PassQueen
a.k.a dumpsqueen.com

Red Hat Linux System Administration

RedHat RH133

Version Demo

Total Demo Questions: 38

Total Premium Questions: 389

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	125
Topic 2, Volume B	110
Topic 3, Volume C	134
Total	389

QUESTION NO: 2

John works as a Network Administrator for Perfect Solutions Inc. The company has a Linux-based network. John is working as a root user on the Linux operating system. He wants to print a file named print.txt. Which of the following commands can John use to accomplish the task? Each correct answer represents a complete solution. Choose all that apply.

A. lpr print.txt

B. lprm

C. lp print.txt

D. a2ps print.txt

ANSWER: A C D

Explanation:

`lpr print.txt` and `lp print.txt` are standard command-line interfaces for submitting a file as a print job through the configured printing system, commonly CUPS on Red Hat Enterprise Linux systems. With no printer explicitly selected, each command submits to the system's configured default destination; a specific destination can also be selected with the appropriate printer option. Both utilities accept a file pathname as an argument, so supplying `print.txt` directly creates the required print request. CUPS documents the `lp` interface as a mechanism for printing files and provides comparable compatibility support for traditional `lpr` workflows. See the CUPS lp manual page and the CUPS lpr manual page.

`a2ps print.txt` is also a complete printing solution when the a2ps package is installed. The utility formats text and other supported input for PostScript printing and sends the formatted result to the default printer unless another output destination is selected. This is useful when printing plain-text source or documentation with pagination and other formatting controls. Its behavior and printer-selection features are documented in the GNU a2ps manual.

QUESTION NO: 3

Which of the following commands can be used to power off a Linux computer?

Each correct answer represents a complete solution. Choose all that apply.

A. init 6

B. init 0

C. shutdown

D. halt

ANSWER: B C

Explanation:

On Red Hat Enterprise Linux systems using systemd, `init 0` requests runlevel 0. systemd maps this legacy runlevel request to `poweroff.target`, which stops services and shuts the machine down for power-off. This syntax remains available for compatibility, although Red Hat generally recommends using explicit `systemctl` commands for routine administration. The command `shutdown` is also a complete power-off command on current systemd-based systems: when no action is specified, its default action is to power off the system, and when no time is supplied it acts immediately. Administrators can additionally use options such as `-P` and an explicit time when they need to make the intended action or schedule clear. Both commands perform an orderly shutdown rather than abruptly removing power, allowing system services to stop and filesystems to be unmounted cleanly. Red Hat's system shutdown guidance is available in Configuring basic system settings. The documented systemd behavior and defaults for `shutdown` are described in the shutdown manual page.

QUESTION NO: 4

Which of the following is the only domain in Xen environment from which an administrator can send command to the Hypervisor?

A. DomainU

B. Domain2

C. Domain0

D. Domain1

ANSWER: C

Explanation:

Domain0 is correct because it is Xen's privileged management domain, commonly written as `dom0`. Xen starts this domain first during host boot and gives it the hardware access and elevated privileges required to control the Xen hypervisor and manage guest virtual machines. An administrator normally logs in to Domain0 and uses Xen management tooling, such as the `xl` toolstack, to create, configure, start, stop, pause, migrate, and inspect guest domains. Domain0 also owns the management responsibilities needed to support guests, including access to physical devices and the services through which unprivileged guest domains obtain virtual disk and network I/O.

The distinction is architectural: the hypervisor is a thin layer responsible for CPU, memory, and isolation, while Domain0 is the privileged operating-system environment that performs administrative control on its behalf. Therefore, commands intended to manage the hypervisor or other Xen domains originate from Domain0 rather than from ordinary guest domains. The Xen Project's xl manual page describes `xl` as the management tool for Xen domains, and the Xen overview manual describes the privileged role of domain zero in the Xen architecture.

QUESTION NO: 5

John works as a Network Administrator for Perfect Solutions Inc. The company has a Linux-based network. John is working as a root user on the Linux operating system. He wants to print a file named print.txt. Which of the following commands can John use to accomplish the task? Each correct answer represents a complete solution. Choose all that apply.

A. lpr print.txt

B. lprm

C. lp print.txt

D. a2ps

ANSWER: A C

Explanation:

The commands `lpr print.txt` and `lp print.txt` are complete ways to submit the named text file to the system's configured print service. On Red Hat Enterprise Linux systems using CUPS, both commands act as print-job submission clients. If no destination is specified with an option such as `-P` or `-d`, they submit the job to the configured default printer. The `lpr` command uses the traditional Berkeley printing interface, while `lp` uses the System V-style printing interface; CUPS supports both interfaces for compatibility. Supplying `print.txt` as the operand ensures that the requested file, rather than data entered through standard input, is submitted as the print job. The submitting account must have permission to use the selected printer, but being root does not change the basic command syntax. For CUPS command details, see the CUPS lpr manual page and the CUPS lp manual page.

QUESTION NO: 6

What is the PID for init?

A. 3

B. 0

C. 2

D. 1

ANSWER: D

Explanation:

The correct answer is **1**. The init process is assigned process ID (PID) 1 when the Linux kernel starts the first userspace process. On current Red Hat Enterprise Linux systems, `systemd` normally fulfills the init role, so it runs as PID 1. This process is responsible for bringing the system into its configured state, starting and supervising services, and coordinating shutdown or reboot operations.

