PassQueen
a.k.a dumpsqueen.com

LPIC-1 Exam 102, Part 2 of 2, version 5.0

LPI 102-500

Version Demo

Total Demo Questions: 32

Total Premium Questions: 328

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Shells and Shell Scripting	55
Topic 2, User Interfaces and Desktops	36
Topic 3, Administrative Tasks	59
Topic 4, Essential System Services	75
Topic 5, Networking Fundamentals	61
Topic 6, Security	42
Total	328

QUESTION NO: 1

Which command displays the UUID and filesystem type of available block devices?

A. `fdisk`

B. `blkid`

C. `df`

D. `mountpoint`

E. `fsstat`

ANSWER: B

Explanation:

The correct command is `blkid`. It probes block devices and displays metadata such as filesystem type, UUID, partition UUID, and labels. This information is commonly used when preparing persistent entries in `/etc/fstab`. See the blkid(8) manual page.

QUESTION NO: 2

Which of the following commands will delete the default gateway from the system's IP routing table? (Choose two.)

A. ifconfig unset default

B. route del default

C. ip route del default

D. netstat -r default

E. sysctl ipv4.default_gw=0

ANSWER: B C

Explanation:

`route del default` and `ip route del default` delete the default route from the kernel IP routing table. A default route is the route used when no more-specific route matches a destination address; in IPv4 routing terms, it corresponds to the destination prefix `0.0.0.0/0`. It commonly specifies the system's default gateway, although a default route can also be configured through another interface or routing mechanism.

The `route` command is the traditional net-tools utility. Its deletion syntax accepts `default` as the route destination, so `route del default` requests removal of that route. The `ip` command from iproute2 is the modern Linux routing-management utility. Its `route del` (or `route delete`) operation also accepts `default` as the shorthand destination for the default prefix. In both cases, the command modifies the active kernel routing table and normally requires administrative privileges, such as running through `sudo`. If multiple default routes exist with differing attributes, a more specific deletion command may be needed to select the intended route. See the ip-route(8) manual page and the route(8) manual page.

QUESTION NO: 3

What is Irue aboul Network Manager on a Linux syslem that uses its distribution's mechanisms lo configure network interfaces? (Choose TWO correct answers.)

A. NetworkManager reconfigures all network interfaces to use DHCP unless they are specifically managed by NetworkManager

B. NetworkManager musl be explicitly enabled for each interface it should manage.

C. NetworkManager by default does not change interfaces which are already configured.

D. NetworkManager disables all interfaces which were not configured by NetworkManager.

E. NetworkManager can be configured to use the distribution's network interface configuration.

ANSWER: C E

Explanation:

NetworkManager is designed to coexist with distribution-specific network configuration mechanisms rather than automatically replacing every existing interface setup. On a system where interfaces have already been configured through the distribution's native mechanism, NetworkManager normally leaves those existing configurations unchanged unless it is directed to manage them. This prevents an installed or started NetworkManager service from unexpectedly disrupting a working static configuration, server connection, or other pre-existing network arrangement.

NetworkManager can also be configured to use a distribution's network interface configuration format through its settings plugins. For example, the ifcfg-rh plugin supports the traditional Red Hat-family `ifcfg-*` files, allowing NetworkManager to read and manage connections represented in that distribution-specific format. The available plugins and preferred configuration source depend on the distribution and its installed NetworkManager components. These capabilities let administrators adopt NetworkManager while retaining compatible, distribution-provided configuration files and workflows. See the NetworkManager configuration documentation and the ifcfg-rh settings plugin documentation.

QUESTION NO: 4 - (SIMULATION)

What command displays all aliases defined in the current shell? (Specify the command without any path information)

ANSWER: See the explanation for the answer

Explanation:

The command is:

```
alias
```

Run `alias` with no arguments to display all aliases currently defined in the active shell.

QUESTION NO: 5

Which of the following connection types, as seen in nmcli connection show, may exist in NetworkManager? (Choose three.)

A. tcp

B. ethernet

C. wifi

D. ipv6

E. bridge

ANSWER: B C E

Explanation:

NetworkManager stores network configurations as connection profiles. Each profile has a connection type that identifies the underlying networking technology or virtual device managed by that profile. Consequently, **ethernet**, **wifi**, and **bridge** may appear in the TYPE column of `nmcli connection show`. An Ethernet profile configures a wired network interface, while a

Wi-Fi profile configures wireless connectivity, including such settings as the SSID and wireless security. A bridge profile represents a software bridge device that can join interfaces into one Layer 2 network segment; it can itself have IP configuration and can be used with bridge-port profiles.

