ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

What output will the following command sequence produce?

```
echo '1 2 3 4 5 6' | while read a b c; do
  echo result: $c $b $a;
done
```

A. result: 3 4 5 6 2 1
B. result: 1 2 3 4 5 6
C. result: 6 5 4
D. result: 6 5 4 3 2 1
E. result: 3 2 1

Answer: E

Explanation:
The while loop reads a line from the standard input and splits it into words using the IFS variable, which by default contains spaces, tabs, and newlines. The read command assigns the first word to the variable a, the second word to the variable b, and the rest of the line to the variable c. Therefore, in this case, a=1, b=2, and c=3 4 5 6. The echo command prints the values of c, b, and a in reverse order, separated by spaces. The output is result: 3 2 1. The loop terminates after reading the first line, since there is no more input to read. Reference: Bash while Loop | Linuxize, Bash Scripting - While Loop - GeeksforGeeks

Question: 2

When the command echo $ outputs 1, which of the following statements is true?

A. It is the process ID of the echo command.
B. It is the process ID of the current shell.
C. It is the exit value of the command executed immediately before echo.
D. It is the exit value of the echo command.

Answer: C

Explanation:
The $? variable in bash is a special parameter that holds the exit status of the last command executed in the current shell. The exit status is a numerical value that indicates whether the command was successful (zero) or failed (non-zero). The echo command simply prints its arguments to the standard output. Therefore, when the command echo $? outputs 1, it means that the previous command failed with an exit status of 1. Reference:

[LPI Linux Essentials - Topic 103: Command Line Basics]
[Bash Special Parameters]
[Exit status - Wikipedia]

Question: 3

Which command makes the shell variable named VARIABLE visible to subshells?

A. export $VARIABLE
B. export VARIABLE
C. set $VARIABLE
D. set VARIABLE
E. env VARIABLE

Answer: B

Explanation:
The export command makes the shell variable named VARIABLE visible to subshells. This means that any child process that is spawned from the current shell will inherit the value of VARIABLE. The export command does not need a dollar sign ($) before the variable name, as that would expand the variable to its value. The set command only affects the current shell and does not export the variable to subshells. The env command can be used to run a command in a modified environment, but it does not export the variable to subshells either. Reference:
[LPI Linux Essentials - Topic 105: Shells, Scripting and Data Management]
[LPI Linux Administrator - Exam 102 Objectives - Topic 105: Shells and Shell Scripting]

Question: 4

What output will the command seq 10 produce?

A. A continuous stream of numbers increasing in increments of 10 until stopped.
B. The numbers 1 through 10 with one number per line.
C. The numbers 0 through 9 with one number per line.
D. The number 10 to standard output.

Answer: B

Explanation:
The seq command in Linux is used to print a sequence of numbers, which can be piped to other commands or used in for loops and bash scripts1. The command can generate a list of integers or real numbers, with options to control the start, end, and increment of the sequence. The general syntax of the command is seq [options] specification1.
If you launch seq with a single number as a command-line parameter, it counts from one to that number. It then prints the numbers in the terminal window, one number per line2. For example, seq 10 will produce the following output:

1
2
3
4
5
6
7
8
9
10

Therefore, the correct answer is B. The numbers 1 through 10 with one number per line.

Reference: 1: 10+ Seq Commands with Examples in Linux – LinuxWizardry 2: How to Use the seq Command on Linux - How-To Geek

Question: 5

By default, the contents of which directory will be copied to a new user's home directory when the account is created by passing the -m option to the useradd command? (Specify the full path to the directory.)

Answer: /etc/skel

Explanation:

The /etc/skel directory contains files and directories that are used as a template for creating a new user's home directory. The useradd command uses the -m (or --create-home) option to create the user home directory as /home/username and copy the files from /etc/skel to it. The files in /etc/skel are typically initialization files such as .bashrc, .profile, and .bash_logout that set the user's environment variables, aliases, and other preferences. The system administrator can customize the /etc/skel directory to provide a consistent and convenient initial setup for new users. Reference: https://www.howtouselinux.com/post/create-new-user-with-home-directory-in-linux

https://linuxize.com/post/how-to-create-users-in-linux-using-the-useradd-command/

Question: 6

After issuing:

function myfunction { echo $1 $2 ; }

in Bash, which output does:

myfunction A B C

Produce?

A. A B
B. A B C
C. A C
D. B C
E. C B A

Answer: A

Explanation:
In Bash, a function is a block of code that can be invoked by its name. A function can take arguments, which are passed to the function as positional parameters. The $1 variable refers to the first argument, $2 to the second argument, and so on. The function can access the number of arguments passed to it by using the $# variable. In this case, the function myfunction simply echoes the first and second arguments to the standard output. Therefore, when the command myfunction A B C is executed, the output is A B, since the third argument C is ignored by the function. Reference: [LPI Linux Essentials - Topic 103: Command Line Basics]
[Bash Functions]

Question: 7

Which of the following commands puts the output of the command date into the shell variable mydate?

A. mydate="$(date)" B. mydate="exec date" C. mydate="$((date))" D. mydate="date" E. mydate="${date}"

Answer: A

Explanation:

(date)" Comprehensive Thecorrectwaytoputtheoutputofthecommanddateintotheshellvariablemy dateistousecommandsubstitutionwiththesyntax(command). This will execute the command in a subshell and replace the expression with its standard output. The double quotes around the expression will prevent word splitting and globbing of the output. The other options are incorrect because they will either assign a literal string to the variable, use an invalid syntax, or try to execute the command as an arithmetic expression. Reference:
[LPI Linux Essentials - Topic 105: Shells, Scripting and Data Management]
[LPI Linux Administrator - Exam 102 Objectives - Topic 105: Shells and Shell Scripting]

Question: 8

Which of the following files, when existing, affect the behavior of the Bash shell? (Choose TWO correct answers.)

A. ~/.bashconf
B. ~/.bashrc
C. ~/.bashdefaults
D. ~/.bash_etc
E. ~/.bash_profile

Answer: B, E

Explanation:
The Bash shell can be configured by various files that affect its behavior, such as setting environment variables, aliases,

functions, options, and prompts. Some of these files are global, meaning they apply to all users of the system, and some are local, meaning they apply to individual users. The global files are usually located in the /etc directory, while the local files are usually located in the user's home directory, which is denoted by the tilde (~) symbol1.
The local files that affect the Bash shell are:
~/.bash_profile: This file is executed when a user logs in to the system. It is used to set up the user's environment, such as the PATH, the default editor, the umask, and other variables. It can also run commands that are needed only once per login session, such as ssh-agent or fortune. This file can also source other files, such as ~/.bashrc, to inherit their settings12.
~/.bashrc: This file is executed when a user starts a new interactive shell, such as opening a terminal window or running a script with the shebang #!/bin/bash. It is used to set up the user's shell preferences, such as aliases, functions, options, and prompts. It can also source other files, such as /etc/bashrc, to inherit their settings12.
~/.bash_logout: This file is executed when a user logs out of the system. It is used to perform any cleanup tasks, such as clearing the screen, deleting temporary files, or printing a farewell message1. The other files listed in the question are not valid Bash configuration files and do not affect the behavior of the shell. Therefore, the correct answer is B. ~/.bashrc and E. ~/.bash_profile.
Reference: 1: Bash Shell Configuration Files - Land of Linux 2: Bash Startup Files - GNU Project

Question: 9

What is the difference between the commands test -e path and test -f path?

A. They are equivalent options with the same behaviour.
B. The -f option tests for a regular file. The -e option tests for an empty file.
C. Both options check the existence of the path. The -f option also confirms that it is a regular file.
D. The -f option tests for a regular file. The -e option tests for an executable file.

Answer: C

Explanation:
The test command is used to perform checks and comparisons on files and values. The -e option tests if a given path exists, regardless of its type (file, directory, link, etc.). The -f option tests if a given path exists and is a regular file, not a directory or a special file. For example, if we have a directory named dir and a file named file, we can use the test command as follows:
test -e dir && echo "dir exists" dir exists test -f dir && echo "dir is a regular file" (no output) test -e file && echo "file exists" file exists test -f file && echo "file is a regular file" file is a regular file Reference: https://www.howtoforge.com/linux-test-command/ https://www.computerhope.com/unix/bash/test.htm

Question: 10

How can the existing environment variable FOOBAR be suppressed for the execution of the script./myscript only?

A. unset -v FOOBAR;./myscript
B. set -a FOOBAR="";./myscript
C. env -u FOOBAR./myscript
D. env -i FOOBAR./myscript

Answer: C

Explanation:

The env command can be used to run a utility or command in a custom environment without having to modify the currently existing environment1. The -u or --unset option can be used to remove a variable from the environment12. Therefore, the command env -u FOOBAR./myscript will run the script./myscript in an environment where the variable FOOBAR is suppressed. The other options are incorrect for the following reasons:

A . unset -v FOOBAR;./myscript: This will unset the variable FOOBAR in the current shell, not just for the script execution. The semicolon (;) separates two commands, so the script will run in the same environment as the unset command.

B . set -a FOOBAR=“”;./myscript: This will set the variable FOOBAR to an empty string, not suppress

it. The -a option means that the variable will be exported to the environment of subsequent commands, so the script will still see the variable FOOBAR, but with no value.

D . env -i FOOBAR./myscript: This will run the script in an empty environment, not just suppress the variable FOOBAR. The -i or --ignore-environment option means that no environment variables will be passed to the command12. Reference: env command in Linux with Examples - GeeksforGeeks, env - Wikipedia.

Question: 11

When the command echo $$ outputs 12942, what is the meaning of 12942?

A. It is the process ID of the echo command.
B. It is the process ID of the current shell.
C. It is the process ID of the last command executed.
D. It is the process ID of the last command which has been placed in the background.