PID 1 has a special role in Linux process management. It is the ancestor of normal userspace processes, either directly or indirectly, and it adopts orphaned processes whose original parent exits. Because it is foundational to system operation, terminating PID 1 is not an ordinary process-management action and can cause the system to panic or otherwise become unusable. Administrators can confirm the active init implementation and PID with commands such as `ps -p 1 -o pid,comm,args` or `systemctl status`.

Red Hat documents `systemd` as the system and service manager used by Red Hat Enterprise Linux, and the Linux kernel documentation describes the special handling associated with the init process. See Red Hat systemd documentation and Linux kernel init documentation.

QUESTION NO: 7

Which of the following statements about Logical Volume Management (LVM) is true?

Each correct answer represents a complete solution. Choose all that apply.

A. It can resize volume groups online by absorbing new physical volumes (PV) or ejecting existing ones.

B. It can resize logical volumes (LV) online by concatenating extents onto them or truncating extents from them.

C. It cannot suffer from internal fragmentation.

D. It can stripe whole or parts of logical volumes across multiple PVs, in a fashion similar to
R AID0.

ANSWER: A B D

Explanation:

LVM manages storage through physical volumes, volume groups, and logical volumes, allowing storage capacity to be allocated and adjusted flexibly. “It can resize volume groups online by absorbing new physical volumes (PV) or ejecting existing ones.” is correct because a volume group can be extended by adding initialized physical volumes. A physical volume can also be removed after any allocated extents have been relocated to other physical volumes, allowing the volume group’s available capacity to be changed without rebuilding the storage layout.

“It can resize logical volumes (LV) online by concatenating extents onto them or truncating extents from them.” is correct at the LVM layer: logical-volume extents can be added or removed. When a filesystem exists on the logical volume, its resizing requirements must also be respected; in particular, online growth depends on filesystem support.

“It can stripe whole or parts of logical volumes across multiple PVs, in a fashion similar to
R AID0.

” is correct because LVM supports striped logical volumes, distributing extents across multiple physical volumes to provide parallel I/O behavior. Red Hat documents volume-group and physical-volume management at Managing LVM physical volumes and logical-volume resizing at Extending logical volumes.

QUESTION NO: 8

What does MBR stand for?

A. Master boot record

B. Memory buffer register

C. Master buffer register

D. Memory boot record

ANSWER: A

Explanation:

Master boot record is correct. The MBR is a boot-sector structure traditionally located in the first sector of an MBR-partitioned storage device. It contains small bootstrap code that firmware can load and execute during the legacy BIOS boot process, along with the disk partition table and a boot signature. The bootstrap code generally transfers control to the next stage of a boot loader, allowing the operating system boot process to continue. In Red Hat Enterprise Linux administration, understanding the MBR is useful when working with legacy BIOS booting, partitioning schemes, and boot loader installation or recovery. Modern systems can instead use UEFI firmware with a GUID Partition Table (GPT) and an EFI System Partition, but MBR remains an important foundational term and partition-table format. Red Hat documentation describes selecting an MBR partition table for BIOS-based booting and identifies the boot loader's relationship to the system boot process. See the Red Hat manual partitioning documentation and the Red Hat Linux boot process overview.

QUESTION NO: 9

Which of the following run levels start the system in single user mode?

Each correct answer represents a complete solution. Choose all that apply.

A. single

B. S

C. 0

D. 1

E. s

ANSWER: B D E

Explanation:

In the traditional SysV init runlevel model used for this administration topic, **1** is the standard numeric runlevel for single-user or maintenance mode. This mode starts only the essential services needed for administrative recovery and maintenance, rather than bringing up the normal multiuser environment. The symbolic runlevel names **S** and **s** are aliases for single-user mode and are treated equivalently to runlevel 1 by SysV init. Consequently, specifying **1**, **S**, or **s** requests a boot or transition into the single-user operating state. These aliases are useful because boot loaders, init implementations, and administrative procedures may use either the numeric form or the symbolic single-user form. Modern Red Hat Enterprise Linux releases use systemd targets internally, but retain runlevel compatibility mappings; the single-user state maps to `rescue.target`. See the init(8) manual page for runlevel behavior and Red Hat's systemd administration documentation for runlevel compatibility information.

QUESTION NO: 10

Which of the following commands is used to modify a user account?

A. USERADD

B. MOUNT

C. MODUSER

D. USERMOD

ANSWER: D

Explanation:

The `usermod` command is used to modify the attributes of an existing local user account. A system administrator can use it to change important account properties, such as the user's login name, numerical UID, primary or supplementary group membership, home directory, login shell, account expiration date, or password-aging settings. For example, `usermod -s /bin/bash username` sets an existing user's login shell, while `usermod -aG groupname username` appends that user to a supplementary group. Because these changes affect account records such as `/etc/passwd`, `/etc/shadow`, and `/etc/group`, the command is ordinarily run by the root user or through `sudo`. This is the standard account-management utility supplied by the shadow-utils package on Red Hat Enterprise Linux systems. Consult the Red Hat documentation on managing users and groups and the usermod manual page for its supported options and usage.