The `nmcli connection show` command lists the saved profiles known to NetworkManager and displays fields including NAME, UUID, TYPE, and DEVICE. NetworkManager's documented connection settings include `802-3-ethernet`, `802-11-wireless`, and `bridge`; nmcli presents the corresponding user-facing connection types as ethernet, wifi, and bridge. This distinction is important: the connection TYPE describes the profile/device technology, whereas IP addressing is configured as settings within an applicable connection profile. See the nmcli documentation and the NetworkManager settings specification.

QUESTION NO: 6

Which of the following may occur as a consequence of using the command ifconfig? (Choose THREE correct answers.)

A. New name servers may be added to the resolver configuration.

B. Network interfaces may become active or inactive.

C. The routing table may change.

D. IP addresses may change.

E. The system ' s host name may change.

ANSWER: B C D

Explanation:

`ifconfig` is a legacy but still commonly available utility for viewing and changing a network interface's runtime configuration. "Network interfaces may become active or inactive" is correct because the `up` and `down` operations set or clear the interface UP state. This directly controls whether the kernel can use that interface for normal network communication.

"IP addresses may change" is also correct. An address can be supplied as part of an `ifconfig` invocation, and interface address configuration can therefore be added, replaced, or removed. Addressing details such as the netmask and broadcast address can be configured at the same time.

"The routing table may change" is correct as a consequence of modifying an interface address or its netmask. The kernel maintains routes associated with directly connected networks; changing the address/prefix configuration can add, remove, or alter the effective connected route for that interface. Consequently, traffic that previously matched a route through the interface can be handled differently after the configuration change. Although modern Linux systems generally prefer the `ip` command from iproute2 for these tasks, these effects accurately describe operations supported by `ifconfig`. See the ifconfig(8) manual page and the Linux network device documentation.

QUESTION NO: 7

Given the following routing table:

```
Kernel IP routing table
Destination     Gateway         Genmask         Flags   Metric  Ref    Use Iface
0.0.0.0         192.168.178.1   0.0.0.0         UG      0       0        0 wlan0
192.168.1.0     0.0.0.0         255.255.255.0   U       0       0        0 eth0
192.168.2.0     192.168.1.1     255.255.255.0   U       0       0        0 eth0
192.168.178.0   0.0.0.0         255.255.255.0   U       9       0        0 wlan0
```

How would an outgoing packet to the destination 192.168.2.150 be handled?

A. It would be passed to the default router 192.168.178.1 on wlan0.

B. It would be directly transmitted on the device eth0.

C. It would be passed to the default router 255.255.255.0 on eth0.

D. It would be passed to the router 192.168.1.1 on eth0.

E. It would be directly transmitted on the device wlan0.

ANSWER: D

Explanation:

It would be passed to the router 192.168.1.1 on eth0. The routing table contains a route for the 192.168.2.0 network that uses 192.168.1.1 as its next-hop gateway and eth0 as the outgoing interface. Since 192.168.2.150 belongs to that destination network, the kernel selects this matching network route rather than using a default route. The host does not send the packet directly to 192.168.2.150 at layer two, because that address is not on the directly connected eth0 network. Instead, it sends an Ethernet frame over eth0 to the MAC address of the next-hop router, 192.168.1.1. That router is then responsible for forwarding the IP packet toward the 192.168.2.0 network. Linux route selection uses the most specific applicable route, commonly called longest-prefix matching; a default route is considered only when no more-specific route matches the destination. The ip-route manual page describes route lookup and next-hop gateways, while the Requirements for IP Version 4 Routers (RFC 1812) documents the forwarding behavior performed through a next-hop router.

QUESTION NO: 8

Which of the following statements is valid in the file /etc/nsswitch.conf?

A. multi on

B. 192.168.168.4 dns-server

C. hosts: files dns

D. include /etc/nsswitch.d/

ANSWER: C

Explanation:

The valid statement is `hosts: files dns`. The `/etc/nsswitch.conf` file controls the Name Service Switch (NSS), which determines the sources Linux consults for several system databases, including host-name resolution. Its standard syntax begins with a database name, followed by a colon and one or more service sources. Here, `hosts` is the database being configured, while `files` and `dns` are NSS service sources.

This configuration directs the resolver to look first in local files, principally `/etc/hosts`, for a host-name-to-address mapping. If the requested name is not resolved through that source, NSS then uses DNS. DNS server details and related resolver settings are normally obtained from `/etc/resolv.conf` (or are managed by a resolver service that supplies equivalent configuration). The order is significant: placing `files` before `dns` permits locally defined host entries to take precedence over DNS results. This is a common and valid default-style configuration on Linux systems. The NSS configuration-file manual documents the database-and-service format and identifies both `files` and `dns` as supported services for the hosts database. See the nsswitch.conf(5) manual page and the resolv.conf(5) manual page.

QUESTION NO: 9

What is the main difference between the batch and at commands?