Answer: B

Explanation:

In bash, the PID of a shell script’s subshell process is stored in a special variable called $$. This variable is read-only, and you cannot modify it in a shell script1. You can use echo $$ to get the PID of the current bash shell you are using2. Therefore, when the command echo $$ outputs 12942, it means that the PID of the current shell is 12942.

Reference:

[LPI Linux Essentials - Topic 103: Command Line Basics]

[Bash Special Parameters]

How to get the process ID (PID) of a shell script

How to know the process id of current bash session?

Question: 12

What output will the following command produce?

seq 1 5 20

A. 1
6
1
1
1
6

B. 1
5
10
15

C. 1
2
3
4

D. 2
3
4
5

E. 5
10
15
20

Answer: B

Explanation:

The seq command in Linux is used to generate a sequence of numbers from FIRST to LAST in steps of INCREMENT1. The syntax for the seq command is:

seq [OPTION]... LAST or seq [OPTION]... FIRST LAST or seq [OPTION]... FIRST INCREMENT LAST

In this case, the command seq 1 5 20 has three arguments: FIRST = 1, INCREMENT = 5, and LAST = 20.

This means that the command will produce numbers from 1 to 20 in steps of 5. The output will be: 1 5 10 15

The output will not include 20 because it is not a multiple of 5. The output will be printed on separate lines by default, unless a different separator is specified with the -s option2. Reference: Seq Command in Linux [Explained With Examples]

seq Man Page - Linux - SS64.com - SS64 Command line reference

Question: 13

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:

The SQL WHERE clause is used to restrict the records that are returned from a SELECT SQL query based on a supplied criteria for the values in the records12. The WHERE clause follows the SELECT and FROM clauses and contains one or more conditions that must be true for a record to be included in the result set. The general syntax of the WHERE clause is: SELECT column1, column2, ...
FROM table_name
WHERE condition;
The condition can be a comparison, a logical operation, a pattern matching, a subquery, or a combination of these using various operators12. For example, the following query selects all the records from the customers table where the country is 'USA': SELECT * FROM customers
WHERE country = 'USA';
The other words listed in the question are not used to filter records based on values. They have different meanings and purposes in SQL:
CASE: This is a conditional expression that returns a value based on a set of conditions3. It can be used in SELECT, UPDATE, DELETE, or WHERE statements. For example, the following query uses a CASE expression to assign a rating to each customer based on their credit limit:
SELECT customer_name, credit_limit, CASE WHEN credit_limit > 10000 THEN 'High' WHEN credit_limit > 5000 THEN 'Medium' ELSE 'Low' END AS rating FROM customers;
FROM: This is a clause that specifies the table (s) or view (s) from which the data is retrieved. It follows the SELECT clause and precedes the WHERE clause. For example, the following query selects the customer name and order date from the customers and orders tables:
SELECT customer_name, order_date FROM customers JOIN orders ON customers.customer_id = orders.customer_id;
IF: This is a control flow statement that executes a block of code based on a condition. It can be used in stored procedures, functions, triggers, or batch files. For example, the following code snippet uses an IF statement to check if a variable is positive or negative:
DECLARE @num INT; SET @num = -10; IF @num > 0 BEGIN PRINT 'Positive'; END ELSE BEGIN PRINT 'Negative'; END
Reference: 1: SQL WHERE Clause - W3Schools 2: How to Write a WHERE Clause in SQL | LearnSQL.com 3: [SQL CASE Statement - W3Schools] : [SQL FROM Clause - W3Schools] : [SQL IF...ELSE Statement - W3Schools]

Question: 14

Which of the following commands lists all defined variables and functions within Bash?

A. env
B. set
C. env -a
D. echo $ENV

Answer: B

Explanation:
The set command lists all defined variables and functions within Bash, including local, environment, and shell variables, as well as aliases and functions. The output of set can be very long, so it is often piped to less, grep, or other commands for filtering or paging. The set command can also be used to set or unset shell options and positional parameters. The -o posix option to set limits the output to only variables, as defined by the POSIX standard123.
The env command lists only the environment variables, which are a subset of the shell variables that are passed to child processes. The env command can also be used to run a command in a modified environment, or to print or set environment variables. The -a option to env is not valid in most implementations45.
The echo command prints a line of text to the standard output. The $ENV variable is not a predefined variable in Bash, but it can be set by the user or by other programs. If it is not set, echo $ENV will print a blank line1 .
Reference:

Question: 15

Which of the following SQL queries counts the number of occurrences for each value of the field order_type in the table orders?

A. SELECT order_type,COUNT(*) FROM orders WHERE order_type=order_type;
B. SELECT order_type,COUNT(*) FROM orders GROUP BY order_type;
C. COUNT(SELECT order_type FROM orders);
D. SELECT COUNT(*) FROM orders ORDER BY order_type;
E. SELECT AUTO_COUNT FROM orders COUNT order_type;

Answer: B

Explanation:
The correct SQL query to count the number of occurrences for each value of the field order_type in the table orders is:
SELECT order_type,COUNT(*) FROM orders GROUP BY order_type;
This query uses the SELECT statement to retrieve the values of the order_type field and the COUNT(*) function to count the number of rows for each order_type. The GROUP BY clause groups the rows by the order_type field, so that the count is calculated for each distinct value of order_type. The result of this query is a table with two columns: order_type and count, where each row shows the number of orders for a specific order_type.
The other options are incorrect for the following reasons:
A: This query uses a WHERE clause that is always true, since order_type=order_type for every row. Therefore, this query returns the same result as SELECT order_type,COUNT(*) FROM orders;, which is a table with one row that shows the total number of orders, regardless of the order_type.
C: This query is syntactically invalid, since the COUNT function cannot take a subquery as an argument. The correct way to use a subquery with COUNT is COUNT((SELECT order_type FROM orders));, which returns the total number of orders, regardless of the order_type.
D: This query uses the ORDER BY clause to sort the rows by the order_type field, but it does not group them by order_type. Therefore, this query returns the same result as SELECT COUNT(*) FROM orders;, which is a table with one row that shows the total number of orders, regardless of the order_type.
E: This query is syntactically invalid, since there is no such function as AUTO_COUNT in SQL, and the COUNT function cannot take a field name as an argument. The correct way to use COUNT with a field name is

COUNT(order_type);, which returns the number of non-null values in the order_type field.

Reference:

[SQL COUNT Function]

[SQL GROUP BY Statement]

[SQL SELECT Statement]

Question: 16

What is the purpose of the file /etc/profile?

A. It contains the welcome message that is displayed after login.
B. It contains security profiles defining which users are allowed to log in.
C. It contains environment variables that are set when a user logs in.
D. It contains default application profiles for users that run an application for the first time.

Answer: C

Explanation:

The file /etc/profile is a configuration file that is read by the Bash shell when a user logs in. It contains commands and settings that apply to all users of the system, such as environment variables, PATH information, terminal settings, and security commands. Environment variables are variables that affect the behavior of programs and processes. For example, the PATH variable defines the directories where the shell looks for executable files, and the JAVA_HOME variable defines the location of the Java installation. The /etc/profile file can also source other files from the /etc/profile.d/ directory, which can contain additional scripts for setting environment variables or other system-wide settings. The /etc/profile file is not the only file that can set environment variables for a user. There are also user-specific files, such as ~/.profile, ~/.bash_profile, and ~/.bashrc, that are read by the shell after /etc/profile. These files can override or append to the settings in /etc/profile, or define new variables for the user. The order and precedence of these files depend on the type of shell (login or interactive) and the options used to start the shell. You can learn more about the difference between these files here1 and here2. Reference: https://www.thegeekdiary.com/understanding-etc-profile-configuration-file-in-linux/

https://unix.stackexchange.com/questions/704610/what-does-the-etc-profile-do

Question: 17

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

Answer: alias, alias -p

Explanation:

The alias command is used to create, list, or remove aliases in the current shell. An alias is a short name that refers to another command, usually with some options or arguments. Aliases are useful for saving typing time, avoiding spelling errors, or customizing the behavior of commands. To list all the aliases defined in the current shell, we can use the alias command without any arguments. This will print the aliases in the format of alias name='command'123. For example:

$ alias alias cp='cp -i' alias l='ls -CF' alias la='ls -A' alias ll='ls -alF' alias mv='mv -i' alias rm='rm -i' The output shows that some

common commands, such as cp, mv, and rm, have aliases that add the -i option, which prompts the user before overwriting or deleting files. The l, la, and ll aliases are shortcuts for different variations of the ls command, which lists files and directories123.
Reference: 1: List All Available Commands and Aliases in Linux - Baeldung 2: get all aliases in linux shell - Stack Overflow 3: How to list all aliases on Linux - Linux Tutorials - Learn Linux Configuration

Question: 18

Which of the following are requirements in order to run a shell script like a regular command from anywhere in the filesystem? (Choose THREE correct answers.)

A. The user issuing the command must be in the group script.
B. The script file must be found in the $PATH.
C. The script file must have the executable permission bit set.
D. The script must begin with a shebang-line (#!) that points to the correct interpreter.
E. The file system on which the script resides must be mounted with the option scripts.

Answer: B, C, D

Explanation:
In order to run a shell script like a regular command from anywhere in the filesystem, the following requirements must be met:
The script file must be found in the $PATH. The $PATH is a variable that contains a list of directories where the shell looks for executable files when a command is issued. If the script file is not in one of these directories, the shell will not be able to find it unless the full path is specified.
The script file must have the executable permission bit set. This is a file attribute that determines whether the file can be executed by the user, the group, or others. The executable permission bit can be set using the chmod command, for example: chmod +x script.sh.
The script must begin with a shebang-line (#!) that points to the correct interpreter. This is a special line at the beginning of the script that tells the shell which program to use to run the script, such as #!/bin/bash for bash scripts, or #!/usr/bin/perl for perl scripts. The shebang-line must match the exact path of the interpreter, otherwise the script will not run.
The other options are not requirements for running a shell script like a regular command. There is no such group as script, and the file system mount option scripts does not exist. Reference: [LPI Linux Essentials - Topic 105: Shells, Scripting and Data Management]
[LPI Linux Professional - Exam 102 Objectives - Topic 105: Shells and Shell Scripting]

Question: 19

Which of the following SQL statements will select the fields name and address from the contacts table?

