PassQueen
a.k.a dumpsqueen.com

SUSE Certified Linux Administrator 12

Novell 050-733

Version Demo

Total Demo Questions: 19

Total Premium Questions: 191

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

Which statements about a RAID 1 array are correct? (Choose two.)

A. Data is mirrored across member devices.

B. A two-disk array can survive failure of one member disk.

C. The usable capacity of two equal disks is the sum of both disk capacities.

D. RAID 1 uses distributed parity blocks.

E. RAID 1 eliminates the need for backups.

ANSWER: A B

Explanation:

RAID 1 stores mirrored copies of data on its member devices. Because the same data is present on more than one member, an array with two healthy members can continue operating after one member device fails.

The usable capacity of a two-device RAID 1 array is approximately the capacity of one member device, not the sum of both devices. Mirroring improves availability but does not replace backups: accidental deletion, filesystem corruption, and application-level data loss can be replicated to all mirrors. See the Linux kernel MD documentation.

QUESTION NO: 2

Which statements about the LVM components are correct? (Choose two.)

A. Logical volumes are grouped in a master group.

B. A volume group always consists of one physical volume.

C. A volume group can be reduced in size by removing physical volumes.

D. The operating system accesses the volume groups like conventional physical partitions.

E. A volume group is part of a logical volume.

F. A physical volume can be a partition or an entire hard disk.

ANSWER: C F

Explanation:

A volume group can be reduced in size by removing physical volumes. In LVM, a volume group is a storage pool built from one or more physical volumes. When a physical volume is no longer needed in that pool, its allocated extents can first be relocated to other physical volumes in the volume group, typically with `pvmove`, and the physical volume can then be removed with `vgreduce`. This permits controlled reduction of the storage pool while preserving the logical volumes it contains.

A physical volume can be a partition or an entire hard disk. LVM initializes a block device as a physical volume, and that device may be a whole disk, such as `/dev/sdb`, or a disk partition, such as `/dev/sdb1`. Physical volumes provide the underlying allocatable storage from which volume groups are assembled; logical volumes are then created from the available extents in those groups. This layered design is what enables flexible allocation and later storage reconfiguration. See the SUSE LVM administration documentation at SUSE LVM configuration and the vgreduce manual page.

QUESTION NO: 3

What is the output of the echo '$HOME' command when user geeko (home directory /home/geeko) enters it at the bash shell prompt?

A. geeko

B. /home/geeko

C. $HOME: No such file or directory

D. $HOME

ANSWER: D

Explanation:

The correct output is `$HOME`. In Bash, text enclosed in single quotation marks is preserved literally. The shell does not perform parameter expansion on `$HOME` when it appears inside single quotes, so the value of the user's HOME environment variable is not substituted. The command therefore passes one literal argument, `$HOME`, to the `echo` command, which prints those characters followed by a newline. Although geeko's home directory establishes that the HOME variable would ordinarily contain `/home/geeko`, that value matters only when Bash is permitted to expand the parameter—for example, with `echo "$HOME"` or `echo $HOME`. Single quotes are specifically used when a command must prevent interpretation of shell metacharacters and variable references. Bash's quoting rules state that single quotes preserve the literal value of every character between the quote delimiters. See the GNU Bash Reference Manual: Single Quotes and SUSE's Bash and Bash Scripts documentation for shell quoting and expansion behavior.

QUESTION NO: 4

How long does "Long Term Service Pack Support (LTSS) for SLES" extend the span of getting updates for SLES packages belonging to a specific Service pack?