A. The batch command will run multiple times.The at command will only run once.

B. The batch command will run when system load is low. The at command runs at a specific time.

C. The at command reads commands from standard input. The batch command requires a command line argument.

D. The at command e-mails results to the user. The batch command logs results to syslog.

ANSWER: B

Explanation:

The batch command will run when system load is low. The at command runs at a specific time. Both commands are part of the at job scheduling facility and accept commands for later, non-interactive execution. The key distinction is the condition that determines when the queued job starts. An `at` job is assigned a requested date and time, such as a particular clock time today or a specified future date; the `atd` daemon runs it when that scheduled time arrives. A `batch` job, in contrast, is deferred until the system's load average falls below the configured threshold. This makes `batch` suitable for work that is not time-critical and should avoid competing with normal interactive or production workloads. The default load threshold is commonly 1.5, although it can be changed when starting `atd`, so it should not be described as a universal fixed value. In both cases, submitted jobs are normally placed in the same at-job queueing system, and job output is mailed to the submitting user when output is produced. The distinction being tested is therefore scheduled clock time versus execution when system load permits. See the at(1) manual page and the atd(8) manual page.

QUESTION NO: 10

Which of the following actions can be performed using `passwd`? (Choose two.)

A. Change an account password.

B. Lock an account password.

C. Change an account's UID.

D. Add an account to a supplementary group.

E. Change an account's home directory.

ANSWER: A B

Explanation:

The `passwd` command can change the password for an account. When invoked by a regular user without an account name, it changes that user's own password, subject to the system's password policy.

It can also lock an account password with the `-l` option. Locking prepends a marker to the encrypted-password field, preventing password authentication using that stored password. It does not necessarily disable all other authentication methods. See the passwd(1) manual page.

QUESTION NO: 11

Which of the following configuration files should be modified to globally set shell variables for all users?

A. /etc/profile

B. /etc/bashrc

C. ~/.bash_profile

D. /etc/.bashrc

E. /etc/shellenv

ANSWER: A

Explanation:

`/etc/profile` is the appropriate system-wide startup file for setting shell variables that should apply to all users of Bourne-compatible shells when they start a login shell. It is read during login-shell initialization, before each user's own profile file,

such as `~/.bash_profile`. Administrators commonly use it to establish globally applicable environment settings, including `PATH`, locale-related variables, editor preferences, and other exported variables needed consistently across user login sessions.

For example, placing `export EDITOR=vim` or an appropriate system-wide `PATH` adjustment in `/etc/profile` makes the setting available to users who log in through a terminal, console, or remote login service and whose login shell reads this standard profile file. Shell startup behavior is an important LPIC-1 topic because the scope of a setting depends on whether it is configured for an individual user or globally, and whether the shell is invoked as a login shell. The Bash Reference Manual documents that Bash reads `/etc/profile` for login shells, while the LPIC-1 objectives include managing user and global shell environments. See the GNU Bash startup-files documentation and the LPI Exam 102 objectives.

QUESTION NO: 12

Which of the following words is used to restrict the records that are returned from a SELECT SQL query based on a supplied criteria for the values in the records?

A. CASE

B. FROM

C. WHERE

D. IF

ANSWER: C

Explanation:

WHERE restricts the rows returned by a SQL `SELECT` statement according to a specified condition. It is placed after the `FROM` clause, which identifies the source table or tables, and before clauses such as `GROUP BY`, `HAVING`, or `ORDER BY` when those are present. A row is included in the query result only when the expression in `WHERE` evaluates to true. For example, `SELECT name FROM customers WHERE country = 'USA';` returns names only for customer rows whose `country` value is `'USA'`.

The condition may use comparison operators such as =, <, and >; logical operators including `AND`, `OR`, and `NOT`; pattern matching with `LIKE`; range checks using `BETWEEN`; membership tests using `IN`; and null checks using `IS NULL`. This filtering behavior is fundamental to querying relational databases because it enables a query to return only the data matching the requested criteria rather than every available row. PostgreSQL's SQL reference describes `WHERE` as the clause that eliminates rows not satisfying its search condition, and SQLite likewise documents it as the mechanism for filtering result rows. See the PostgreSQL SELECT documentation and the SQLite SELECT documentation.

QUESTION NO: 13

Which of the following lines are valid in the file /etc/hosts? (Choose TWO correct answers.)

