PassQueen
a.k.a dumpsqueen.com

SUSE Certified Linux Administrator 11

Novell 050-720

Version Demo

Total Demo Questions: 15

Total Premium Questions: 159

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Installation	17
Topic 2, Administration	124
Topic 3, Networking	18
Total	159

QUESTION NO: 1

You want to secure a POP3 connection to a server that does not support SSL. While you do not have ssh access to the POP3 server itself, you at least have SSH access to a server near the POP3 server. Given the following information:

POP3 server: pop3.digitalairlines.com

SSH server: ssh.digitalairlines.com

Your host: da10.digitalairlines.com

POP3 port: 110

Local port: 11110

Which command will forward connections from the local port 11110 to port 110 of the POP3 server, using a secure tunnel between your host and the SSH server?

A. ssh -L 11110:ssh.digitalairlines.com:110 geeko@pop3.digitalairlines.com

B. ssh -L 110:ssh.digitalairlines.com:11110 geeko@pop3.digitalairlines.com

C. ssh -L 11110:pop3.digitalairlines.com:110 geeko@ssh.digitalairlines.com

D. ssh -L 110:pop3.digitalairlines.com:11110 geeko@ssh.digitalairlines.com

ANSWER: C

Explanation:

The command `ssh -L 11110:pop3.digitalairlines.com:110 geeko@ssh.digitalairlines.com` creates the required local SSH port-forwarding tunnel. SSH local forwarding uses the form `-L [bind_address:]local_port:destination_host:destination_port`. Here, SSH listens on local port `11110` on the client host, so a local POP3 client can connect to `localhost:11110`. Each such connection is then carried through the encrypted SSH session established to `ssh.digitalairlines.com` as user `geeko`.

After the traffic reaches the SSH server, that server opens the onward connection to `pop3.digitalairlines.com` on TCP port `110`. This works even though there is no SSH account on the POP3 server itself, because the SSH server is the endpoint that makes the connection to the POP3 service and is assumed to have appropriate network access to it. The POP3 traffic is protected over the potentially untrusted path between the local host and the SSH server; the short server-side hop to the POP3 service uses ordinary POP3. Configure the mail client to use `localhost` and port `11110`. OpenSSH documents local forwarding and its destination-host/port syntax in the ssh(1) manual; SUSE also documents SSH tunneling concepts in its SSH security documentation.

QUESTION NO: 2

Which actions can be performed using YaST Software Management on SLES 11? (Choose 3.)

A. Install packages from configured repositories.

B. Remove installed packages.

C. Update installed packages.

D. Create disk partitions.

E. Configure an Ethernet interface address.

F. Edit user crontabs.

ANSWER: A B C

Explanation:

YaST Software Management provides a graphical or text-based interface for software package administration. It can install packages from configured repositories, remove installed packages, and update installed packages when newer versions are available in enabled repositories. It also resolves package dependencies as part of a transaction. YaST Software Management does not partition disks or configure network interfaces; those tasks use separate YaST modules. See the SUSE SLES 11 SP4 software management documentation.

QUESTION NO: 3

Which parts can be combined in the syslog-ng configuration file (/etc/syslog-ng/syslog-ng.conf) to specify what information is logged? (Choose 3.)

A. Alerts

B. Users

C. Filters

D. Sources

E. Log Paths

F. Destinations

G. Date and Time

ANSWER: C D F

Explanation:

syslog-ng defines logging behavior by combining **Sources**, **Filters**, and **Destinations** in a log path. A source identifies where syslog-ng receives messages, such as the local system socket, kernel messages, a file, or a network listener. A filter selects the messages of interest from that source stream, for example by facility, severity, program name, host, or a matching message expression. A destination specifies where the selected messages are delivered, including a local log file, a remote syslog server, a user terminal, or another supported output.

These definitions are connected in a `log { ... };` statement. A typical path therefore references a source, optionally references one or more filters, and references one or more destinations. This separation lets an administrator collect a broad set of events once, then route different subsets to suitable files or remote systems. For example, messages received from a local source can be filtered for authentication-related events and written to a dedicated destination file. See the syslog-ng.conf manual page and the syslog-ng Administration Guide for configuration syntax and log-path examples.

QUESTION NO: 4

What information do you need to set up an LDAP client? (Choose 4.)