A. SELECT (name, address) FROM contacts;
B. SELECT (name address) FROM contacts;
C. SELECT name, address FROM contacts;

D. SELECT name address FROM contacts;

Answer: C

Explanation:

The correct syntax for selecting specific columns from a table in SQL is to use the SELECT keyword followed by a comma-separated list of column names and then the FROM keyword followed by the table name. Therefore, the only option that follows this syntax is C. SELECT name, address FROM contacts; The other options are incorrect because they either use parentheses around the column names, which are not needed, or they omit the comma between the column names, which causes a syntax error. Reference: https://www.sqltutorial.org/sql-select/

https://www.w3schools.com/mysql/mysql_select.asp

Question: 20

Which directory in /etc is used to keep a sample copy of files and directories for when a new user has a home directory created? (Please provide the full path)

Answer: /etc/skel, /etc/skel/

Explanation:

The /etc/skel directory is used to keep a sample copy of files and directories for when a new user has a home directory created. The /etc/skel directory contains files and directories that are automatically copied over to a new user's home directory when such user is created by the useradd or adduser command. The /etc/skel directory allows the system administrator to create a standard environment for all new users on the system. For example, the /etc/skel directory may contain a default .bashrc file that sets some aliases and environment variables for the new user, or a default .profile file that executes some commands at login. The /etc/skel directory may also contain subdirectories such as .ssh or .config that store configuration files for various applications or services. The name /etc/skel comes from the word "skeleton", as it provides a basic structure for the new user's home directory. Reference: [Linux User Administration] [Linux Directory Structure]

Question: 21

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

A. /etc/bashrc
B. /etc/profile
C. ~/.bash_profile
D. /etc/.bashrc

Answer: B

Explanation:

The /etc/profile file is a configuration file that is read by the Bash shell when a user logs in. It contains

commands and settings that apply to all users of the system, such as environment variables, PATH information, terminal settings, and security commands. Environment variables are variables that affect the behavior of programs and processes. For example, the PATH variable defines the directories where the shell looks for executable files, and the JAVA_HOME variable defines the location of the Java installation. The /etc/profile file can also source other files from the /etc/profile.d/ directory, which can contain additional scripts for setting environment variables or other system-wide settings. The /etc/profile file is the best option for setting default shell variables for all users, as it is executed before any user-specific files. The other options are not suitable for this purpose, because:

/etc/bashrc is a configuration file that is read by the Bash shell when it is started as an interactive non-login shell. It contains commands and settings that apply to all interactive shells of the system, such as aliases, functions, and prompt settings. It is not executed when the shell is started as a login shell, which is the case when a user logs in. Therefore, it is not a good place to set default shell variables for all users.

~/.bash_profile is a configuration file that is read by the Bash shell when it is started as a login shell for a specific user. It contains commands and settings that apply only to that user, such as environment variables, PATH information, and startup programs. It can also source other files, such as ~/.bashrc, which is read by the shell when it is started as an interactive non-login shell for that user. It is not a good place to set default shell variables for all users, as it only affects the user who owns the file.

/etc/.bashrc is not a valid configuration file for the Bash shell. The dot (.) at the beginning of the file name indicates that it is a hidden file, which means that it is not visible by default in the file system. The Bash shell does not look for this file when it is started, and it does not execute any commands or settings from it. Therefore, it is not a good place to set default shell variables for all users. Reference: 1 2

Question: 22

Which of the following is the best way to list all defined shell variables?

A. env
B. set
C. env -a
D. echo $ENV

Answer: B

Explanation:

The set command is used to display or modify the shell variables and functions in the current shell. When used without any arguments, it prints the names and values of all shell variables, including environment variables and user-defined variables, in alphabetical order. The output also includes the shell options and the positional parameters. The set command can be used in any POSIX-compliant shell, such as bash, zsh, ksh, etc123.

The other options are not correct because:

env is used to print or modify the environment variables, not the shell variables. It does not show the user-defined variables or the shell options. It can also be used to run a command in a modified environment45.

env -a is an invalid option for the env command. The -a option is not supported by the env command in any standard or common implementation45.

echo $ENV is used to print the value of the environment variable ENV, not the list of all shell variables. The ENV variable is usually set to the name of a file that contains commands or aliases to be executed by the shell. It is mainly used by the ksh and some versions of bash .

Reference: 1: How can I list all shell variables? - Unix & Linux Stack Exchange 2: 2.1 Command Line Basics - Linux Professional Institute Certification ... 3: set - The Open Group Base Specifications Issue 7, 2018 edition 4: How to set and list environment

variables on Linux 5: env - The Open Group Base Specifications Issue 7, 2018 edition : What is the difference between .bash_profile and .bashrc? - Unix & Linux Stack Exchange : ENV - The Open Group Base Specifications Issue 7, 2018 edition

Question: 23

Which command allows you to make a shell variable visible to subshells?

A. export $VARIABLE
B. export VARIABLE
C. set $VARIABLE
D. set VARIABLE
E. env VARIABLE

Answer: B

Explanation:
The command that allows you to make a shell variable visible to subshells is export VARIABLE. This command turns the variable into a global or environment variable, which means it can be accessed by any child process or subshell that inherits the environment of the parent shell. The syntax of the export command does not require a dollar sign ($) before the variable name, unlike when referencing the value of the variable. The other commands are either invalid or do not affect the visibility of the variable to subshells. The set command can be used to assign values to variables, but it does not export them. The env command can be used to run a command in a modified environment, but it does not change the environment of the current shell. Reference:
[LPI Linux Essentials - Topic 105: Shells, Scripting and Data Management]
[LPI Linux Professional - Exam 102 Objectives - Topic 105: Shells and Shell Scripting]
What is a Subshell? - Linux Bash Shell Scripting Tutorial Wiki - nixCraft
What is Subshell in Linux? [Explained]

Question: 24

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

A. LIMIT
B. FROM
C. WHERE
D. IF

Answer: C

Explanation:
The correct keyword for restricting the records that are returned from a SELECT query based on a supplied criteria for the values in the records is WHERE. The WHERE clause is used to filter records based on one or more conditions. The

syntax of the WHERE clause is: SELECT column1, column2, ... FROM table_name WHERE condition;
The condition can be a logical expression that evaluates to true, false, or unknown. The condition can also use comparison operators, logical operators, and wildcards to specify the criteria. For example, the following query selects all the records from the employees table where the salary is greater than 50000:
SELECT * FROM employees WHERE salary > 50000;
The other options are incorrect because they have different purposes in SQL:
LIMIT is used to specify the maximum number of records to return from a query. For example, the following query returns only the first 10 records from the employees table: SELECT * FROM employees LIMIT 10;
FROM is used to specify the table or tables from which to retrieve data. For example, the following query selects all the columns from the employees table:
SELECT * FROM employees;
IF is used to execute a block of code conditionally. For example, the following query updates the salary of an employee based on their performance:
UPDATE employees SET salary = IF(performance = 'excellent', salary * 1.1, salary) WHERE employee_id = 123;
Reference:

https://bing.com/search?q=SQL+statements+restrict+records+based+on+criteria
https://stackoverflow.com/questions/11611931/sql-query-to-select-records-based-on-criteria

Question: 25

What benefit does an alias in bash provide?

A. It provides faster lookups for commands in the system directory.
B. It creates a local copy of a file from another directory.
C. It hides what command you are running from others.
D. It allows a string to be substituted for the first word of a simple command.

Answer: D

Explanation:
An alias in bash provides the benefit of allowing a string to be substituted for the first word of a simple command. This means that you can create a shortcut or alternative name for a command that is already installed on your system, and use the new name to run the command instead of the original name. For example, you can create an alias for the ls -la command, which lists all files and directories in the current directory with detailed information, by running the following command: alias ll='ls -la'
After defining the alias, you can use the ll command to execute the ls -la command. The alias will be active for the duration of the current shell session, unless you make it persistent by adding it to your shell startup file (such as ~/.bashrc for the Bash shell).
The other options are incorrect for the following reasons:
A: An alias does not provide faster lookups for commands in the system directory. The system directory is where the executable files for the commands are stored, and the shell uses the PATH variable to search for them. An alias does not affect the PATH variable or the system directory. B: An alias does not create a local copy of a file from another directory. An alias is a way to rename a command, not a file. To create a local copy of a file, you can use the cp command.
C: An alias does not hide what command you are running from others. An alias is a way to simplify the use of a command, not to conceal it. Anyone can see what command an alias represents by using the type command or the alias command without any arguments.

Reference:
LPI E - alias
105.1 Lesson 1 - Linux Professional Institute Certification Programs
How to Create Bash Aliases | Linuxize
How to create a permanent Bash alias on Linux/Unix - nixCraft bash - How do create an alias in shell scripts? - Stack Overflow

Question: 26

You are looking into a new script you received from your senior administrator. In the very first line you notice a #! followed by a file path. This indicates that:

A. The file at that location was used to make the script.
B. This script provides identical functionality as the file at that location.
C. This script will self-extract into a file at that location.
D. The program at that location will be used to process the script.