A. 2001:db8::15 www.example.com www

B. www.example.com www 203.0.13.15

C. 203.0.113.15 www.example.com www

D. www.example.com,www 203.0.13.15,2001:db8::15

E. 2003.0.113.15,2001:db8::15 www.example.com www

ANSWER: A C

Explanation:

The entries `2001:db8::15 www.example.com www` and `203.0.113.15 www.example.com www` use the required `/etc/hosts` layout: an IP address first, followed by the host's canonical name and then any optional aliases, with fields separated by whitespace. The first entry uses an IPv6 address, while the second uses an IPv4 address; both address families are supported in hosts-file entries. In each case, `www.example.com` is the primary hostname associated with the address and `www` is an additional alias for the same address. The address prefixes shown are documentation ranges, but they are syntactically valid addresses and therefore make valid hosts-file records. The hosts file is a local, static name-resolution source, commonly consulted through the Name Service Switch configuration before or alongside DNS. A hosts-file record maps one address to one or more whitespace-separated names; it does not require a special delimiter between the canonical hostname and aliases. See the Linux hosts(5) manual page for the file format and the RFC 2606 documentation-domain definitions for the example domain name.

QUESTION NO: 14 - (SIMULATION)

Which file contains a set of services and hosts that will be allowed to connect to the server by going through a TCP Wrapper program such as tcpd? (Specify the full name of the file, including

ANSWER: See the explanation for the answer

Explanation:

The file is `/etc/hosts.allow`.

TCP Wrapper-controlled services check this file first for hosts that are explicitly permitted to connect. Rules not allowed here can subsequently be evaluated against `/etc/hosts.deny`.

QUESTION NO: 15

Which TWO statements about crontab are true?

A. Every user may have their own crontab.

B. Changing a crontab requires a reload/restart of the cron daemon.

C. The cron daemon reloads crontab files automatically when necessary.

D. hourly is the same as "0 * * * *".

E. A cron daemon must run for each existing crontab.

ANSWER: A C

Explanation:

`Every user may have their own crontab.` is correct in the normal cron model: user crontabs are separate schedules associated with individual user accounts. A user edits their own schedule using `crontab -e`, and commands in that crontab execute with that user's identity and environment. System policy can limit access through mechanisms such as `cron.allow` or `cron.deny`, but cron supports distinct per-user crontabs rather than requiring all scheduled work to be placed in a shared system file.

`The cron daemon reloads crontab files automatically when necessary.` is also correct. Cron monitors its crontab sources and recognizes updates made through the `crontab` command. Consequently, after a user saves a modified crontab, the new schedule is used without an administrator manually restarting or reloading the cron service. This is important operationally because it allows users to manage their scheduled jobs independently while the single running cron daemon continues dispatching jobs for all eligible users. The `crontab(5)` manual documents user crontab management and cron access controls, while the `cron(8)` documentation describes cron's handling of changed crontab files. See crontab(5) and cron(8).

QUESTION NO: 16

Which of the following IPv4 networks are reserved by IANA for private address assignment and private routing? (Choose three.)

A. 10.0.0.0/8

B. 127.0.0.0/8

C. 169.255.0.0/16

D. 172.16.0.0/12

E. 192.168.0.0/16

ANSWER: A D E

Explanation:

The private IPv4 address blocks defined for internal addressing are **10.0.0.0/8**, **172.16.0.0/12**, and **192.168.0.0/16**. These are the three address ranges specified by RFC 1918 for use within private networks. Organizations may assign addresses from these ranges to hosts, servers, routers, and other internal devices without obtaining globally unique public addresses from a regional Internet registry.

Private-address use is intended for networks whose internal addressing does not need to be globally reachable. When devices using these ranges need Internet access, a border gateway commonly uses Network Address Translation to map internal private addresses to one or more public addresses. The three blocks provide substantially different address capacities: 10.0.0.0/8 supports the largest private allocation, 172.16.0.0/12 covers the contiguous range from 172.16.0.0 through 172.31.255.255, and 192.168.0.0/16 is commonly used in home and small-office networks. RFC 1918 also states that these addresses are not to be routed across the public Internet. The authoritative specification is RFC 1918; IANA's registry also identifies these allocations as special-purpose IPv4 address space.

QUESTION NO: 17

Which of the following features are provided by SPICE? (Choose two.)

A. Connecting local USB devices to remote applications.

B. Accessing graphical applications on a remote host.

C. Replacing Xorg as local X11 server.

D. Downloading and locally installing applications from a remote machine.

E. Uploading and running a binary program on a remote machine.

ANSWER: A B

Explanation:

SPICE (Simple Protocol for Independent Computing Environments) is a remote-display protocol and software stack designed primarily to provide responsive graphical access to virtual machines. "Accessing graphical applications on a remote host." is therefore a core use case: a SPICE client such as `remote-viewer` displays the guest's graphical desktop remotely while transmitting user input back to the guest. SPICE is commonly used with QEMU/KVM and libvirt-managed virtual machines.

"Connecting local USB devices to remote applications." is also correct because SPICE provides USB redirection. This lets a USB device attached to the client machine be redirected into the remote guest, where the guest can use it as though it were locally attached. Typical redirected devices include storage devices, smart-card readers, and other USB peripherals, subject to client, guest, and policy support. SPICE's project documentation identifies both remote display functionality and USB redirection among its capabilities. See the SPICE project site and the SPICE user manual for details on SPICE clients, remote desktops, and device redirection.