QUESTION NO: 11

Which of the following options is used with yum command to update a currently installed package?

A. -U

B. update

C. provides

D. -u

ANSWER: B

Explanation:

The `update` subcommand is used with `yum` to update installed packages to newer versions available from enabled software repositories. For example, `yum update package_name` updates the named installed package, including any required dependency updates. Running `yum update` with no package name checks all installed packages and updates those for which newer versions are available. Yum resolves dependencies and uses the configured repositories to select suitable package versions, making it the standard administrative command for applying package updates on Red Hat Enterprise Linux systems that use yum. In current RHEL versions, `yum` is provided as a compatibility interface to DNF, so the familiar command syntax remains applicable. Administrators commonly review the transaction summary before confirming it, particularly on production systems, because updating a package can also require related package changes. Red Hat documents package updates through the `yum update` command in its RHEL 7 system administration guidance, while current documentation describes the equivalent DNF package-management workflow. See Red Hat's RHEL 7 package update documentation and Red Hat's current software update guidance.

QUESTION NO: 12

Which of the following tools can be used to manage Red Hat Virtualization environment?

A. xen

B. chroot

C. sed

D. xm

E. Red Hat Virtualization Manager

ANSWER: E

Explanation:

Red Hat Virtualization Manager is the centralized management component for a Red Hat Virtualization environment. It provides the Administration Portal and a REST API for managing the virtualized data center, including hosts, storage domains, networks, virtual machines, templates, users, roles, and high-availability policies. Administrators connect to the Manager's web-based Administration Portal to perform routine operational tasks, monitor resource usage and events, configure the environment, and control the lifecycle of virtual machines. The Manager also coordinates actions across the underlying virtualization hosts, providing an environment-wide control plane rather than a command limited to an individual host or unrelated operating-system utility. In Red Hat Virtualization architecture, the Manager is therefore the appropriate tool to identify when the question asks how to manage the overall virtualization environment. Red Hat's documentation describes the Manager as the central component used to manage resources in the environment and documents use of its Administration Portal for administrative work. See the Red Hat Virtualization Administration Guide and the Red Hat Virtualization technical reference.

QUESTION NO: 13

You work as a System Administrator for McNeil Inc. The company has a Linux-based network. You are a root user on the Red Hat operating system. You have added a new swap partition on /dev/hdb2. After adding the new swap partition and enabling it, you want to ensure that it is enabled. Which of the following commands will you use to accomplish the task?

Each correct answer represents a complete solution. Choose two.

A. cat /proc/swaps

B. mkswap

C. swapon

D. free

ANSWER: A D

Explanation:

`cat /proc/swaps` is correct because the kernel maintains the `/proc/swaps` pseudo-file to report every swap area currently active. Its output identifies each active swap device or file, its type, size, amount in use, and priority. Therefore, after enabling `/dev/hdb2`, an administrator can confirm that the partition is active by checking that it appears in this file.

`free` is also correct because it reports a system-wide memory summary, including total configured swap, used swap, and available swap. Once the newly added partition has been enabled, the reported total swap capacity reflects that active swap space. This makes `free` a useful confirmation that the host has swap enabled and that the additional capacity is available to the kernel. For detailed device-level verification, `/proc/swaps` provides the partition name; for a concise overall capacity check, `free` provides the memory-and-swap summary. Red Hat documents inspection of swap information through procfs and memory-reporting tools in its storage and performance documentation. See Red Hat: Getting information about swap spaces and free(1) manual page.

QUESTION NO: 14

Which command creates a database that describes kernel's modules dependencies?

A. depmod

B. lsmod

C. rmmod

D. modprobe

ANSWER: A

Explanation:

The `depmod` command creates the module-dependency database for the currently installed Linux kernel. It scans the kernel module files, normally under `/lib/modules/$(uname -r)/`, examines their exported and required symbols, and generates dependency metadata such as `modules.dep`. This metadata records which modules must be available or loaded to satisfy a module's dependencies. Tools that manage modules, particularly `modprobe`, use this generated dependency information to load a requested module together with any required supporting modules in the proper order. Administrators commonly run `depmod` after manually installing, removing, or updating kernel module files; package-management and kernel-installation processes also run it automatically as needed. The command can be directed at a particular kernel release with an argument such as `depmod 5.x.y`, while `depmod -a` explicitly processes all modules for the target release. The depmod manual page describes its role in generating `modules.dep` and related map files. Red Hat's kernel module management documentation also explains the dependency-based module-loading workflow.

QUESTION NO: 15 - (SIMULATION)

John works as a Network Administrator for Perfect Solutions Inc. The company has a Linux-based network. John is working as a root user on the Linux operating system. John wants to remove some installed kernel modules. Which of the following commands will John use to accomplish his task?

ANSWER: See the explanation for the answer

Explanation:

Use the `rmmod` command to remove (unload) an installed/running kernel module.

```
rmmod module_name
```

For example:

```
rmmod e1000
```

The module must not be in use and must have no dependent modules loaded. Alternatively, `modprobe -r module_name` can unload a module and its unused dependencies, but the direct command requested is `rmmod`.

QUESTION NO: 17

Which of the following run levels start the system in single user mode?

Each correct answer represents a complete solution. Choose all that apply.