Answer: D

Explanation:
The #! followed by a file path is called a shebang or a hashbang. It is a special notation that tells the operating system which interpreter to use to execute the script. For example, if the first line of a script is #!/bin/bash, it means that the script will be run by the Bash shell, which is located at /bin/bash. Similarly, if the first line of a script is #!/usr/bin/python3, it means that the script will be run by the Python 3 interpreter, which is located at /usr/bin/python3. The shebang must be the very first line of the script, and it must start with #! without any spaces. The file path after the #! must be an absolute path, not a relative path or a symbolic link. The shebang allows the script to be executed as a standalone program, without specifying the interpreter explicitly. For example, if a script named hello.sh has a shebang of #!/bin/bash, and it has the executable permission, it can be run as ./hello.sh instead of bash hello.sh. The shebang also allows the script to be associated with a specific interpreter, regardless of the default interpreter of the system or the user. For example, if a script named hello.py has a shebang of #!/usr/bin/python3, it will always be run by Python 3, even if the system or the user has Python 2 as the default Python interpreter. The shebang is not a comment, although it looks like one. It is a special instruction that is only recognized by the operating system when the script is executed. It is ignored by the interpreter when the script is read. Therefore, the shebang does not indicate that the file at that location was used to make the script, or that the script provides identical functionality as the file at that location, or that the script will self-extract into a file at that location. The correct answer is that the program at that location will be used to process the script. You can learn more about the shebang here1 and here2. Reference: 1 2

Question: 27

What keyword is missing from this code sample of a shell script?

```
    i in *.txt; do
  echo $i
done
```

A. for
B. loop
C. until

D. while

Answer: B

Explanation:

The set command is used to display or modify the shell variables and functions in the current shell. When used without any arguments, it prints the names and values of all shell variables, including environment variables and user-defined variables, in alphabetical order. The output also includes the shell options and the positional parameters. The set command can be used in any POSIX-compliant

shell, such as bash, zsh, ksh, etc123.

The other options are not correct because:

env is used to print or modify the environment variables, not the shell variables. It does not show the user-defined variables or the shell options. It can also be used to run a command in a modified environment45.

env -a is an invalid option for the env command. The -a option is not supported by the env command in any standard or common implementation45.

echo $ENV is used to print the value of the environment variable ENV, not the list of all shell variables.

The ENV variable is usually set to the name of a file that contains commands or aliases to be executed by the shell. It is mainly used by the ksh and some versions of bash .

Reference: 1: How can I list all shell variables? - Unix & Linux Stack Exchange 2: 2.1 Command Line Basics - Linux Professional Institute Certification ... 3: set - The Open Group Base Specifications Issue 7, 2018 edition 4: How to set and list environment variables on Linux 5: env - The Open Group Base Specifications Issue 7, 2018 edition : What is the difference between .bash_profile and .bashrc? - Unix & Linux Stack Exchange : ENV - The Open Group Base Specifications Issue 7, 2018 edition

Question: 28

What word is missing from the following SQL statement?

count(*) from tablename;

(Please specify the missing word using lower-case letters only.)

Answer: select

Explanation:

The missing word is select, which is the keyword used to query data from a table in SQL. The select statement has the following syntax:

select column_list from table_name where condition;

The column_list can be one or more columns separated by commas, or an asterisk () to indicate all columns. The table_name is the name of the table that contains the dat

a. The where clause is optional and specifies a condition to filter the rows. The count() function is an aggregate function that returns the number of rows in the table or in a group. Therefore, the complete statement is:

select count(*) from tablename;

This statement will return the number of rows in the table named tablename. Reference: SQL COUNT() Function - W3Schools, SQL COUNT: The Ultimate Guide To SQL COUNT Function - SQL Tutorial, The SQL Count Function Explained With 7 Examples.

Question: 29

Which file used by XDM specifies the default wallpaper?

A. /etc/X11/xdm/Xsetup
B. /etc/X11/xdm.conf
C. /etc/X11/xdm/Defaults
D. /etc/X11/defaults.conf

Answer: A

Explanation:

The file that specifies the default wallpaper for XDM is /etc/X11/xdm/Xsetup. XDM is a display manager for the X Window System that provides a graphical login screen and manages user sessions. The /etc/X11/xdm/Xsetup file is executed when XDM starts the X server and before any user login or session starts. This file can be used to configure the X server, set X resources, and perform any other system-wide setup tasks, such as setting the wallpaper. To set the wallpaper, one can use a command like qiv -z /usr/local/share/backgrounds/wallpaper.jpg in the /etc/X11/xdm/Xsetup file, where qiv is an image viewer and /usr/local/share/backgrounds/wallpaper.jpg is the path to the desired wallpaper image1.

The other options are not correct/etc/X11/xdm.conf is the configuration file for XDM, which specifies how XDM should behave, such as the access control, the login window, and the session types2/etc/X11/xdm/Defaults is the directory where the default XDM configuration files are stored, such as Xresources, Xsession, and Xwilling2.

/etc/X11/defaults.conf is not a valid file or directory related to XDM or X Window System. Reference: XDM - ArchWiki

Configuring XDM - Linux Documentation Project

Question: 30

Which command can be used to investigate the properties for a particular window in X by clicking that window? (Specify ONLY the command without any path or parameters.)

Answer: /usr/bin/xwininfo, xwininfo

Explanation:

The command that can be used to investigate the properties for a particular window in X by clicking that window is xwininfo. xwininfo is a command-line tool that provides information about X windows. When executed, it opens a small window and waits for the user to select a window by clicking on it. Then, it displays various characteristics about the window in question, such as its geometry, position, size, depth, class, name, id, and more. xwininfo is part of the X Window System, which is a graphical user interface system for Unix-like operating systems. xwininfo can be useful for debugging, testing, or scripting

purposes. Reference: https://bing.com/search?q=command+to+investigate+properties+of+a+windo w+in+X

https://www.exam-answer.com/linux-foundation-certified-system-administrator-lfcs-simulation- investigate-window-properties

Question: 31

The X11 configuration file xorg.conf is grouped into sections. How is the content of the section SectionName associated with that section?

A. It is placed in curly brackets as in Section SectionName { ... }.
B. It is placed between a line containing Section "SectionName" and a line containing EndSection.
C. It is placed between the tags <Section name="SectionName"> and </Section>
D. It is placed after the row [SectionName].
E. It is placed after an initial unindented Section "SectionName" and must be indented by exactly one tab character.

Answer: B

Explanation:
The X11 configuration file xorg.conf is grouped into sections, and the content of the section SectionName is associated with that section by placing it between a line containing Section “SectionName” and a line containing EndSection. For example, the following is a section named ServerLayout that defines the layout of the X server:
Section “ServerLayout” Identifier “X.org Configured” Screen 0 “Screen0” 0 0 InputDevice “Mouse0” “CorePointer” InputDevice “Keyboard0” “CoreKeyboard” EndSection
The other options are incorrect for the following reasons:
A: Curly brackets are not used to delimit sections in xorg.conf. They are used to enclose values that are lists, such as Option “XkbLayout” “{us,fr}”.
C: Tags are not used to delimit sections in xorg.conf. They are used in XML files, which have a different syntax and structure than xorg.conf.
D: Rows are not used to delimit sections in xorg.conf. They are used to define key-value pairs within a section, such as Identifier “Screen0”.
E: Indentation is not required to delimit sections in xorg.conf. It is used to improve readability and clarity, but it does not affect the functionality of the file.
Reference:
xorg.conf - X Window System
Editing basics for the xorg.conf file - Linux.com
106.1 Lesson 1 - Linux Professional Institute Certification Programs

Question: 32

What is the purpose of a screen reader?

A. It reads text displayed on the screen to blind or visually impaired people.
B. It reads the parameters of the attached monitors and creates an appropriate X11 configuration.
C. It displays lines and markers to help people use speed reading techniques.
D. It manages and displays files that contain e-books.

Answer: A

Explanation:
A screen reader is a form of assistive technology that renders text and image content as speech or braille output. Screen readers are essential to people who are blind, and are useful to people who are visually impaired, illiterate, or have a

learning disability. Linux has several screen readers available, such as Orca, Speakup, and Emacspeak. These screen readers can help users interact with the graphical or console interface, read documents and web pages, and perform various tasks on the system. Reference:

Screen reader - Wikipedia
Orca Screen Reader - GNOME
Accessibility in Linux is good (but could be much better)

Question: 33

How is a display manager started?

A. It is started by a user using the command startx.
B. It is started like any other system service by the init system.
C. It is started by inetd when a remote hosts connects to the X11 port.
D. It is started automatically when a X11 user logs in to the system console.

Answer: B

Explanation:
A display manager is a program that provides a graphical login screen for users to access a graphical desktop environment. A display manager is usually started by the init system, which is the first process that runs when the system boots up. The init system is responsible for starting and stopping various system services, including the display manager. The init system can be configured to start a specific display manager by setting the default runlevel or target, or by editing the /etc/X11/default-display-manager file123.
The other options are not correct because:
A . It is started by a user using the command startx. This option is false because the startx command is used to start an X session without a display manager. The startx command launches an X server and runs the user's .xinitrc or .xsession file, which contains the commands to start the desired desktop environment or window manager. The startx command does not invoke a display manager or a graphical login screen .
C . It is started by inetd when a remote hosts connects to the X11 port. This option is false because inetd is a daemon that listens for incoming network connections and launches the appropriate service for each connection. Inetd does not start a display manager, but it can be used to enable remote access to an X session using the XDMCP protocol. XDMCP stands for X Display Manager Control Protocol, and it allows a remote host to request a graphical login screen from a display manager running on another host. However, this is not the same as starting a display manager, and it requires the display manager to be already running on the host that provides the XDMCP service . D . It is started automatically when a X11 user logs in to the system console. This option is false

because a display manager is not started by a user login, but by the init system. A user login can trigger the start of an X session, but not a display manager. A display manager is independent of the user login, and it can run on multiple virtual consoles or display devices. A display manager can also allow multiple users to log in to different X sessions simultaneously123.
Reference: 1: LPI Linux Certification/Setup A Display Manager - Wikibooks 2: Working with Display Managers - LPIC-1 102 Linux certification - Linux ... 3: How to Change the Default Display Manager in Ubuntu 20.04 : startx - ArchWiki : How to start GUI from command line? - Ask Ubuntu : inetd - Wikipedia : XDMCP - ArchWiki

Question: 34

What is the default name of the configuration file for the Xorg X11 server? (Specify the file name only without any path.)