QUESTION NO: 18

Which of the following statements about systemd-journald are true? (Choose three.)

A. It is incompatible with syslog and cannot be installed on a system using regular syslog.

B. It only processes messages of systemd and not messages of any other tools.

C. It can pass log messages to syslog for further processing.

D. It maintains metadata such as _UID or _PID for each message.

E. It supports syslog facilities such as kern, user, and auth.

ANSWER: C D E

Explanation:

`systemd-journald` is the system service that collects, stores, and indexes log records from numerous sources, including the kernel, services started by systemd, standard output and standard error streams, and applications using syslog-compatible interfaces. “It can pass log messages to syslog for further processing.” is correct because journald can forward received records to a traditional syslog daemon. This behavior is controlled by the `ForwardToSyslog=` setting in `journald.conf`, allowing tools such as rsyslog to continue providing filtering, routing, and remote log delivery.

“It maintains metadata such as _UID or _PID for each message.” is also correct. Journal records are structured entries containing trusted fields added by journald, including the originating process ID (`_PID`), user ID (`_UID`), executable path, systemd unit, boot ID, and timestamp. These fields make it possible to query logs precisely with `journalctl`, rather than relying only on unstructured text matching.

“It supports syslog facilities such as kern, user, and auth.” is correct because journald stores the syslog facility associated with applicable messages in the `SYSLOG_FACILITY` field. This preserves conventional syslog classification when messages originate through syslog-compatible logging mechanisms. See the systemd-journald documentation and the journal field reference.

QUESTION NO: 19

Which of the following comparison operators for test work on elements in the file system? (Choose two.)

A. -z

B. -eq

C. -d

D. -f

E. -lt

ANSWER: C D

Explanation:

The `-d` and `-f` primaries of the `test` command (also available through the `[ ... ]` syntax) evaluate properties of filesystem pathnames. `-d pathname` succeeds when the pathname exists and resolves to a directory. This is useful before operations intended for directories, such as changing into a directory or searching its contents. For example, `test -d /etc` succeeds because `/etc` is a directory.

`-f pathname` succeeds when the pathname exists and resolves to a regular file. A regular file is an ordinary data file rather than a directory, symbolic link, device, FIFO, or socket. Scripts commonly use it to verify that a configuration file, script, or other expected file is present before reading or processing it; for example, `[ -f /etc/passwd ]`.

These are filesystem test primaries defined by POSIX and supported by common shells. Their result is communicated through the command's exit status, allowing their use directly in shell conditionals such as `if [ -d "$dir" ]`. See the POSIX test utility specification and the GNU Bash conditional expressions manual.

QUESTION NO: 20

On a system using system-journald, which of the following command add the message Howdy to the system log? (Choose two correct answers.)

A. Append Howdy

B. logger Howdy

C. systemd-cat echo Howdy

D. echo Howdy > /dev/journal

E. journalctl and Howdy

ANSWER: B C

Explanation:

`logger Howdy` sends the supplied text as a log message through the system logging interface. On systems using systemd, the journal service receives messages submitted through the syslog-compatible logging socket, so the message can subsequently be viewed with `journalctl`. The `logger` utility is specifically intended for adding entries to the system log and can also attach facilities, priorities, tags, and other metadata when needed.

`systemd-cat echo Howdy` runs `echo Howdy` and connects its standard output to the journal. `systemd-cat` is designed to send a command's standard output and standard error to systemd-journald; therefore, the text written by `echo` becomes a journal entry. It is particularly useful when a command does not natively write to syslog or the journal. Both commands create log entries; afterward, an administrator can inspect them using `journalctl`.

The systemd documentation describes `systemd-cat` as a tool for connecting command output to the journal: systemd-cat manual. The util-linux documentation describes `logger` as entering messages into the system log: logger manual.

QUESTION NO: 21

On a regular users workstation the route command takes a long time before printing out the routing table. Which of the following errors does that indicate?

A. The local routing information may be corrupted and must be re-validated using a routing protocol.

B. One of the routers in the routing table is not available which causes the automatic router failure detection mechanism (ARF-D) to wait for a timeout.

C. There may accidentally be more than one default router in which case a default router election has to be done on the network in order to choose one router as the default.

D. DNS resolution may not be working as route by default tries to resolve names of routers and destinations and may run into a timeout.

ANSWER: D

Explanation:

DNS resolution may not be working as route by default tries to resolve names of routers and destinations and may run into a timeout. The traditional `route` utility obtains routing-table data from the kernel, but, unless numeric output is requested, it can attempt reverse name resolution for displayed network addresses and gateway addresses. This resolution uses the system name-service configuration, commonly DNS. When a configured DNS server is unreachable, slow, or unable to

return reverse-lookup results, each lookup can wait until its timeout expires. As a result, the routing data is available locally but the command appears to pause before it prints the table.