A. 0

B. single

C. S

D. s

E. 1

ANSWER: B C D E

Explanation:

On legacy SysV-init-based Red Hat systems, runlevel `1` is the standard single-user runlevel. It starts the system with a minimal environment for administrative maintenance, typically without normal multiuser services or network access. SysV init also recognizes the special runlevel names `S` and `s` as single-user mode selectors. These names are commonly used when invoking `init` or `telinit`, and are treated as aliases for entering the maintenance-oriented single-user environment.

The `single` selector is also accepted as a legacy boot-time request for single-user mode. For example, an administrator can add it to the kernel command line when booting to request maintenance mode. Therefore, `single`, `S`, `s`, and `1` all identify valid ways to start a SysV-init system in single-user mode. In modern Red Hat Enterprise Linux releases that use systemd, the comparable maintenance target is generally `rescue.target`; however, this question uses the traditional SysV runlevel terminology. Red Hat documents the relationship between traditional runlevels and systemd targets in its systemd administration guide and describes booting into rescue or emergency maintenance environments in the RHEL rescue-shell documentation.

QUESTION NO: 18

You work as a system administrator for Tech Perfect Inc. The company has a Linux-based network. You are a root user on the Linux operating system. A user, Jetson, wants to view the speed and the duplex for his Ethernet card. His ethernet card is activated on interface eth0. Which of the following commands can you use to accomplish this task?

A. ifup eth0

B. ifconfig eth0

C. ip link show eth0

D. sudo ethtool eth0

ANSWER: D

Explanation:

`sudo ethtool eth0` is correct because `ethtool` is the standard Linux utility for querying and changing settings of a network interface's Ethernet controller and driver. When it is run against `eth0`, its output includes link-related fields such as `Speed:` and `Duplex:`, for example, `Speed: 1000Mb/s` and `Duplex: Full`. This directly provides the requested operational speed and duplex mode for the Ethernet card. The output also commonly reports useful related information, including whether link detection is active, the port type, supported link modes, advertised link modes, and auto-negotiation status. Because the scenario states that the administrator is already the root user, the `sudo` prefix is unnecessary but does not prevent the command from being valid; it simply invokes `ethtool` with elevated privileges when sudo is configured. The `ethtool` manual documents that invoking the utility with an interface name displays the current Ethernet device settings, including speed and duplex. See the ethtool(8) Linux manual page and the Red Hat networking diagnostics documentation.

QUESTION NO: 19

John works as a Network Administrator for Perfect Solutions Inc. The company has a Linux-based network. He is working as a root user on the Linux operating system. He wants to add a soft limit quota warning for the users, in which, after exceeding the quota value, a user will receive e-mail warnings about being over quota. Which of the following commands will John use to accomplish his task?

Each correct answer represents a complete solution. Choose all that apply.

A. repquota

B. quotaon

C. warnquota

D. edquota

E. quotaoff

ANSWER: C D

Explanation:

The `edquota` command is used by the root user to create or modify filesystem quota limits for individual users or groups. In its editor interface, the administrator can define soft limits for block usage and inode usage, as well as hard limits and grace periods. A soft limit is the threshold that a user may temporarily exceed during the applicable grace period, making it the appropriate quota setting for this requirement.

The `warnquota` command provides the notification portion of quota administration. It checks configured filesystems for users who are over their quota limits and sends warning email messages to those users. Administrators commonly configure `warnquota` to run periodically through a scheduler so that users are notified while they remain over a soft limit. Thus, `edquota` establishes the soft quota threshold, and `warnquota` sends the requested email warnings after the threshold is exceeded. For command behavior and quota-management details, see the edquota manual page and the warnquota manual page.

QUESTION NO: 20

You work as a Network Administrator for Tech Perfect Inc. The company has a Linux-based network. You have assigned a permission set of 440 on a folder. Which of the following permissions are added to the folder? Each correct answer represents a part of the solution. Choose all that apply.

A. No permission to others.

B. Read permission to the owner.

C. Read and Execute permissions to the group.

D. Read, Write, and Execute permissions to the owner and the group.

E. Read permission to the group.

F. Read and Execute permissions to the owner.

ANSWER: A B E

Explanation:

Linux numeric permissions are expressed as three octal digits in the order owner, group, and others. Each digit is the sum of the relevant permission values: read is 4, write is 2, and execute is 1. Therefore, a mode of `440` assigns the first `4` to the file or directory owner, granting read permission. The second `4` applies to the owning group and likewise grants read permission. The final `0` applies to all other users and grants no permissions. Consequently, “Read permission to the owner,” “Read permission to the group,” and “No permission to others” accurately describe the permissions represented by `440`. On a directory, these are still the stored mode bits; practical traversal into the directory additionally requires execute permission, but that does not change the meaning of the numeric mode itself. Red Hat documents the owner/group/other permission classes and the read, write, and execute permission values in its file-permission guidance. See Red Hat Enterprise Linux documentation on managing file-system permissions and Red Hat Storage Administration Guide: permissions.

QUESTION NO: 22

Which of the following commands displays the version number of the running Linux kernel?