Answer: xorg.conf

Explanation:

The default name of the configuration file for the Xorg X11 server is xorg.conf. This file is used to store initial setup for X, such as settings for video cards, monitors, input devices, and other options. The Xorg X11 server is a display server that uses a configuration file called xorg.conf and files ending in the suffix .conf for its initial setup1. The xorg.conf file is typically located in /etc/X11/xorg.conf, but its location may vary across operating system distributions2. The xorg.conf file is not mandatory, as the Xorg X11 server can automatically configure most hardware and settings. However, it can be created and edited manually if needed3. Reference: Xorg - ArchWiki xorg.conf - Wikipedia
How to Configure X11 in Linux: 10 Steps (with Pictures) - wikiHow

Question: 35

Which of the following commands shows the current color depth of the X Server?

A. xcd
B. xcdepth
C. xwininfo
D. xcolordepth
E. cat /etc/X11

Answer: C

Explanation:

The command that can be used to show the current color depth of the X Server is xwininfo. xwininfo is a command-line tool that provides information about X windows. When executed, it opens a small window and waits for the user to select a window by clicking on it. Then, it displays various characteristics about the window in question, such as its geometry, position, size, depth, class, name, id, and more. The depth value indicates the number of bits per pixel used to represent the colors of

the window. xwininfo is part of the X Window System, which is a graphical user interface system for Unix-like operating systems. xwininfo can be useful for debugging, testing, or scripting purposes. The other options are incorrect because they are either invalid commands or do not show the color depth of the X Server:

xcd is not a valid command in Linux. It may be confused with cd, which is used to change the current working directory.

xcdepth is not a valid command in Linux. It may be confused with xrandr, which is used to change the screen resolution and orientation.

xcolordepth is not a valid command in Linux. It may be confused with xcalib, which is used to load, alter, and query the color profile of the X display.

cat /etc/X11 is not a command, but a directory. cat is used to concatenate files and print them to the standard output. /etc/X11 is a directory that contains configuration files for the X Window System. However, these files do not necessarily show the current color depth of the X Server, as it may be overridden by other settings or options. Reference:

https://bing.com/search?q=command+to+show+color+depth+of+X+Server

https://x.org/releases/X11R7.5/doc/man/man5/xorg.conf.5.html

Question: 36

For accessibility assistance, which of the following programs is an on-screen keyboard?

A. xkb
B. atkb
C. GOK
D. xOSK

Answer: C

Explanation:

GOK stands for GNOME On-screen Keyboard, and it is a program that provides a virtual keyboard for users who have difficulty using a physical keyboard. GOK is designed to be accessible and customizable, and it supports different keyboard layouts, input methods, and modes. GOK can also generate mouse and gesture events, and it can be controlled by various input devices, such as switches, joysticks, or head trackers. GOK is part of the GNOME desktop environment, and it can be enabled from the Universal Access settings panel123.

The other options are not correct because:

A . xkb is not a program, but a component of the X Window System that handles keyboard configuration and mapping. XKB stands for X Keyboard Extension, and it allows users to define the behavior and appearance of their keyboards, such as the layout, the modifiers, the symbols, and the actions. XKB does not provide an on-screen keyboard, but it can be used by other programs that do45.

B . atkb is not a valid name for any known program or component related to on-screen keyboards. There is no such program or component in the LPI Linux certification program or in the common Linux distributions. The closest match is ATK, which stands for Accessibility Toolkit, and it is a library that provides a set of interfaces for accessibility support in GNOME applications. ATK does not provide an on-screen keyboard, but it can be used by GOK and other programs that do6 .

D . xOSK is a program that provides an on-screen keyboard, but it is not the one that is mentioned in the LPI Linux certification program or in the common Linux distributions. xOSK stands for X On-Screen Keyboard, and it is a simple and lightweight virtual keyboard that can be used with any X11 application. xOSK is not part of any desktop environment, and it has to be installed and launched manually. xOSK is not as accessible and customizable as GOK, and it does not support different input methods or modes .

Reference: 1: GOK - GNOME Wiki! 2: How to Set Up a Virtual On-Screen Keyboard in Linux 3: Working With On-Screen Keyboards - Oracle Help Center 4: X keyboard extension - Wikipedia 5: XKB Configuration Guide 6: Accessibility Toolkit - GNOME Developer : Accessibility - ArchWiki : xosk - X On-Screen Keyboard : How to use on-screen virtual keyboard on Linux - Xmodulo

Question: 37

What is the name of the simple graphical login manager that comes with a vanilla X11 installation? (Specify ONLY the

command without any path or parameters.)

Answer: xdm

Explanation:

The name of the simple graphical login manager that comes with a vanilla X11 installation is xdm. XDM is the traditional graphical login manager for the X Window System, independent of any window manager or environment the user might choose. When it is run at system startup, it displays a graphical login prompt rather than the text-based login prompt at the console1. XDM is part of the xorg-x11-apps package, which provides the basic applications for the X Window System2. XDM is also one of the topics covered by the LPI Linux Professional - Exam 102 Objectives - Topic 111: Graphical Desktops3.

Reference:

xorg-x11-apps - Linux Man Pages (1) - SysTutorials
LPI Linux Professional - Exam 102 Objectives - Topic 111: Graphical Desktops
GitHub - iwamatsu/slim: SLiM (Simple Login Manager) is a graphical login manager for X11 slim-fork download | SourceForge.net
Using the XDM Graphical Login Manager | FreeBSD 6 Unleashed - Flylib
Xorg - ArchWiki
How to remotely log in with full graphical desktop over X11 - Unix & Linux Stack Exchange

Question: 38

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:

The tasks that are handled by a display manager like XDM or KDM are to start and prepare the desktop environment for the user and to handle the login of a user. A display manager is a software component that manages the graphical user interface of an operating system. It provides a login screen where the user can enter their credentials and choose their preferred desktop environment or window manager. After the user is authenticated, the display manager launches the selected desktop environment or window manager and sets up the graphical session. The display manager also handles the logout, shutdown, and reboot of the system.

The other options are incorrect because they are not tasks handled by a display manager:

B . Configure additional devices like new monitors or projectors when they are attached. This task is handled by the X server, which is the core component of the X Window System. The X server is responsible for communicating with the hardware devices, such as the keyboard, mouse, monitor, and graphics card. The X server can detect and configure new devices dynamically using tools like xrandr or xorg.conf.

D . Lock the screen when the user was inactive for a configurable amount of time. This task is handled by the screensaver program, which is a utility that runs in the background and activates when the user is idle for a certain period of time. The screensaver can display various animations or images on the screen, or it can blank the screen entirely. The screensaver can also lock the screen and require the user to enter their password to resume the session. The screensaver can be configured by the user using tools like xscreensaver or gnome-screensaver.

E . Create an X11 configuration file for the current graphic devices and monitors. This task is handled by the X server,

which is the core component of the X Window System. The X server can create an X11 configuration file, which is a text file that contains the settings for the X server and the devices it communicates with. The X11 configuration file is usually located at /etc/X11/xorg.conf or /etc/X11/xorg.conf.d/. The X server can generate a default configuration file using the command Xorg -configure, or it can be edited manually by the user or the system administrator. Reference https://www.baeldung.com/linux/display-managers-explained https://quizlet.com/185979426/lx0-104-flash-cards/

Question: 39

X is running okay but you're concerned that you may not have the right color depth set. What single command will show you the running color depth while in X?

A. xcd
B. xcdepth
C. xwininfo
D. xcolordepth
E. cat /etc/X11

Answer: C

Explanation:
The xwininfo command is a utility for displaying information about windows on an X server. One of the information it displays is the depth of the window, which is the number of bits per pixel used to represent the color of the window. The depth of the root window, which is the background window of the X server, is the same as the color depth of the X server. To display the depth of the root window, one can use the command xwininfo -root and look for the line that says “depth of root window”. Alternatively, one can use the command xdpyinfo, which displays information about the X server, and look for the line that says “depths of root window”. Reference: xwininfo(1) - Linux man page xdpyinfo(1) - Linux man page
[LPI Linux Certification/Configure the X Window System, Xorg and ...]

Question: 40

Your senior administrator asked you to change the default background of his machine, which uses XDM. Which file would you edit to achieve this?

A. /etc/X11/xdm/Xsetup
B. /etc/X11/xdm.conf
C. /etc/X11/xdm/Defaults
D. /etc/X11/defaults.conf

Answer: A

Explanation:
The file /etc/X11/xdm/Xsetup contains commands that are executed by XDM before displaying the login screen. This file can

be used to set the background image, color, or run other programs on the X display. The other files are either not related to XDM or do not exist by default. Reference: XDM - ArchWiki
Customizing the XDM Login Screen | Linux Journal

Question: 41

What is the purpose of the Sticky Keys feature in X?

A. To assist users who have difficulty holding down multiple keys at once
B. To prevent repeated input of a single character if the key is held down
C. To ignore brief keystrokes according to a specified time limit
D. To repeat the input of a single character

Answer: A

Explanation:
The Sticky Keys feature in X is an accessibility option that allows users to press modifier keys (such as Ctrl, Alt, Shift, or the Windows key) one at a time, instead of holding them down simultaneously, to perform keyboard shortcuts. For example, to copy something, a user can press Ctrl, release it, and then press C, instead of pressing Ctrl+C together. This can be helpful for users who have difficulty pressing multiple keys at once, or who prefer not to do so. Reference: https://www.howtogeek.com/739764/how-to-turn-off-sticky-keys-on-windows-10/ https://geekflare.com/using-sticky-keys-in-windows/

Question: 42

Why is the xhost program considered dangerous to use?