This behavior can be confirmed by running `route -n`. The `-n` option instructs `route` to show numeric addresses rather than resolving host names. If numeric output appears immediately, the delay points to name resolution rather than a routing-table retrieval problem. Checking the configured resolvers in `/etc/resolv.conf`, DNS reachability, and reverse-DNS records is then appropriate. The route(8) manual page documents that `-n` inhibits name resolution, and the resolv.conf(5) manual page describes the resolver configuration used for DNS queries.

QUESTION NO: 22

In case neither cron.allow nor cron.deny exist in /etc/, which of the following is true?

A. Without additional configuration, no users may have user specific crontabs.

B. Without additional configuration, all users may have user specific crontabs.

C. The cron daemon will refuse to start and report missing files in the system ' s logfile.

D. When a user creates a user specific crontab the system administrator must approve it explicitly.

ANSWER: B

Explanation:

Without additional configuration, all users may have user specific crontabs. The `crontab` command uses `/etc/cron.allow` and `/etc/cron.deny` as access-control lists. When neither list is present, the usual default policy assumed by this question is that ordinary local users are permitted to install and manage their own per-user crontab files. Each permitted user can use commands such as `crontab -e` to edit scheduled jobs, and the cron service reads those jobs under that user's identity when their scheduled times occur.

The precise fallback behavior when both files are absent can be implementation- or distribution-dependent; the Debian `crontab` documentation explicitly notes that it depends on site-specific configuration parameters. However, for the default behavior represented in this LPIC-1 question, the absence of both access-control files does not itself impose a user restriction, so user-specific crontabs are available to all users. Administrators who need to limit this capability should create and maintain `/etc/cron.allow`, or use `/etc/cron.deny` where appropriate. See the Debian crontab(1) manual and the crontab(5) manual page for cron-table operation and access-control context.

QUESTION NO: 23

FILL BLANK

Which command included in systemd supports selecting messages from the systemd journal by criteria such as time or unit name? (Specify only the command without any path or parameters.)

A. journalctl

ANSWER: A

Explanation:

The correct command is `journalctl`. It is the systemd utility for querying and displaying entries stored in the systemd journal. It supports filtering journal messages using a wide range of selection criteria, including time ranges and the unit that generated a message. For example, `journalctl --since "2024-01-01"` displays entries starting at a specified time, while `journalctl -u ssh.service` restricts output to messages associated with a particular systemd unit. These filters make it practical to investigate service startup failures, review recent activity, or isolate events from a specific daemon without manually searching a general log file. The command can also select entries by boot, priority, process, user, executable path, and other journal fields. Although systemd-journald collects and stores the journal data, `journalctl` is the

command-line interface intended to read, search, and filter that data. The systemd manual documents unit filtering with `-u` and time filtering with `--since` and `--until`. See the journalctl manual and the systemd-journald documentation.

QUESTION NO: 24

Which of the commands below might have produced the following output?

```
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 40997
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; ENDS: version: 0, flags:; udp:1280
;; QUESTION SECTION:
; www.example.org.        IN    NS

;; AUTHORITY SECTION:
example.org.        3600 IN    SOA  sns.dns.icann.org noc.dns.icann.org.
2016110776 7200 3600 1209600 3600

;; Query time: 144 msec
;; SERVER: 127.0.1 1#53(127.0.1.1)
;; WHEN: Sat Apr 01 21:22:46 BST 2017
;; MSG SIZE rcvd: 114
```

A. dig -t mx www.example.org

B. dig www.example.org

C. dig -t ns www.example.org

D. dig -t a www.example.org

E. dig -t soa www.example.org

ANSWER: B

Explanation:

`dig www.example.org` is the appropriate command because, when no query type is supplied, `dig` performs an Internet-class query for an IPv4 address record, whose DNS resource-record type is `A`. Its output includes the standard sections produced by `dig`, such as the query header, the QUESTION section naming `www.example.org`, an ANSWER section when the name resolves, and authority or additional data where applicable. The command invocation printed at the top of normal `dig` output also corresponds to a lookup issued without an explicit `-t` argument. This behavior is defined in the BIND `dig` manual: if no type is specified, the default query type is `A`. Therefore, a normal lookup of the host name using `dig www.example.org` can produce the displayed address-lookup output. See the BIND 9 dig manual and the DNS specification (RFC 1035) for the definition of DNS resource records and address queries.

QUESTION NO: 25

What the echo $$ command?