A. LDAP version

B. LDAP base DN

C. Kind of encryption

D. Synchronization frequency

E. Address of the LDAP server

F. File name of the LDAP database

ANSWER: A B C E

Explanation:

An LDAP client must know the **Address of the LDAP server** so it can contact the directory service. This is typically supplied as a hostname or IP address, together with the relevant LDAP or LDAPS service port. The client also needs the **LDAP base DN**, which defines the point in the directory information tree from which it searches for users, groups, and other directory entries. For example, a directory serving an organization might use a base DN such as `dc=example,dc=com`.

The **LDAP version** is needed because the client and server must communicate using a compatible LDAP protocol version; LDAP version 3 is the normal choice for modern deployments. Finally, the **Kind of encryption** must be selected to establish whether connections use no transport encryption, TLS, or SSL/LDAPS. This choice determines the connection URI or security settings and, when TLS is enabled, whether server certificates must be validated. These connection, directory-location, protocol, and transport-security settings are the core information required when configuring an LDAP client in YaST or through LDAP client configuration files. See the SUSE LDAP client documentation at SUSE SLES 11 SP4: LDAP—A Directory Service and the OpenLDAP client configuration reference at ldap.conf(5).

QUESTION NO: 5

You created a script (/root/bin/backup.sh) that backs up your files when executed. Using cron, you want to execute this script once every Saturday at 3:00 pm. Which line would you have to enter after invoking crontab - e in a terminal window?

A. 15 0 * * 6 /root/bin/backup.sh

B. 0 15 * * 6 /root/bin/backup.sh

C. 0 15 * 6 * /root/bin/backup.sh

D. 15 0 * * 6 root /root/bin/backup.sh

ANSWER: B

Explanation:

`0 15 * * 6 /root/bin/backup.sh` is the correct crontab entry because a user crontab uses five time-and-date fields, followed by the command to run. In order, the fields specify minute, hour, day of the month, month, and day of the week. The minute value `0` schedules execution at the start of the hour, while `15` in the hour field specifies 15:00, which is 3:00 pm in cron's 24-hour time notation. The asterisks in the day-of-month and month fields mean the job is not restricted to a particular calendar date or month. The value `6` in the day-of-week field represents Saturday in standard cron implementations, where Sunday may be represented by either 0 or 7. Because this entry is installed through `crontab -e`, it is a per-user crontab and the command is written directly after the five scheduling fields. Cron will invoke the specified script at 15:00 every Saturday, provided the cron service is running and the script is executable or otherwise runnable by the user owning that crontab. See the crontab(5) manual and SUSE's scheduled tasks documentation.

QUESTION NO: 6

Which commands can be used to obtain routing information or manage routes on a SLES 11 system? (Choose 3.)

A. `route -n`

B. `ip route show`

C. `ip route add`

D. `ifconfig`

E. `netstat -l`

ANSWER: A B C

Explanation:

`route -n` displays the kernel routing table without attempting name resolution. `ip route show` displays routing information using the iproute2 utility suite. The `ip route add` command can add a route when supplied with appropriate

route parameters and privileges. In contrast, `ifconfig` is used primarily to display or configure network interfaces, and `netstat -l` lists listening sockets rather than routes. See the route(8) manual page and the ip-route(8) manual page.

QUESTION NO: 7

You want to determine which package owns the file `/usr/bin/ssh`. Which command should you use?

A. `rpm -qf /usr/bin/ssh`

B. `rpm -ql /usr/bin/ssh`

C. `rpm -V /usr/bin/ssh`

D. `rpm -e /usr/bin/ssh`

ANSWER: A

Explanation:

The command `rpm -qf /usr/bin/ssh` queries the RPM database for the package that owns the specified installed file. The `-q` option performs a query and `-f` treats its argument as a filename. This is useful when an administrator needs to identify the package responsible for a binary or configuration file. See the rpm(8) manual page.