A. 3 years

B. 1 year

C. 5 years

D. 7 years

E. 10 years

ANSWER: A

Explanation:

3 years is correct. Long Term Service Pack Support is an add-on offering for SUSE Linux Enterprise Server that extends the update and maintenance period for a particular service pack after its regular support period ends. The extension is measured per service pack rather than extending the entire major SLES release lifecycle by that amount. During this additional three-year period, customers with the appropriate subscription can continue to receive maintenance updates, including security fixes and selected bug fixes, for packages associated with that service pack. This support model is designed for environments in which an upgrade to a newer service pack cannot be completed before normal service-pack maintenance ends, such as systems subject to application certification, change-control windows, or operational constraints. LTSS therefore provides a controlled extra maintenance window while organizations plan and validate their migration to a supported service pack. SUSE describes LTSS as providing an additional three years of support for eligible SUSE Linux Enterprise service packs. Details of the offering are available on the SUSE Long Term Service Pack Support page and in the SLES support and lifecycle documentation.

QUESTION NO: 5

Which group name or group ID is stored in /etc/passwd?

A. Primary group.

B. Secondary group.

C. Effective group.

D. No group ID is stored there.

ANSWER: A

Explanation:

The primary group is correct. Each local user account has an entry in `/etc/passwd` whose fields include the user name, password placeholder, numeric user ID, numeric primary group ID, user-information field, home directory, and login shell. The group-related field in this file is therefore the account's primary GID, rather than a group name. The system resolves that numeric GID to a group name by consulting the configured group database, commonly the local `/etc/group` file or another Name Service Switch source such as LDAP.

This association establishes the user's default group membership at login. New files created by the user normally receive this primary group as their group ownership, subject to directory settings such as the set-group-ID permission bit and to any access-control policies in effect. A user can belong to additional groups, but those memberships are recorded through group database entries rather than in the primary-GID field of the user's `/etc/passwd` record. The Linux `passwd(5)` manual documents the GID field as "the numeric primary group ID of this user," and the SUSE documentation describes local users and groups as separate account databases. See passwd(5) — Linux manual page and SUSE documentation on managing users and groups.

QUESTION NO: 6

Which statement is not true about symbolic links?

A. Symbolic links can point to files on a different file system.

B. Symbolic links can be created for directories.

C. Symbolic links always have the same size as the original file.

ANSWER: C

Explanation:

"Symbolic links always have the same size as the original file." is not true. A symbolic link is a distinct filesystem object whose contents are the pathname used to reach its target, rather than a second copy of the target file's data. Consequently, the link's apparent size is normally related to the number of characters in that stored pathname. For example, a link containing the target path `../data/report` has a size corresponding to that path string, regardless of whether the target is an empty file, a multi-gigabyte database, or a directory. The target may also be removed or moved, leaving a dangling symbolic link; the link still retains its own pathname and metadata. On Linux, commands such as `ls -l` display the symbolic link and its target, while `stat` can inspect link metadata; using `stat -L` follows the link and reports information about the target instead. This distinction is fundamental when administering filesystems and troubleshooting links. The Linux `symlink(7)` manual documents that a symbolic link contains a text string automatically interpreted as a pathname, and the GNU Coreutils documentation describes how `stat` treats links.

References: Linux symlink(7) manual page; GNU Coreutils stat invocation.

QUESTION NO: 7

Within less, how do you search for a string downwards from where the cursor is positioned?

A. :string

B. ?string

C. /string

D. =string

ANSWER: C

Explanation:

Within `less`, enter `/string` and press Enter to search forward (downward) from the current cursor or viewing position for the specified text. The slash introduces a forward search pattern, and `less` moves to the next matching line in the file. This is useful for quickly locating a command, configuration setting, error message, or other text while examining output interactively. Search patterns in `less` are regular expressions, so the text after the slash can be a simple literal string or a more expressive pattern when needed. After a search has been performed, pressing `n` repeats the search in the same forward direction, allowing successive matches to be located efficiently; `N` repeats it in the opposite direction. The current search position is therefore the natural starting point for finding later occurrences rather than restarting at the beginning of the file. The GNU `less` manual documents the slash command as a forward search and describes the follow-up search navigation commands. See the less documentation and FAQ and the less(1) manual page.

QUESTION NO: 8

What is described as follows:

If content of a block changes, the changed block is put beside the non-changed block.