A. The process ID of the current shell.

B. The process ID for the following command.

C. The process ID of the last command executed.

D. The process ID of the last command which has been placed in the background.

E. The process ID of the echo command.

ANSWER: A

Explanation:

The process ID of the current shell. In a POSIX-compatible shell, `$$` is a special parameter that expands to the decimal process ID of the shell process that is executing the command. Therefore, when the shell parses `echo $$`, it replaces `$$` with its own PID before `echo` receives its argument. For example, if the active shell has PID `1234`, running `echo $$` displays `1234`. This is useful in scripts and interactive sessions when a shell needs to identify itself, such as when creating a PID-specific temporary-file name, recording diagnostic information, or signalling the running shell from another process. The value identifies the shell context rather than the `echo` utility itself. The POSIX Shell Command Language specifies that the special parameter `$$` expands to the process ID of the invoked shell. Bash documents the same behavior, describing `$$` as expanding to the process ID of the shell. See the POSIX Shell Command Language and the GNU Bash manual: Special Parameters.

QUESTION NO: 26

Given the following excerpt of the sudo configuration:

jane ANY=NOPASSWD: /bin/kill, /bin/id, PASSWD: /sbin/fdisk

Which of the following statements are true? (Choose three.)

A. Jane can run /bin/id only after specifying her password.

B. Jane can run /sbin/fdisk after specifying root's password.

C. Jane can run /sbin/fdisk after specifying her password.

D. Jane can run /bin/kill without specifying a password.

E. Jane can run /bin/id without specifying her password.

ANSWER: C D E

Explanation:

The sudoers entry authorizes user `jane` to run the listed commands on any host, as denoted by `ANY`. Sudo tags apply to the commands that follow them in a command list, continuing until another relevant tag changes the setting. Therefore, `NOPASSWD:` applies to both `/bin/kill` and `/bin/id`. Jane can run `/bin/kill` without specifying a password, and she can run `/bin/id` without specifying her password.

The `PASSWD:` tag then changes the authentication requirement for `/sbin/fdisk`. Consequently, Jane can run `/sbin/fdisk` after specifying her password. In the normal sudo configuration, sudo authenticates the invoking user, Jane, rather than requesting the target user's password. This arrangement is useful when low-risk or routine administrative commands are permitted without an authentication prompt, while a more consequential disk-management command requires Jane to authenticate first. The sudoers manual documents that `NOPASSWD` and `PASSWD` tags control whether authentication is required, and that a tag's effect continues through the command list until overridden. See the sudoers manual and the sudo manual.

QUESTION NO: 27

How many IP addresses can be used for unique hosts inside the IPv4 subnet 192.168.2.128/26?

A. 6

B. 14

C. 30

D. 62

E. 126

ANSWER: D

Explanation:

The correct answer is **62**. An IPv4 address contains 32 bits. The `/26` CIDR prefix reserves 26 bits for the network portion, leaving 6 bits for host addressing. Six host bits provide 2^6 = 64 total addresses in the subnet. In a conventional IPv4 subnet, two of those addresses have reserved purposes: the all-zero host-bit address is the network address and the all-one host-bit address is the directed broadcast address. Therefore, 64 minus 2 leaves 62 addresses that can be assigned to unique hosts.

For this specific subnet, `192.168.2.128` is the network address, and the subnet spans through `192.168.2.191`, which is its broadcast address. The assignable host range is consequently `192.168.2.129` through `192.168.2.190`, inclusive: exactly 62 usable host addresses. This calculation is the standard subnetting approach used for traditional IPv4 LAN addressing. CIDR prefix notation and IPv4 addressing are defined in RFC 4632; the special network and broadcast address conventions are described in RFC 1812.

QUESTION NO: 28

FILL BLANK

What option to useradd creates a new user's home directory and provisions it with a set of standard files? (Specify only the option name without any values or parameters.)

A. -D

B. -m

ANSWER: B

Explanation:

The correct option is `-m`. When used with `useradd`, this option instructs the system to create the new account's home directory if it does not already exist. During that creation, `useradd` copies the standard initialization files from the system skeleton directory, conventionally `/etc/skel`, into the new home directory. These files commonly include default shell configuration files such as `.profile`, `.bashrc`, or similar site-defined files, giving the account a usable initial environment.

The long-form equivalent is `--create-home`, but `-m` is the requested option name and is the conventional concise answer for this task. The specific source directory for the standard files can be configured through the `SKEL` setting in `/etc/default/useradd`; administrators may customize it to provide organization-wide default files for newly created accounts. The behavior and syntax are documented in the useradd(8) manual page. The related configuration defaults, including the skeleton-directory setting, are described in the useradd configuration documentation.

QUESTION NO: 29

Which of the following protocols is designed to access the video card output of a virtual machine?