A. w

B. uname -r

C. motd

D. lsmod

ANSWER: B

Explanation:

`uname -r` is the appropriate command because it prints the kernel release string for the kernel currently running on the system. This release value is commonly what administrators need when checking the active kernel version, such as `5.14.0-503.35.1.el9_5.x86_64`. It identifies the kernel release, including distribution-specific build information where applicable. The `uname` utility retrieves system information from the running kernel, and the `-r` option specifically requests the kernel release. This is useful for confirming that a newly installed kernel is active after a reboot, documenting a system's operating environment, or checking compatibility requirements for kernel modules and software. Running the command requires no special privileges and reports information about the live kernel rather than merely listing kernel packages installed on disk. The GNU Coreutils documentation describes `uname` and its release option, while the Linux manual page documents the system information exposed by the command: GNU Coreutils: uname invocation and uname(1) Linux manual page.

QUESTION NO: 23

Fill in the blank with the appropriate command's option.

_____option is used with chage command to change the number of days of warning before a password change is required.

A. -W warndays

ANSWER: A

Explanation:

The `-W warndays` option is used with `chage` to set the number of days during which a user receives a warning that their password is approaching its expiry date. The `warndays` value is a nonnegative integer representing how many days before password expiration the warning period begins. For example, `chage -W 7 username` configures a seven-day warning period for `username`, so the user is notified during the final seven days before a password change becomes required. This setting is useful for account-lifecycle administration because it gives users adequate notice to update passwords before an enforced expiration disrupts access. The setting changes password-aging information stored for the specified local account and can be reviewed with `chage -l username`. The `chage` manual documents `-W` as `--warndays` and defines it specifically as the number of days of warning before a password change is required. See the chage(1) manual page and the Red Hat Enterprise Linux documentation on managing user accounts.

QUESTION NO: 24

John works as a Network Administrator for Perfect Solutions Inc. The company has a Linux-based network. John is working as a root user on the Linux operating system. He wants to change the quota setting of the user named MariA. Which of the following commands will he use to accomplish the task?

A. edquota -u MariA

B. edquota -g

C. edquota -p

D. edquota â€“T

ANSWER: A

Explanation:

`edquota -u MariA` is the appropriate command because `edquota` opens the quota editor for a specified account, and the `-u` option explicitly selects a user quota. When run by root (or an appropriately privileged administrator), it opens the configured text editor with MariA's quota records for quota-enabled file systems. The administrator can then set or modify the

soft and hard limits for disk blocks and inodes, save the file, and have the updated limits applied. The username must be supplied as an argument; Linux account names are case-sensitive, so `MariA` must be entered with the capitalization shown in the question. This command assumes that the relevant filesystem has user quotas enabled and quota accounting is active. The `edquota` manual documents `-u` as the user-quota selection option and describes editing quota limits for named users. See the edquota manual page and Red Hat's quota management documentation.

QUESTION NO: 25

Which of the following commands is used to specify that the gpm service should be executed at run level 3?

A. chkconfig -- level 3 gpm on

B. chkconfig --level 3 gpm

C. chkconfig level 3 gpm

D. chkconfig

E. chkconfig --level 3 gpm on

ANSWER: E

Explanation:

`chkconfig --level 3 gpm on` is the correct command because it enables the SysV init service named `gpm` specifically for runlevel 3. The `--level` option accepts one or more runlevel numbers and limits the change to those levels; `gpm` identifies the service; and `on` creates or enables the appropriate startup links for that service. Runlevel 3 is traditionally the multi-user, network-enabled text-mode operating state on SysV-init-based Red Hat Enterprise Linux systems. Thus, after this configuration, the system starts the gpm service when entering runlevel 3, assuming the service is installed and its init script supports chkconfig management. The option must be written as the single token `--level`, with no space between the two hyphens and the word `level`. Red Hat documents that `chkconfig --level levels service_name on` configures a service for the specified runlevels. See the Red Hat chkconfig documentation and the chkconfig manual page.

QUESTION NO: 26

The ____ command is used to view the TCP/IP configuration on a Windows 9x computer.

A. IFCONFIG

B. WINIPCFG

C. CHKDSK

D. IPCONFIG

ANSWER: B

Explanation:

`WINIPCFG` is the correct command for viewing TCP/IP configuration on Windows 9x systems, including Windows 95, Windows 98, and Windows Me. Its name is short for Windows IP Configuration. When run from the Run dialog box or a command prompt, it opens a graphical utility that displays network details for the selected adapter, such as the IP address, subnet mask, default gateway, DHCP status, and DNS server information. The utility could also be used to release and renew a DHCP lease and to inspect detailed adapter information. This was particularly appropriate to the Windows 9x networking interface, which commonly exposed configuration information through GUI tools rather than the later command-line interface used by Windows NT-based systems. The historically correct executable name is `WINIPCFG`, without the extra "IN" in "WINIPCONFIG." For background on the utility and the Windows versions that included it, see Winipcfg. Microsoft's current documentation for the later `ipconfig` command is available at Microsoft Learn: ipconfig.

QUESTION NO: 27 - (DRAG DROP)

Drag and drop the appropriate command in front of its respective function.