A. Rollback

B. Journaling

C. Subvolume

D. Copy on write

ANSWER: D

Explanation:

Copy on write describes this behavior. In a copy-on-write filesystem such as Btrfs, existing data blocks are not overwritten in place when their contents are modified. Instead, the filesystem writes the altered data to a newly allocated block while retaining the original block unchanged. Metadata is then updated atomically to reference the newly written block. Thus, the old and new versions of the data can coexist until the older version is no longer needed.

This design is fundamental to Btrfs data consistency and enables efficient snapshots. A snapshot can initially share the same extents as its source subvolume without duplicating all data. When a block belonging to either the original subvolume or the snapshot is changed, Btrfs writes the changed version elsewhere, preserving the version referenced by the other subvolume. This permits snapshots to retain a point-in-time view while consuming additional disk space only for changed data. Copy on write also helps maintain filesystem integrity because updates can be committed through new blocks rather than destructively modifying the only existing copy. See the SUSE Btrfs documentation and the Btrfs introduction.

QUESTION NO: 9

Where are the SuSEfirewall2 service definition files located?

A. /etc/sysconfig/SuSEfirewall2.d/

B. /etc/SuSEfirewall2/service/

C. /etc/SuSEfirewall2/config/

D. /etc/SuSEfirewall2/service/def/

E. /etc/sysconfig/SuSEfirewall2.d/services/

ANSWER: E

Explanation:

SuSEfirewall2 service definition files are stored in `/etc/sysconfig/SuSEfirewall2.d/services/`. This directory is part of the firewall's sysconfig configuration hierarchy and contains individual service definitions that describe the ports, protocols, and optional related settings associated with a named service. These definitions allow an administrator to refer to a service by name when configuring firewall rules, instead of repeatedly entering protocol and port details. A locally created service definition is normally placed in this directory using the expected service-file format; SuSEfirewall2 then reads it along with its main configuration when firewall settings are applied or reloaded. The parent directory `/etc/sysconfig/SuSEfirewall2.d/` contains the service-definition directory, but it is not the specific location requested for the definition files themselves. SUSE documents the use of custom service definitions under the `services` subdirectory in its firewall configuration guidance. See the SUSE Linux Enterprise Server 12 SP5 Security and Hardening Guide: Firewall and the SuSEfirewall2 manual page for the configuration layout and service-definition conventions.

QUESTION NO: 10

Which statements about exporting directories with NFS are correct? (Choose two.)

A. `/etc/exports` defines directories exported by an NFS server.

B. `exportfs -r` rereads and reapplies export definitions.

C. NFS exports are configured in `/etc/fstab` on the server.

D. An NFS export can only be mounted by the local host.

E. The `mount` command permanently creates server exports.

ANSWER: A B

Explanation:

The file `/etc/exports` defines directories that an NFS server makes available to clients and specifies client access rules and export options. After changing this file, `exportfs -r` re-exports the entries, causing the NFS server to reread the export table without requiring a reboot.

An export definition identifies a directory and the client hosts or networks allowed to mount it. See the exports(5) manual page and the exportfs(8) manual page.

QUESTION NO: 11

Which command displays block devices together with filesystem type, UUID, and mount point information?

A. `lsblk -f`

B. `df -f`

C. `mount -u`

D. `blkid -m`

ANSWER: A

Explanation:

`lsblk -f` lists block devices and adds filesystem-related columns, including filesystem type, label, UUID, and mount point where available. It is useful when identifying a partition for a persistent `/etc/fstab` entry, because UUID values are generally more stable than device names.

See the lsblk(8) manual page and the fstab(5) manual page.

QUESTION NO: 12

Which statements about the Linux boot process are correct? (Choose two.)

