination makes it appropriate for technically proficient workstation users who want flexibility and control, as well as for developers who need a dependable platform for creating, testing, and building software.

Leap emphasizes a balanced release model: it offers a stable base and benefits from shared technology with SUSE Linux Enterprise, while remaining a distribution suited to individual users, contributors, and development work. Its broad software repositories and YaST administration tools further support productive workstation and development use cases. The openSUSE Project describes Leap as a stable distribution for a range of users and identifies its connection to SUSE Linux Enterprise technology. See the openSUSE Leap download and overview page and the openSUSE Leap portal for release and usage information.

QUESTION NO: 10

Which command locks the password of the local user account tux?

A. passwd -l tux

B. passwd -d tux

C. userdel -l tux

D. usermod -e tux

E. loginctl lock-user tux

ANSWER: A

Explanation:

`passwd -l tux` locks the password for the specified local account. It modifies the password entry so password authentication cannot use the existing password hash. Locking a password does not necessarily prevent all access to the

account; for example, SSH public-key authentication can still be possible if separately configured and permitted. To reverse the operation, use `passwd -u tux`. See the passwd(1) manual page and the SUSE user-management documentation.