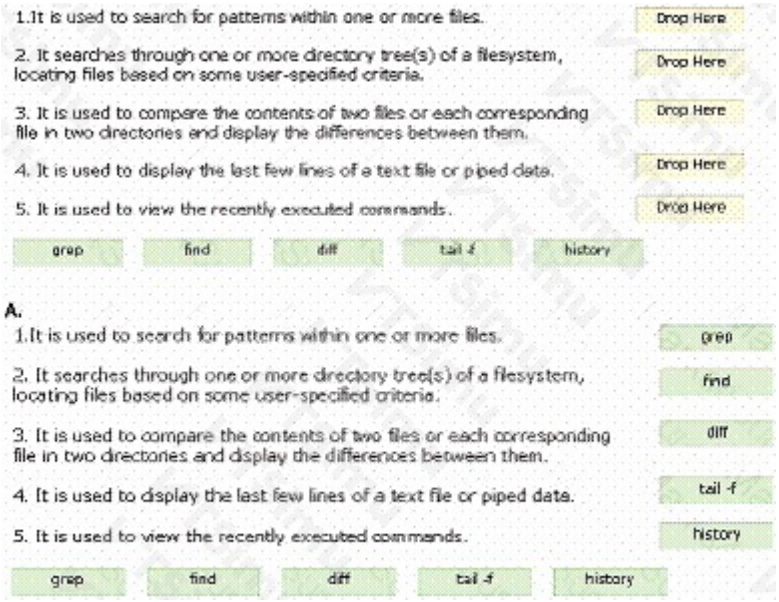

ANSWER:

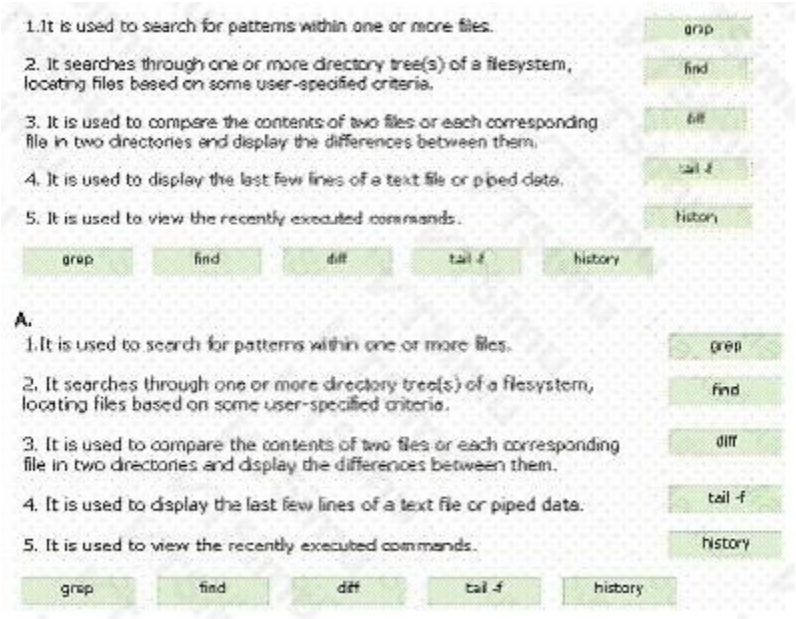

Explanation:

The correct command-to-function mappings use standard Red Hat Enterprise Linux shell utilities. `grep` searches input for lines matching a text pattern or regular expression, so it is the appropriate command for searching patterns in one or more files. The `find` command walks a directory hierarchy and selects files according to criteria such as name, ownership, type, permissions, time, or size. This makes it the command used to search through filesystem locations based on user-specified criteria.

`diff` compares the contents of two files, and it can also compare corresponding files in two directories. Its output identifies the differences between the compared content. For showing the final portion of file or stream content, the supplied `tail -f` command is the correct match. It displays trailing lines and, because of `-f`, continues to follow the file as new lines are appended; this is commonly used for monitoring logs. Finally, `history` displays commands previously entered in the current shell session. These are core command-line tools covered in Red Hat system administration. See the Red Hat terminal documentation and the GNU tail documentation for command behavior details.

QUESTION NO: 28

You work as a Network Administrator for Perfect Solutions Inc. You are required to configure a hard disk for a Linux workstation. You are using the FDISK utility to configure it. Which of the following options will you use to show all the existing partitions of the hard disk?

A. n

B. w

C. 1

D. p

ANSWER: D

Explanation:

The `p` command is used from the interactive `fdisk` prompt to print the current partition table. It displays the existing partitions on the selected disk, including each partition's device identifier, start and end sectors, size, type, and bootable status where applicable. This is particularly useful before making partitioning changes, because it lets an administrator inspect the disk's current layout and confirm which partitions already exist.

For example, after starting the utility with a command such as `fdisk /dev/sda`, entering `p` shows the partition-table information currently held for that disk. The displayed layout can be reviewed repeatedly during an `fdisk` session; no changes are written merely by printing the table. Red Hat documentation describes using `fdisk` to examine and manage disk partition tables, while the upstream manual explicitly identifies `p` as the command to print the partition table. See the Red Hat partitioning documentation and the fdisk manual page.

QUESTION NO: 29

Which of the following Linux commands is used to assign privileges over a particular file to a designated user?