QUESTION NO: 8

What does the following command do? tar -czvf /tmp/backup.tgz /etc/HOSTNAME

A. It extracts the /etc/HOSTNAME file from the /tmp/backup.tgz archive.

B. It includes the /etc/HOSTNAME file in the newly created /tmp/backup.tgz archive.

C. It produces an error message because the options are in the wrong sequence.

D. It displays a line similar to the following, but otherwise does nothing: -rw-r--r-- root/root 23 2005-03-11 14:20 etc/HOSTNAME

ANSWER: B

Explanation:

The command creates a new compressed tar archive at `/tmp/backup.tgz` and adds the specified `/etc/HOSTNAME` file to it. In the option string, `c` selects archive creation, `z` filters the archive through gzip compression, `v` requests verbose output while tar processes the file, and `f` means that the next command-line argument is the archive filename. Therefore, `/tmp/backup.tgz` is the output archive and `/etc/HOSTNAME` is the file being archived. Although the source pathname is written with an absolute path, GNU tar normally removes the leading slash when storing it, producing an archive member such as `etc/HOSTNAME`; this protects against unintentionally extracting files directly into absolute system locations. The `.tgz` suffix is a conventional indication of gzip-compressed tar content, consistent with the `z` option. Verbose mode can print the processed pathname, but it does not change the operation from archive creation. GNU tar documents the create operation and archive-file handling in its Creating Archives section, and documents gzip compression through the gzip option.

QUESTION NO: 9 - (SIMULATION)

SIMULATION

In which file is the host name of a Linux computer configured? (Hint: You do not need to include the entire path.)

ANSWER: See the explanation for the answer

Explanation:

The hostname is configured in:

`/etc/HOSTNAME`

Because the question says the full path is not required, enter:

`HOSTNAME`

On SUSE Linux Enterprise Server 11, `/etc/HOSTNAME` contains the system's persistent hostname.

QUESTION NO: 10

The /etc/sysconfig/network/ifcfg-eth-id-macaddress configuration file contains a BOOTPRO option. Which are possible values of BOOTPROTO?

A. static or dhcp

B. master or slave

C. onboot, ifpluged, or manual

D. ethernet, wireless, or manual

ANSWER: A

Explanation:

`static` or `dhcp` are valid BOOTPROTO values for the traditional SUSE network-interface configuration files used in SLES 11. BOOTPROTO selects how the interface receives its IPv4 configuration when the network service brings the interface up. With `BOOTPROTO='static'`, the administrator supplies addressing details in the same ifcfg file, typically including `IPADDR`, `NETMASK`, and, where needed, routing information. This is appropriate where a host must retain a predictable address, such as a server or an infrastructure system. With `BOOTPROTO='dhcp'`, the interface requests its address and related network settings dynamically from a DHCP server. DHCP can provide an address, prefix or netmask, default route, DNS information, and other site-specific parameters, reducing manual configuration effort.

In the sysconfig/ifcfg model, BOOTPROTO is specifically the protocol or method used to obtain interface addressing. This aligns with the SUSE documentation for configuring network cards through files under `/etc/sysconfig/network/` and with the documented ifcfg variable conventions. See the SUSE Linux Enterprise Server 11 SP4 network configuration documentation and the sysconfig network(5) manual page.

QUESTION NO: 11

You want to copy the mytext file from your computer to the /tmp/ directory on the remote computer da20. digitalairlines.com. Which command would you use to accomplish this?

A. ssh mytext da20.digitalairlines.com:/tmp

B. scp da20.digitalairlines.com:/tmp mytext

C. scp mytext da20.digitalairlines.com:/tmp

D. sftp mytext da20.digitalairlines.com:/tmp

ANSWER: C

Explanation:

The command `scp mytext da20.digitalairlines.com:/tmp` copies the local file named `mytext` to the `/tmp` directory on the remote host `da20.digitalairlines.com`. Secure Copy Protocol (`scp`) uses SSH for transport and authentication, so the user is normally prompted for credentials unless SSH keys or another configured authentication mechanism are available. Its basic command syntax is `scp source destination`. Here, `mytext` is the source in the current local directory, while `da20.digitalairlines.com:/tmp` identifies a remote destination: the hostname is followed by a colon and the destination path. Because the destination is a directory, the transferred file is created there with its existing filename, resulting in `/tmp/mytext` on the remote system. A username can be supplied explicitly when the local and remote usernames differ, such as `scp mytext user@da20.digitalairlines.com:/tmp`. This command is appropriate for a noninteractive, one-file transfer and is a standard SSH-based administration tool on SUSE Linux Enterprise systems. See the OpenSSH scp manual page and the SUSE Linux Enterprise Server 11 OpenSSH documentation.