A. It makes it difficult to uniquely identify a computer on the network.
B. It allows easy access to your X server by other users.
C. It logs sensitive information to syslog.
D. It makes your computer share network resources without any authentication.
E. It is a graphical DNS tool with known exploits.

Answer: B

Explanation:
The xhost program is used to add and delete host names or user names to the list allowed to make connections to the X server1. In the case of hosts, this provides a rudimentary form of privacy control and security. It is only sufficient for a workstation (single user) environment, although it does limit the worst abuses1. However, if xhost is used to grant access to everyone, even if they aren't on the list (i.e., access control is turned off), then any user on the network can connect to your X server and monitor your keystrokes, capture your screen, or run malicious programs2. This is why xhost is considered dangerous to use and should be avoided in favor of more secure methods, such as xauth or ssh23. Reference:
xhost linux command man page - commandlinux.com
Linux Xhost Command Help and Examples - Computer Hope xhost(1) — Arch manual pages

Question: 43

An administrator wants to determine the geometry of a particular window in X, so she issues the -metric command and then clicks on the window.

Answer: /usr/bin/xwininfo, xwininfo

Explanation:

The xwininfo command is a utility for displaying information about windows in X. It can show various attributes of a window, such as its location, size, depth, border width, visual class, colormap, map state, and event masks. The -metric option specifies that all dimensions should be displayed in metric units (millimeters) rather than pixels. By issuing the xwininfo -metric command and then clicking on a window, the administrator can determine the geometry of that window, including the decorations, in millimeters. Reference: xwininfo(1) — Arch manual pages [command line -

Question: 44

On a system running the KDE Display Manager, when is the /etc/kde4/kdm/Xreset script automatically executed?

A. When KDM starts
B. When a user's X session exits
C. When KDM crashes
D. When X is restarted
E. When X crashes

Answer: B

Explanation:

The /etc/kde4/kdm/Xreset script is a script that runs as root after a user's X session exits. It can be used to perform some cleanup tasks or other actions that need to be done when the user logs out of the graphical environment. For example, it can reassign the ownership of the console to root, or shut down the system if desired. The /etc/kde4/kdm/Xreset script is part of the KDE Display Manager (kdm), which is a graphical login manager for X. KDM can be configured to run this script by setting the Reset key in the [X-*-Core] section of the /etc/kde4/kdm/kdmrc configuration file. Reference: kdm.options - configuration options for X display manager kdm(1) — kdm — Debian jessie — Debian Manpages debian - How to get system to shutdown when Xorg is quit? - Unix ...

Question: 45

Which of the following lines is an example of a correct setting for the DISPLAY environment variable?

A. hostname:displayname
B. hostname:displaynumber
C. hostname/displayname
D. hostname/displaynumber

E. hostname

Answer: B

Explanation:

The correct format for the DISPLAY environment variable is hostname:displaynumber.screennumber, where hostname is the name of the computer where the X server runs, displaynumber is a sequence number (usually 0) that identifies a display, and screennumber is the number of the screen within that display (usually 0). The screennumber can be omitted if it is 0. For example, localhost:0 or myhost:1.0 are valid values for the DISPLAY variable. The other options are either missing the colon, using the wrong separator, or not specifying the display number. Reference:
X11 - DISPLAY (environment variable) - Datacadamia
x11 - How can I specify a display? - Stack Overflow
What is the $DISPLAY environment variable? - Ask Ubuntu

Topic 3, Administrative Tasks

Question: 46

Which of the following steps prevents a user from obtaining an interactive login session?

A. Run the command chsh -s /bin/false with the user name.
B. Set the UID for the user to 0.
C. Remove the user from the group staff.
D. Add the user to /etc/noaccess.
E. Create a .nologin file in the user's home directory.

Answer: A

Explanation:

Running the command chsh -s /bin/false with the user name will change the user's login shell to /bin/false, which is a program that does nothing and returns a non-zero exit code. This means that the user will not be able to execute any commands or start an interactive shell session. This is a common way to disable a user's login without disabling the account completely, which can be useful for users who only need to access the system via scp, sftp, or other non-interactive services.

However, this method does not prevent the user from authenticating with the system, and it may not work with some services that do not rely on the login shell, such as ssh with a forced command. Therefore, it is not a foolproof way to secure the system from unauthorized access. Reference: 1234

Question: 47

Which file specifies the user accounts that can NOT submit jobs via at or batch? (Provide the full path and filename)

Answer: /etc/at.deny

Explanation:

The /etc/at.deny file specifies the user accounts that can NOT submit jobs via at or batch. The format of the file is a list of usernames, one on each line. Whitespace is not permitted. The superuser may always use at. If the file /etc/at.allow exists, only usernames mentioned in it are allowed to use at. If /etc/at.allow does not exist, /etc/at.deny is checked12. The at and batch commands use the files /usr/lib/cron/at.allow and /usr/lib/cron/at.deny to restrict usage on some systems3.

Reference: at.allow(5) - Linux man page at.deny(5) [linux man page] - The UNIX and Linux Forums

The at.allow and at.deny files - IBM

Question: 48

Which character in the password field of /etc/passwd is used to indicate that the encrypted password is stored in /etc/shadow?

A. *

B. - C. S

D. X

Answer: D

Explanation:

The password field of /etc/passwd is used to store the user's encrypted password or a special character that indicates how the password is stored. In older Linux systems, the user's encrypted password was stored in the /etc/passwd file. On most modern systems, this field is set to x, and the user password is stored in the /etc/shadow file12. The /etc/shadow file is more secure than the /etc/passwd file because it is readable only by the root user and not by regular users1. The other options are not valid characters for the password field of /etc/passwd. Reference: Understanding the /etc/passwd File | Linuxize Understanding the /etc/passwd File - GeeksforGeeks

Question: 49

The system's timezone may be set by linking /etc/localtime to an appropriate file in which directory? (Provide the full path to the directory, without any country information)

Answer: /usr/share/zoneinfo/

Explanation:

The /usr/share/zoneinfo directory contains the binary time zone files that are used by the system to determine the local time for any region. The files are organized in subdirectories by continent, country, or ocean. Some files represent the standard time zones, while others may have historical or political variations. To set the system's timezone, one can create a symbolic link from /etc/localtime

to the appropriate file in the /usr/share/zoneinfo directory. For example, to set the timezone to America/New_York, one can use the command sudo ln -sf /usr/share/zoneinfo/America/New_York /etc/localtime. Alternatively, one can use the

timedatectl command to set the timezone without creating the link manually. Reference:
How to Set or Change the Time Zone in Linux | Linuxize
4 Ways to Change the Timezone in Linux - wikiHow

Question: 50

Which of the following fields are available in both the global /etc/crontab file as well as in user-specific crontab files? (Select TWO correct answers)

A. Year
B. Minute
C. Username
D. Command

Answer: B, D

Explanation:
The crontab file format consists of six fields: minute, hour, day of month, month, day of week, and command. The user-specific crontab files have the same format as the global /etc/crontab file, except that they do not have the username field. The username field is only present in the system-wide crontab files and specifies which user will run the cron job. The year field is not a valid crontab field and is not supported by cron. Reference: Scheduling Cron Jobs with Crontab | Linuxize Crontab Explained in Linux [With Examples]

Question: 51

Which command can be used to delete a group from a Linux system?

A. groupdel
B. groupmod
C. groups
D. groupedit

Answer: A

Explanation:

The groupdel command is used to delete a group from a Linux system. It removes the group name from the /etc/group and /etc/gshadow files, but not the group's configuration files, entries, or account files. The groupdel command requires root or sudo privileges and does not accept any

options except for one for chroot. The groupdel command does not print any output on success, but it will display an error message if the group does not exist or if it is the primary group of an existing user. The groupdel command is part of the shadow-utils package, which provides tools for managing user and group accounts. The groupdel command is also

compatible with the Linux Standard Base (LSB) specification, which defines a common set of commands and utilities for Linux distributions. Reference: 1234

Question: 52

What is the purpose of the iconv command?

A. It converts bitmap images from one format to another such as PNG to JPEG.
B. It verifies that the root directory tree complies to all conventions from the Filesystem Hierarchy Standard (FHS).
C. It displays additional meta information from icon files ending in .ico.
D. It changes the mode of an inode in the ext4 file system.
E. It converts files from one character encoding to another.

Answer: E

Explanation:

The iconv command is used to convert the encoding of a file from one character set to another. A character set is a collection of characters that are assigned numerical values called code points. Different character sets may use different numbers of bytes to represent each character, and may have different mappings of code points to characters. For example, ASCII is a single-byte character set that encodes 128 characters, while UTF-8 is a variable-length character set that can encode over a million characters. The iconv command can convert between many different character sets, such as ASCII, UTF-8, ISO-8859-1, etc. The basic syntax for using the command is as follows: iconv [options] -f from-encoding -t to-encoding input-file > output-file
The -f option specifies the encoding of the input file, and the -t option specifies the encoding of the output file. The input file is read from standard input, and the output file is written to standard output, unless specified otherwise. The iconv command can also list all the supported character sets with the -l option1234. Reference: How To Use the iconv Command on Linux - How-To Geek iconv command in Linux with Examples - GeeksforGeeks iconv - convert file encoding from one character set to another | Linux ... Using iconv to change character encodings - FileFormat.Info

Question: 53

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:

The /etc/cron.allow and /etc/cron.deny files are used to control access to the crontab command and cron jobs for individual users. If neither of these files exists, then depending on site-dependent configuration parameters, only the superuser (root user) will be allowed to use this command, or all users will be able to use this command1. The default behavior of most Linux distributions is to allow all users to use the crontab command and have user specific crontabs if neither /etc/cron.allow nor /etc/cron.deny exists23. Therefore, option B is the correct answer. The other options are not true because:

Option A is false because it contradicts the default behavior of most Linux distributions.