A. chgrp

B. chmod

C. chown

D. useradd

ANSWER: C

Explanation:

`chown` is correct because it changes the ownership of a file or directory, allowing a specified user to become the file's owner. File ownership is a core part of Linux discretionary access control: each filesystem object has a user owner and a group owner, and the owner's permission bits determine the access available to that owner. For example, an administrator can use `chown alice report.txt` to make `alice` the owner of `report.txt`. After ownership is assigned, the designated user has the access represented by the file's owner permissions and can ordinarily modify those permissions with `chmod` for files they own. Administrators may also use `chown alice:developers report.txt` to set both the user owner and group owner in one command. This use of `chown` is the standard Red Hat Linux administration method for assigning a file to a designated user. The chown manual page defines the command as changing file owner and group, while Red Hat's documentation describes ownership and permissions as fundamental controls for access to filesystem objects: Managing file system permissions.

QUESTION NO: 30

Which of the following images provide support for different types of network interface cards and other network devices?

A. drvblock.img

B. bootdisk.img

C. drvnet.img

D. pcmciadd.img

ANSWER: C

Explanation:

`drvnet.img` is the network-driver supplemental image used by older Red Hat installation media. Its purpose is to make additional network-interface-card drivers and related network-device drivers available during installation, particularly when the installer does not include the driver required for the system's network hardware. The image can be supplied when prompted for an appropriate driver disk, allowing the installer to detect and configure the device so that network-based installation methods or network connectivity can be used.

This naming convention is meaningful: the `drvnet` portion identifies the image as containing network drivers, rather than boot files or storage-controller drivers. Red Hat installation documentation describes the use of driver disks to load hardware support that is not already available to the installer, including drivers needed for network adapters. See the Red Hat Enterprise Linux Installation Guide: Driver Disks and Red Hat's RHEL 5 Installation Guide for historical installer and driver-disk guidance.

QUESTION NO: 31

John works as a Network Administrator for Perfect Solutions Inc. The company has a

Linux-based network.

John is working as a root user on the Linux operating system. He wants to check the errors of an ext2 filesystem. Which of the following commands can John use to accomplish the task? Each correct answer represents a complete solution. Choose all that apply.

A. fsck.ext2

B. fsck

C. e2fsck

D. fsck -N

ANSWER: A B C

Explanation:

The commands `fsck.ext2`, `fsck`, and `e2fsck` can be used to check an ext2 filesystem when supplied with the filesystem device, such as `/dev/sdb1`. The `e2fsck` utility is the filesystem-checking program for the ext2/ext3/ext4 filesystem family. It examines filesystem metadata, directory structures, inode and block allocation information, and can report or repair detected inconsistencies. `fsck.ext2` invokes the ext2-specific checker and is normally provided as a link or wrapper to `e2fsck`, making it an appropriate command name for checking an ext2 filesystem.

The general `fsck` command is a front end that determines the filesystem type and calls the appropriate checker, including the ext-family checker for an ext2 filesystem. For safe administration, the target filesystem should generally be unmounted before performing checks or repairs, because checking a mounted filesystem can produce unreliable results or risk filesystem changes while it is active. Administrators can use options such as `-f` to force a full check when needed and should review the utility's prompts or use appropriate repair-policy options in accordance with operational procedures. See the e2fsck manual page and the fsck manual page.

QUESTION NO: 32

You work as a Network Administrator for Tech Perfect Inc. You want to configure quotas for everyone on the /home directory. What will you add to the options in /etc/fstab?

Each correct answer represents a complete solution. Choose two.

A. usrquota

B. grpquota

C. groupquota

D. userquota

ANSWER: A B

Explanation:

To enable traditional disk quota accounting on the filesystem mounted at `/home`, add the `usrquota` and `grpquota` mount options to that filesystem's entry in `/etc/fstab`. The `usrquota` option enables user quota accounting, allowing the administrator to define and enforce per-user limits for disk blocks and inode usage. The `grpquota` option enables group quota accounting, so limits can also be applied collectively to members of a group. Using both options provides quota support for both types of principals, which is appropriate when quota control is required for all home-directory users and their groups.

After changing `/etc/fstab`, the filesystem must be remounted, or the system restarted, for the mount options to take effect. Quota database files must then be created and quota checking enabled with the applicable quota administration tools. Red Hat documents the use of `usrquota` and `grpquota` as filesystem mount options when configuring user and group disk quotas. See the Red Hat disk quota documentation and the quotaon manual page.

QUESTION NO: 34

You want to shrink a Logical Volume manager partition named /dev/disk1/lvm from 4GB to

3GB. Which of the following commands can you use to accomplish the task?

Each correct answer represents a complete solution. Choose all that apply.

A. lvreduce

B. resize2fs

C. lvsize

D. reduce2fs

E. lvreduce --resizefs --size 3G /dev/disk1/lvm

ANSWER: A E

Explanation:

`lvreduce` is the LVM command used to reduce the allocated size of a logical volume. For a complete and safer shrink operation, it can be used with the `--resizefs` (or `-r`) option, which requests filesystem resizing before the logical volume is reduced. For example, `lvreduce --resizefs --size 3G /dev/disk1/lvm` reduces the target logical volume to 3 GiB while coordinating the supported filesystem resize operation. This ordering is essential: the filesystem must fit within the new, smaller logical-volume boundary before storage is removed.