QUESTION NO: 12

Which fields are stored in each traditional local-user entry in `/etc/passwd`? (Choose 3.)

A. Numeric user ID

B. Primary group ID

C. Login shell

D. Supplementary group list

E. Encrypted password hash when shadow passwords are used

F. User disk quota limit

ANSWER: A B C

Explanation:

A traditional `/etc/passwd` entry is colon-separated and includes the login name, numeric user ID (UID), numeric primary group ID (GID), optional GECOS/comment information, home directory, and login shell. With shadow passwords in use, the password field in this file normally contains an `x`, while the password hash is stored in `/etc/shadow`. Supplementary group membership is recorded in `/etc/group`, not in a passwd entry. See the passwd(5) manual page and the group(5) manual page.

QUESTION NO: 13

You can schedule jobs to be executed by cron in several ways. Which directories are used for scripts that are to be executed at regular intervals? (Choose 2.)

A. /etc/cron.d/

B. /etc/crontabs/

C. /etc/cron.daily/

D. /var/spool/cron/tabs/

ANSWER: A C

Explanation:

`/etc/cron.daily/` is a standard system cron directory intended for executable maintenance scripts that must run once per day. On SUSE systems, the system crontab invokes the appropriate periodic-processing mechanism, which runs the executable files placed in this directory according to the configured daily schedule. This arrangement lets administrators install a script without creating a separate user crontab entry for every daily task.

`/etc/cron.d/` is also a system-wide cron configuration directory. Files placed there use crontab syntax and define scheduled commands, commonly scripts or program jobs owned by packages or administrators. It is useful for services that need their own centrally managed schedules while avoiding direct modification of `/etc/crontab`. Cron reads these files as additional system crontabs, including the user field required in system-wide entries. Together, these directories provide standard mechanisms for regular automated work: periodic script execution through `/etc/cron.daily/` and explicitly scheduled system jobs through `/etc/cron.d/`.

For the cron file format and the role of system-wide cron files, see the openSUSE crontab(5) manual page. SUSE administration guidance is available in the SUSE Linux Enterprise Server Common Tasks documentation.

QUESTION NO: 14

What is the effect of selecting the option "Enable IP Forwarding" in the routing configuration in YaST?

A. IP packets that do have this host as their destination are routed to another host.

B. IP packets that do not have this host as their destination are filtered according to the firewall rules.

C. IP packets that do have this host as their destination are dropped to protect the host from attacks.

D. IP packets that do not have this host as their destination are transferred from one network connected to the host to another network connected to this host, according to the routing table.

ANSWER: D

Explanation:

Selecting *Enable IP Forwarding* configures the system to act as an IP router. The kernel is permitted to receive a packet on one network interface and, when the packet is addressed to a host other than itself, look up the destination in its routing table and send the packet through the appropriate outgoing interface. This is required when the machine connects separate IP networks, such as a system serving as a gateway between a LAN and an upstream network. In Linux, this behavior corresponds to enabling the IPv4 forwarding kernel setting, commonly exposed through `net.ipv4.ip_forward`. Forwarding uses the routing table to choose the next hop or outgoing interface; it does not change the packet's destination into the local machine's destination. Firewall policy can additionally permit, restrict, or inspect forwarded traffic, but enabling IP forwarding itself is the setting that allows the routing operation. SUSE documents routing and gateway configuration in its Network Management guide, while the Linux kernel documentation describes the IP sysctl forwarding controls.

QUESTION NO: 15

Which statements about file permissions are correct? (Choose 3.)

A. Execute permission on a directory permits pathname traversal through it.

B. Write permission on a directory can permit creating new directory entries.

C. Execute permission on a regular file can permit execution of that file.

D. Read permission on a directory by itself permits access to every file below it.

E. Write permission on a regular file automatically grants execute permission.

F. File permissions apply only to directories.

ANSWER: A B C

Explanation:

The execute permission on a regular file allows the file to be executed when its contents are in an executable format or interpreted by an appropriate interpreter. On a directory, execute permission permits pathname traversal through that directory; it is required to access objects inside it. Write permission on a directory controls creating, removing, or renaming

directory entries, subject to additional restrictions such as the sticky bit. Read permission on a directory permits listing its names, but read permission alone does not permit traversal. See the chmod(1) manual page and the GNU Coreutils file permission documentation.