Option C is false because the cron daemon will not refuse to start or report missing files in the system's logfile if neither /etc/cron.allow nor /etc/cron.deny exists. The cron daemon will start normally and use the default configuration parameters1.

Option D is false because the system administrator does not need to approve user specific crontabs explicitly. The user can create, edit, display, or remove their own crontab files without any intervention from the system administrator1.

Reference:

How cron.allow and cron.deny can be used to limit access to crontab for a particular user | The Geek Search

crontab(1) — cron — Debian bullseye — Debian Manpages

Controlling Access to crontab (System Administration Guide: Basic Administration) - Oracle /etc/cron.allow - Linux Bash Shell Scripting Tutorial Wiki - nixCraft

Question: 54

Which of the following commands can remove a user from a group?

A. grouprm
B. groupmod
C. passwd
D. usergroups
E. usermod

Answer: E

Explanation:

The usermod command is a utility for modifying user accounts. One of its options is -G, which allows specifying a list of supplementary groups that the user is a member of. If the user is currently a member of a group that is not listed, the user will be removed from that group. For example, to remove the user alice from the group sales, one can use the command sudo usermod -G admin alice,

assuming that alice is only a member of admin and sales groups. Alternatively, one can use the gpasswd command with the --delete option to remove a user from a specific group without affecting other groups. For example, to remove the user alice from the group sales, one can use the command sudo gpasswd --delete alice sales. The other commands in the options are not used for removing a user from a group. The grouprm command does not exist. The groupmod command is used for modifying group attributes, not membership. The passwd command is used for changing user passwords, not groups. The usergroups command is used for displaying the groups that a user belongs to, not modifying them. Reference:

usermod(8) - Linux man page

gpasswd(1) - Linux man page

How to Remove User From Group in Linux [Quick Tip]

Question: 55

Where are user specific crontabs stored?

A. In the database file /etc/crontab.db which is shared by all users.
B. As individual per-user files within /var/spool/cron.
C. As individual per-user files in /etc/cron.user.d.
D. In the .crontab file in the user's home directory.
E. In the file /var/cron/user-crontab which is shared by all users.

Answer: B

Explanation:
The user-specific crontab files are stored in the /var/spool/cron/crontabs directory, where each file is named after the username of the owner. These files are not meant to be edited directly, but rather through the crontab command. The other options are either incorrect or non-existent locations for user crontab files. Reference:
Where is the user crontab stored?
crontab running as a specific user
Location of the crontab file
Where is the User Crontab Stored?

Question: 56

Which file contains the date of the last change of a user's password?

A. /etc/gshadow
B. /etc/passwd
C. /etc/pwdlog
D. /etc/shadow
E. /var/log/shadow

Answer: D

Explanation:

The /etc/shadow file contains the encrypted passwords and other information for each user account on a Linux system. The third field in each line of this file is the date of the last password change, expressed as the number of days since Jan 1, 1970. This information is used by the system to determine when a user must change their password, based on the password aging policy. The /etc/shadow file can be viewed and modified by the root user or by using the chage command123. The other files listed in the options do not store the date of the last password change. The /etc/gshadow file contains the encrypted passwords for group accounts4. The /etc/passwd file contains the basic information for each user account, such as the user name, user ID, group ID, home directory, login shell, etc., but not the password5. The /etc/pwdlog file does not exist by default on most Linux systems, and it is not related to the password change date. The /var/log/shadow file also does not exist by default on most Linux systems, and it is not related to the password change date. Reference: https://www.redhat.com/sysadmin/password-changes-chage-command

https://www.golinuxcloud.com/check-last-password-change-expiration-linux/

Question: 57

Which environment variable should be set in order to change the time zone for the commands run from within the environment variable's scope? (Specify the variable name only.)

Answer: TZ

Explanation:

The TZ environment variable is used to change the time zone for the commands run from within the environment variable's scope. It specifies the name of a time zone as defined in the /usr/share/zoneinfo directory or a custom time zone in the POSIX format12. The TZ variable can be set either globally in a shell profile file or locally in a shell session. For example, to set the time zone to America/New_York for the current shell session, one can use the following command: export TZ=America/New_York

To verify the change, one can use the date command to display the current date and time according to the TZ variable. The TZ variable can also be used to run a single command with a different time zone without affecting the system's time zone. For example, to run the date command with the Asia/Tokyo time zone, one can use the following syntax: TZ=Asia/Tokyo date

The TZ variable is useful for testing how applications behave in different time zones or for displaying the time in different locations34. Reference:

How to Set or Change the Time Zone in Linux | Linuxize

Linux / UNIX: TZ Environment Variable - nixCraft

Get Current System Time Zone in Linux | Baeldung on Linux

Setting the TZ Environment Variable on Linux | InterSystems Developer

Question: 58

Each entry in a crontab must end with what character?

A. Tab
B. Space
C. Backslash
D. Newline

Answer: D

Explanation:

Each entry in a crontab file consists of six fields, specifying in the following order: minute, hour, day, month, weekday, and command1. Any of these fields can be set to an asterisk (*), which stands for "first through last." So, for example, to run a job every hour, put * in the hour field1. Each entry in a crontab file must end with a newline character (\n), which indicates the end of a line2. A newline character is created by pressing the Enter key on the keyboard. The other options are not valid characters for ending a crontab entry. A tab or a space is used to separate each field, and a backslash is used to escape special characters or continue a long command to the next line2. Reference: How to Use the Cron Job Format to Schedule Task in Linux

Syntax of crontab File Entries - Oracle.

Question: 59

To prevent a specific user from scheduling tasks with at, what should the administrator do?

A. Add the specific user to /etc/at.allow file.
B. Add the specific user to [deny] section in the /etc/atd.conf file.
C. Add the specific user to /etc/at.deny file.
D. Add the specific user to nojobs group.
E. Run the following: atd --deny [user].

Answer: C

Explanation:
The /etc/at.deny file is a file that contains a list of users who are not allowed to use the at command to schedule jobs. If the file exists, any user who is not in the /etc/at.allow file and is in the /etc/at.deny file will be denied access to the at command. To prevent a specific user from scheduling tasks with at, the administrator can simply add the user's name to the /etc/at.deny file. For example, to prevent the user bob from using the at command, the administrator can use the following command:
echo "bob" | sudo tee -a /etc/at.deny
The other options are not correct. The /etc/at.allow file is a file that contains a list of users who are allowed to use the at command. Adding a user to this file will not prevent them from scheduling tasks with at. The /etc/atd.conf file is a configuration file for the at daemon, which does not have a [deny] section. Adding a user to this file will not affect their access to the at command. The nojobs

group is not a predefined group in Linux, and adding a user to this group will not prevent them from scheduling tasks with at. The atd command does not have a --deny option, and running this command will not prevent a user from scheduling tasks with at. Reference:
at Command in Linux with Examples - GeeksforGeeks
How to Use the Linux at Command {9 Examples} - phoenixNAP
at(1) - Linux man page

Question: 60

Which of the following crontab entries will execute myscript at 30 minutes past every hour on Sundays?

A. 0 * * * 30 myscript
B. 30 * * * 6 myscript
C. 30 0 * * 0 myscript
D. 30 0-23 * * 0 myscript
E. 0 0-23 * * 30 myscript

Answer: D

Explanation:

The correct crontab entry for executing myscript at 30 minutes past every hour on Sundays is D. 30 023 * * 0 myscript. This is because the crontab format consists of six fields: minute, hour, day of month, month, day of week, and command. The values for each field can be:

A single number, such as 5 or 10.

A range of numbers, such as 1-5 or 10-15.

A list of numbers separated by commas, such as 1,3,5 or 10,12,14.

An asterisk (*), which means all possible values for that field.

A step value, which means every nth value for that field, such as */5 or 10-20/2.

The day of week field can be either a number from 0 to 6, where 0 and 7 are Sunday, or a three-letter abbreviation, such as SUN or MON. The month field can be either a number from 1 to 12, or a three- letter abbreviation, such as JAN or FEB.

In this case, the crontab entry D. 30 0-23 * * 0 myscript means:

30: Execute the command at the 30th minute of every hour.

0-23: Execute the command for every hour from 0 (midnight) to 23 (11 PM).

*: Execute the command for every day of the month, regardless of the month.

*: Execute the command for every month, regardless of the year.

0: Execute the command only on Sundays.

The other options are either incorrect or do not match the requirement. For example, option A. 0 * * * 30 myscript means:

0: Execute the command at the 0th minute of every hour.

*: Execute the command for every hour of the day.

*: Execute the command for every day of the month, regardless of the month.

*: Execute the command for every month, regardless of the year.

30: Execute the command only on the 30th day of the week, which is invalid.

Reference:

Crontab Explained in Linux [With Examples]

'crontab' in Linux with Examples - GeeksforGeeks

Crontab Syntax on Linux + Useful Examples - Hostinger

Question: 61

Which of the following files assigns a user to its primary group?

A. /etc/pgroup B. /etc/shadow C. /etc/group D. /etc/passwd E. /etc/gshadow

Answer: D

Explanation:

The /etc/passwd file assigns a user to its primary group by specifying the group ID (GID) of the primary group in the fourth field of each line. The /etc/passwd file contains the basic information for each user account on a Linux system, such as the user name, user ID (UID), group ID (GID), home directory, login shell, etc. The format of each line is: username:password:UID:GID:comment:home:shell

For example, the following line assigns the user bob to the primary group bob, which has the GID of 1001:

bob:x:1001:1001::/home/bob:/bin/sh

The /etc/passwd file can be viewed and modified by the root user or by using the useradd, usermod, or userdel

commands123. The other files listed in the options do not assign a user to its primary group. The /etc/pgroup file does not exist by default on most Linux systems, and it is not related to the primary group. The /etc/shadow file contains the encrypted passwords and other information for each user account, but not the primary group4. The /etc/group file contains the information for each group on the system, such as the group name, group password, group ID, and group members, but not the primary group of each user5. The /etc/gshadow file contains the encrypted passwords for group accounts.
Reference: 12345

Question: 62

Which of the following commands should be added to /etc/bash_profile in order to change the language of messages for an internationalized program to Portuguese (pt)?