The operation requires a filesystem that supports shrinking. In particular, XFS filesystems cannot be shrunk; an XFS volume requires a backup, recreation at the desired size, and restoration instead. Before reducing any logical volume, ensure that current backups exist and that the filesystem type and its shrink requirements have been verified. Red Hat documents logical-volume reduction and the `--resizefs` workflow in its logical volume management documentation. Red Hat also documents the XFS limitation in its XFS filesystem resizing guidance.

QUESTION NO: 35

You work as a System Administrator for Tech-Perfect Inc. The company has Linux-based network. You are a root user on the Red Hat operating system. You want to change the speed of your network card 1000Mbps with full duplex on eth1. For this, you run the following command on command-linE. Ethtool-s eth1 speed

1000 duplex full as you run this command, the link goes down and you are unable to ping out with your NIC. What should you do to accomplish the task?

Each correct answer represents a part of the solution. Choose three.

A. ifdown eth1

B. ethtool -s eth1 autoneg off speed 1000 duplex full

C. ethtool eth1

D. ifup eth1

ANSWER: A B D

Explanation:

To apply a forced link configuration safely, first take the interface out of service with `ifdown eth1`. This prevents active traffic and address configuration from interfering while the network-device settings are changed. Next, use `ethtool -s eth1 autoneg off speed 1000 duplex full` to explicitly disable autonegotiation and request the required 1000 Mbps, full-duplex operating mode. The corresponding switch port must be configured compatibly; otherwise, the physical link can remain down or operate unreliably.

After the device settings are applied, run `ifup eth1` to activate the interface again and restore its configured network connectivity, such as its IP addressing and routes. This down/configure/up sequence is the intended operational procedure for the question and ensures that the interface is returned to service after the requested link-mode change. Red Hat documents `ethtool` as the utility used to examine and modify Ethernet NIC settings, including speed, duplex, and autonegotiation. See Red Hat's network troubleshooting documentation and the ethtool manual page for the supported settings and syntax.

QUESTION NO: 36

Which of the following scripts invokes the runlevel scripts?

A. /etc/rc.d/rc.local

B. /etc/rc.d/rc.sysinit

C. /etc/rc.d/rc.serial

D. /etc/rc.d/rc

ANSWER: D

Explanation:

`/etc/rc.d/rc` is the script that invokes the runlevel scripts in the traditional SysV init layout used by older Red Hat Enterprise Linux systems. When init changes to a selected runlevel, it calls this rc script and supplies the runlevel as an argument. The script then processes the corresponding runlevel directory, such as `/etc/rc.d/rc3.d/` or `/etc/rc.d/rc5.d/`. Those directories contain symbolic links to service-control scripts in `/etc/rc.d/init.d/`. Link names beginning with `K` are handled to stop services, and links beginning with `S` are handled to start services, in numeric order. This design lets administrators control which services are active at each runlevel without duplicating the underlying service scripts. Thus, `/etc/rc.d/rc` acts as the runlevel dispatcher that executes the appropriate ordered service actions. Modern RHEL releases use systemd units and targets instead of SysV runlevels, but the described script is the correct

component for the legacy runlevel-script architecture addressed by this question. See the init(8) manual page and Red Hat's init scripts documentation.

QUESTION NO: 37

Which of the following commands can you use to create an ext3 file system?

Each correct answer represents a complete solution. Choose two.

A. mkfs.ext3

B. mke2fs

C. mke2fs -j

D. mkfs.ext2

ANSWER: A C

Explanation:

`mkfs.ext3` and `mke2fs -j` can both create an ext3 filesystem when followed by the target device or filesystem image. For example, an administrator can run `mkfs.ext3 /dev/sdb1` or `mke2fs -j /dev/sdb1`. The `mkfs.ext3` command is the filesystem-specific formatting utility for ext3; it invokes the ext filesystem creation tools with the settings required for ext3. The `-j` option to `mke2fs` creates a journal, which makes the newly created ext filesystem an ext3 filesystem. Journaling records filesystem metadata changes before they are committed, enabling faster and more reliable recovery after an unclean shutdown. In practical administration, verify that the target is the intended unused partition or logical volume before formatting it, because filesystem creation replaces existing filesystem metadata and can make stored data inaccessible. Red Hat's storage documentation describes creating filesystems with the `mkfs` tools, and the `mke2fs` manual documents that `-j` creates an ext3 journal. See Red Hat: Creating a file system and mke2fs(8) manual page.

QUESTION NO: 38

Which of the following switches of the rpm command lists the package that owns the specified file?

A. -a

B. -q

C. -e

D. -f

ANSWER: D

Explanation:

The `-f` switch is correct because it performs an RPM database query for the package that owns a specified file. The usual syntax is `rpm -qf /path/to/file`. For example, running `rpm -qf /usr/bin/bash` returns the installed RPM package that provided that file. This is particularly useful to identify the source package of an existing executable, configuration file, library, or documentation file before investigating package details, verifying files, or reinstalling a package.

In this command, `-q` places RPM in query mode and `-f` supplies the file-based query selector. Although `-f` is normally used together with `-q`, it is the switch that specifically tells RPM to determine ownership from a file path. The specified path must refer to a file recorded in the local RPM database; files created manually or installed outside RPM package management have no RPM owner. Red Hat documents this operation as querying a package by file name with `rpm -qf`. See the Red Hat RPM commands documentation and the rpm(8) manual page.