A. KDE

B. X11

C. Xfce

D. SPICE

E. XDMCP

ANSWER: D

Explanation:

SPICE is correct. The Simple Protocol for Independent Computing Environments is a remote-display protocol designed particularly for virtualized desktop and server environments. A SPICE server running with a virtual machine provides a client with access to that guest's graphical display, including the virtual video device output, and can also carry related interactive desktop functions such as keyboard and mouse input, audio, clipboard integration, and display resizing. It is commonly used with QEMU/KVM and libvirt-managed virtual machines; for example, a guest may be configured with a SPICE graphics device and then accessed using a compatible client such as `remote-viewer`. This purpose directly matches accessing the video-card output of a virtual machine. The protocol is intended to provide responsive remote graphical console access while allowing features appropriate to virtual hardware and guest displays. Red Hat's virtualization documentation describes SPICE as a remote display protocol used to view virtual-machine graphical consoles, while the SPICE project documents its architecture and client/server components. See the SPICE project website and libvirt's graphical framebuffer documentation.

QUESTION NO: 30

What is true about NetworkManager on a Linux system that uses its distribution's mechanisms to configure network interfaces? (Choose two.)

A. NetworkManager reconfigures all network interfaces to use DHCP unless they are specifically managed by NetworkManager.

B. NetworkManager must be explicitly enabled for each interface it should manage.

C. NetworkManager by default does not change interfaces which are already configured.

D. NetworkManager disables all interfaces which were not configured by NetworkManager.

E. NetworkManager can be configured to use the distribution's network interface configuration.

ANSWER: C E

Explanation:

NetworkManager by default does not change interfaces which are already configured is correct in the common configuration where NetworkManager uses the distribution's native network-configuration integration. For example, NetworkManager's ifupdown plugin can read Debian-style `/etc/network/interfaces` configuration. With the plugin's default `managed=false` behavior, interfaces declared through that mechanism remain under the control of the traditional network configuration system rather than being actively reconfigured by NetworkManager. This allows an administrator to retain existing distribution-managed networking while using NetworkManager only where appropriate.

NetworkManager can be configured to use the distribution's network interface configuration is also correct. NetworkManager supports plugins that integrate with distribution-specific configuration formats and tools, including the ifupdown plugin and the ifcfg-rh plugin used by Red Hat-family systems. These plugins allow connection profiles and device handling to be derived from, or coordinated with, the networking mechanisms provided by the distribution. The active plugins and their behavior are configured in `NetworkManager.conf`. See the NetworkManager.conf documentation and the ifupdown plugin documentation.

QUESTION NO: 31

Which of the following commands is used to rotate, compress, and mail system logs?

A. logrotate

B. striplog

C. syslogd –-rotate

D. rotatelog

E. logger

ANSWER: A

Explanation:

`logrotate` is the standard Linux utility for managing log-file rotation. It is normally invoked automatically by a scheduled job, such as cron or a systemd timer, and processes configuration rules from `/etc/logrotate.conf` and commonly from files under `/etc/logrotate.d/`. Those rules determine when a log should be rotated, for example daily, weekly, when it reaches a specified size, or when it is empty.

Its configuration supports compression of archived logs through directives such as `compress`, typically producing gzip-compressed rotated files. It can also email a rotated log to a designated recipient using the `mail` directive. This makes `logrotate` the command that directly covers the three stated log-management operations: retaining and renaming older log files, compressing archives, and sending rotated content by mail. The utility can additionally control retention periods, create replacement log files with appropriate ownership and permissions, and run scripts before or after rotation, making it a core administrative tool for keeping system log storage manageable. See the logrotate manual page and the upstream logrotate documentation.

QUESTION NO: 32

Which of the following are tasks handled by a display manager like XDM or KDM? (Choose TWO correct answers.)

A. Start and prepare the desktop environment for the user.

B. Configure additional devices like new monitors or projectors when they are attached.

C. Handle the login of a user.

D. Lock the screen when the user was inactive for a configurable amount of time.

E. Create an X11 configuration file for the current graphic devices and monitors.

ANSWER: A C

Explanation:

Display managers such as XDM and KDM provide and manage the graphical login process for an X Window System display. They present a graphical greeter in which a user supplies credentials, authenticate the user through the system's configured authentication mechanisms, and then establish that user's graphical session. Thus, "Handle the login of a user." is a core display-manager responsibility.

Once authentication has succeeded, a display manager starts the selected session, commonly by invoking the user's X session configuration or a desktop-session startup command. This begins and prepares the user's graphical working environment, including the desktop environment or window manager selected for the login session. Therefore, "Start and prepare the desktop environment for the user." is also a display-manager task. XDM documentation describes how it manages X displays and starts user sessions after login, while the freedesktop.org Desktop Entry specification documents the session entries that graphical display managers can present for desktop selection. See the XDM manual page and the freedesktop.org Desktop Entry specification.