A. The initrd decompress itself.

B. The Kernel looks for and mounts the initrd.

C. The boot manager loads the kernel and initrd into the memory.

D. The boot manager starts the init process

E. The BIOS starts the kernel.

ANSWER: B C

Explanation:

During the early Linux boot sequence, the boot manager (such as GRUB) loads both the selected kernel image and its associated initial RAM disk image into RAM. It also supplies the kernel with the location and size of that initial RAM disk through the boot parameters. This early userspace image provides the drivers, tools, modules, and scripts that may be needed before the real root filesystem can be accessed—for example, storage-controller drivers, RAID or LVM activation, disk encryption support, or network configuration for a remote root filesystem.

After control transfers to the kernel, the kernel initializes its core subsystems and processes the initial RAM disk supplied by the boot manager. In the traditional initrd design, the kernel locates and mounts the initrd as an initial root filesystem, then executes its early userspace startup program. Modern systems commonly use initramfs instead, which the kernel unpacks into its in-memory root filesystem; the same overall responsibility remains: the boot manager places the early boot image in memory, and the kernel uses it to prepare access to the permanent root filesystem. The Linux kernel's documentation describes this early-root mechanism at Linux kernel: ramfs, rootfs and initramfs. SUSE's boot-process overview is also available in the SUSE Linux Enterprise Server documentation.

QUESTION NO: 13

Which commands can be used to display information about mounted filesystems and their available disk space? (Choose two.)

A. `df -h`

B. `findmnt`

C. `free -h`

D. `du -s /`

E. `mountpoint -q /`

ANSWER: A B

Explanation:

The `df` command reports filesystem disk-space usage, including total space, used space, available space, and mount points. Adding `-h` displays sizes using human-readable units.

The `findmnt` command displays mounted filesystems in a tree or table and can show the source device, target mount point, filesystem type, and mount options. Although it is particularly useful for mount relationships, it is not a replacement for `df` when free-space totals are required. See the df(1) manual page and the findmnt(8) manual page.

QUESTION NO: 14

Which statements about masking a systemd unit are correct? (Choose two.)

A. A masked unit cannot be started until it is unmasked.

B. `systemctl mask` prevents dependency-based starts.

C. Masking removes the RPM package that provides the unit.

D. A masked unit is automatically started at the next boot.

E. Masking is identical to stopping a running service.

ANSWER: A B

Explanation:

A masked unit cannot be started manually or as a dependency because systemd redirects the unit file to `/dev/null`. This is stronger than merely disabling a unit. A disabled unit can still be started manually or by another unit, whereas a masked unit is prevented from starting until it is unmasked.

The command `systemctl mask name.service` masks a unit, and `systemctl unmask name.service` removes that mask. See the systemctl documentation and systemd.unit documentation.

QUESTION NO: 15

Which statements about `systemctl enable` are correct? (Choose two.)

A. It configures a unit to start automatically at boot.

B. It starts the unit immediately in every case.

C. It uses installation information from the `[Install]` section.

D. It permanently prevents the unit from being started manually.

E. It deletes the unit file after reading it.

ANSWER: A C

Explanation:

`systemctl enable` configures a unit to start automatically by creating the links specified by the unit's `[Install]` section. It does not normally start the unit immediately; to both enable and start a unit in one operation, use `systemctl enable --now unit-name`.

The operation generally creates symbolic links under systemd configuration directories, such as target `.wants` directories. See the systemctl documentation and the systemd.unit documentation.

QUESTION NO: 16

Where are the SuSEfirewall2 service definition files located?

A. /etc/sysconfig/SuSEfirewall2.d/

B. /etc/SuSEfirewall2/service/

C. /etc/SuSEfirewall2/config/

D. /etc/SuSEfirewall2/service/def/

E. /etc/sysconfig/SuSEfirewall2.d/services/

ANSWER: E

Explanation:

The correct location is `/etc/sysconfig/SuSEfirewall2.d/services/`. SuSEfirewall2 uses its primary configuration under `/etc/sysconfig` and supports modular configuration through the `SuSEfirewall2.d` directory hierarchy. Service definitions are individual files in the `services` subdirectory. These definitions describe a named service and the network ports or protocols that belong to it, allowing that service to be referenced from the firewall configuration rather than requiring every port definition to be written directly in the main configuration file. This organization makes local firewall customizations clearer and helps preserve service-specific settings independently from the main SuSEfirewall2 configuration. The path must include the `services` component: `/etc/sysconfig/SuSEfirewall2.d/` is the parent directory, not the specific directory containing service definition files. SUSE documentation describes SuSEfirewall2 service definitions and the associated configuration layout in its firewall administration guidance. See the SUSE Linux Enterprise Server Security Guide and the openSUSE Security Guide firewall documentation.

QUESTION NO: 17

How do you start a YaST module (for instance, the sw_single module used to install software) directly from the command line? (Choose two.)

A. yast sw_single

B. yast2 sw_single

C. yast2 –m sw_single

D. yast - -start sw_single

E. yast2 - -module sw_single

ANSWER: A B

Explanation:

Both `yast sw_single` and `yast2 sw_single` directly invoke YaST's Software Management module by supplying its module name as the argument to the YaST command. YaST modules can be started independently rather than opening the full YaST control center and navigating through its categories. The `sw_single` module is YaST's software-management interface, used to search for, select, install, remove, and update packages according to the configured software repositories.

On SUSE systems, `yast` is the standard command-line launcher for YaST modules. The `yast2` command is also available for compatibility with the YaST2 naming convention and accepts the same module invocation form. The actual user interface selected can depend on the execution environment and available YaST UI packages, such as a text-based interface in a terminal or a graphical interface in a desktop session; however, the direct module syntax remains the module name following the launcher command. SUSE documentation describes invoking individual YaST modules from the command line and identifies YaST Software as the package-management tool. See the SUSE YaST documentation and the SUSE software-management documentation.

QUESTION NO: 18

Which command can be used to set the system time? (Choose two.)

A. time

B. timeofday

C. date

D. netdate

E. hctosys

ANSWER: C D

Explanation:

The `date` command can set the operating system's current clock directly when it is run with a time value and sufficient administrative privileges. For example, the administrator can supply a date-and-time string to update the system time immediately. This is useful for manual correction, although systems normally maintain accurate time automatically through a network time service. The `netdate` command can also set the system clock, obtaining the time from a network time source rather than requiring the administrator to enter it manually. This makes it appropriate when the machine needs to synchronize its clock from a configured time server. After setting the system clock, an administrator may separately update the hardware clock when needed so that the correct time persists across reboot; that persistence task is distinct from setting the running system time. See the SUSE time synchronization documentation and the date command manual page.

QUESTION NO: 19

You want to run the /usr/local/sbin/myscript script at 6.00 a.m. every Friday.

Which crontab entry is suitable for this?

A. 5 6 * * 0 /usr/local/sbin/myscript

B. 0 6 * * 5 /usr/local/sbin/myscript

C. 6 0 * 5 * /usr/local/sbin/myscript

D. * * 5 6 0 /usr/local/sbin/myscript

ANSWER: B

Explanation:

The suitable entry is `0 6 * * 5 /usr/local/sbin/myscript`. A standard crontab schedule has five time-and-date fields before the command, in this order: minute, hour, day of the month, month, and day of the week. Here, `0` in the minute field selects the start of the hour, and `6` in the hour field specifies 6:00 a.m. The asterisks for day of month and month mean that no restriction is applied to those fields, so the command can run in every month and on any calendar day that matches the weekday condition. Finally, `5` in the day-of-week field represents Friday in the cron weekday numbering used by cronie on SUSE Linux systems, where Sunday is represented by `0` or `7`. Consequently, cron invokes `/usr/local/sbin/myscript` once each Friday at exactly 06:00. This entry belongs in the appropriate user's crontab, edited with `crontab -e`, provided that the script is executable and has an appropriate interpreter line. See the crontab(5) manual page for the field order and weekday values, and the SUSE scheduling tasks documentation for cron usage.