A. export LANGUAGE="pt"
B. export MESSAGE="pt"
C. export UI_MESSAGES="pt"
D. export LC_MESSAGES="pt"
E. export ALL_MESSAGES="pt"

Answer: D

Explanation:

The LC_MESSAGES environment variable specifies the language to use in diagnostic messages for an internationalized program. It can be set to any value supported by the installation, such as pt for Portuguese, en for English, fr for French, etc. The LC_MESSAGES variable can be set either globally in a shell profile file, such as /etc/bash_profile, or locally in a shell session. For example, to set the language of messages to Portuguese for the current shell session, one can use the following command:

export LC_MESSAGES=pt

To verify the change, one can run an internationalized program, such as man, and see the output in Portuguese. The LC_MESSAGES variable can also be used to run a single command with a different language without affecting the system's language. For example, to run the man command with the Spanish language, one can use the following syntax:

LC_MESSAGES=es man

The LC_MESSAGES variable is useful for testing how programs behave in different languages or for displaying messages in different languages1234. Reference:
Locale Environment Variables in Linux | Baeldung on Linux
Linux / UNIX: TZ Environment Variable - nixCraft
Changing your locale on Linux and UNIX systems - IBM
Selecting message language in gcc and g++ - Stack Overflow

Question: 63

In which file, if present, must all users be listed that are allowed to use the cron scheduling system? (Specify the full name of the file, including path.)

Answer:
/etc/cron.allow

Explanation:

The /etc/cron.allow file is a file that contains a list of users who are allowed to use the cron scheduling system. The cron scheduling system is a way of running commands or scripts at specified times or intervals. Users can create their own cron jobs by using the crontab command, which edits a file called crontab that stores the user's scheduled tasks. However, not all users may have access to the crontab command or the cron system. The access is controlled by two files: /etc/cron.allow and /etc/cron.deny. If the /etc/cron.allow file exists, then only the users listed in this file can use the crontab command and the cron system. The file should have one user name per line. If the /etc/cron.allow file does not exist, then the /etc/cron.deny file is checked. If this file exists, then the users listed in this file are denied access to the crontab command and the cron system. If neither file exists, then the access depends on the configuration of the cron daemon, which is the program that runs the cron jobs. By default, only the root user can use the cron system if no files exist. The root

user can always use the cron system regardless of the existence or content of these files. To create or edit the /etc/cron.allow file, the root user needs to use a text editor such as vi, nano, or emacs. For example, to allow the users alice and bob to use the cron system, the root user can use the following command:

sudo vi /etc/cron.allow

And then add the following lines to the file:

alice bob

And then save and exit the file. Reference:

How cron.allow and cron.deny can be used to limit access to crontab for ...
/etc/cron.allow - Linux Bash Shell Scripting Tutorial Wiki
Linux / UNIX Restrict at / cron Usage To Authorized Users

Question: 64

Which commands can be used to change a user's account aging information? (Choose THREE correct answers.)

A. usermod
B. passwd
C. chattr
D. chage
E. chsh

Answer: A, B, D

Explanation:

The usermod, passwd, and chage commands can be used to change a user's account aging information. These commands can modify the password expiry date, the last password change date, the minimum and maximum number of days between password changes, the number of days of warning before password expiration, and the number of days of inactivity after password expiration. The usermod command is mainly used for modifying a user account, but it also has options for changing the password expiry and aging information, such as -e, -f, -p, and -L1. The passwd command is mainly used for changing the user password, but it also has options for changing the password expiry and aging information, such as -e, -i, -n, -w, and -x2. The chage command is specifically used for changing the user password expiry and aging information, and it

has options such as -d, -E, -I, -m, -M, and -W3.
The other options, chattr and chsh, are not related to changing the user's account aging information. The chattr command is used to change the file attributes on a Linux file system4. The chsh command is used to change the user's login shell5.
Reference:
usermod(8) — Linux manual page passwd(1) — Linux manual page chage(1) — Linux manual page chattr(1) — Linux manual page chsh(1) — Linux manual page

Question: 65

Which command is used to add an empty group to the system? (Specify ONLY the command without any path or parameters.)

Answer: groupadd, /usr/sbin/groupadd

Explanation:

The groupadd command is used to add an empty group to the system. It takes the name of the group as an argument and creates an entry for it in the /etc/group file. The groupadd command also assigns a unique group ID (GID) to the new group. The groupadd command can take various options to specify the GID, the password, and other attributes of the new group. For example, groupadd -g 1000 mygroup will create a new group named mygroup with a GID of 1000.
Reference: Linux Groups - javatpoint groupadd(8) - Linux manual page
How to Add and Delete User Groups on Linux

Question: 66

What is NOT contained in the locale setting of the operating system?

A. currency symbol
B. language
C. timezone
D. thousands separator

Answer: C

Explanation:

The locale setting of the operating system is a set of environmental variables that defines the language, country, and character encoding settings (or any other special variant preferences) for the applications and shell session on a Linux system12. The locale setting usually consists of at least a language code and a country/region code, such as en_US for English (United States) or fr_FR for French (France). The locale setting also affects things such as the currency symbol, the thousands separator, the decimal point, the date and time format, the collation order, the paper size, the telephone number format, and many other values formatted in accordance with the language or region/country12. However, the timezone is not contained in the locale setting of the operating system. The timezone is a separate setting that determines the local time of the system based on the offset from the Coordinated Universal Time (UTC) and the daylight saving time (DST) rules. The timezone can be different from the country/region code in the locale setting, for example, a user can have a locale setting of en_US but a timezone of Asia/Kolkata. The timezone can be viewed and modified by using the date, timedatectl, or

tzselect commands . Reference: 12

Question: 67

What is true about the file /etc/localtime?

A. It is a plain text file containing a string such as Europe/Berlin
B. It is created and maintained by the NTP service based on the location of the system's IP address.
C. It is a symlink to /sys/device/clock/ltime and always contains the current local time.
D. After changing this file, newtzconfig has to be run to make the changes effective.
E. It is either a symlink to or a copy of a timezone information file such as /usr/share/zoneinfo/Europe/Berlin.

Answer: E

Explanation:
The /etc/localtime file is used to configure the system-wide timezone of the local system that is used by applications for presentation to the user. It should be either a symlink to or a copy of a timezone information file that contains the binary data for the configured timezone. The timezone information files are located under /usr/share/zoneinfo/ and are named after the geographic regions and cities, such as Europe/Berlin or Etc/UTC. The timezone identifier is extracted from the symlink target name of /etc/localtime, so it is recommended to use a symlink rather than a copy. The timezone can be changed by using the timedatectl command or by creating a new symlink to the desired timezone file123. Reference: How to Set or Change the Time Zone in Linux | Linuxize localtime(5) - Linux manual page - man7.org localtime(5) — Arch manual pages

Question: 68

What is true regarding the command userdel --force --remove bob? (Choose TWO correct answers.)

A. The user bob is removed from the system's user database.
B. The user bob's home directory is removed.
C. The locate database is updated to drop files owned by bob.
D. All files owned by bob are remove from all mounted filesystems.
E. In case bob was the last member of a group, that group is deleted.

Answer: A, B

Explanation:

The command userdel --force --remove bob is used to delete the user account named bob and all its associated files. The --force option forces the removal of the user account, even if the user is still logged in. The --remove option forces userdel to remove the user's home directory and mail spool,

even if another user uses the same home directory or if the mail spool is not owned by the specified user12. Therefore, options A and B are true regarding this command.

The other options are not true because:

Option C is false because the locate database is not updated by the userdel command. The locate database is updated by the updatedb command, which is usually run by cron as a scheduled job3. Option D is false because the userdel command does not remove all files owned by bob from all mounted filesystems. The userdel command only removes the user's home directory and mail spool, and it does not search for and delete the user files located in other file systems. You have to search for and delete the files manually1.

Option E is false because the userdel command does not delete the group with the same name as the user, unless the USERGROUPS_ENAB parameter is set to yes in the /etc/login.defs file and the group has no other members14.

Reference:

How to Delete/Remove Users in Linux (userdel Command) | Linuxize

userdel(8) — Linux manual page

updatedb(8) — Linux manual page

Understanding the /etc/login.defs File | Linuxize

Question: 69

Which of the following fields can be found in the /etc/group file? (Choose THREE correct answers.)

A. The list of users that belong to the group.
B. The home directory of the group.
C. The name of the group.
D. The description of the group.
E. The password of the group.

Answer: A, C, E

Explanation:

The /etc/group file is a text file that defines the groups and their members on the system. Each line in the file represents a single group, with the following format: group_name:password:GID:user_list

The fields are:

group_name: the name of the group

password: the (encrypted) group password, or empty for no password

GID: the numerical group ID

user_list: a comma-separated list of users who belong to the group

Therefore, the fields that can be found in the /etc/group file are the name of the group, the password of the group, and the list of users that belong to the group. The home directory and the description of the group are not part of the /etc/group file format. Reference:

group(5) - Linux manual page

/etc/group file format | Linux#

Question: 70

On a system using shadowed passwords, the most correct permissions for /etc/passwd are _and the most correct permissions for /etc/shadow are.

A. -rw-r , -r
B. -rw-r--r--, -r--r--r--

C. -rw-r r , -r
D. -rw-r--rw-, -r—r--
E. -rw , -r

Answer: C

Explanation:

The /etc/passwd file stores local accounts of the system. It is a readable text file and uses colons (:) to separate the fields. This file helps with converting user IDs to names (and back). It is fine that all users can read this file, but they should not be able to change fields. Therefore, the most correct permissions for /etc/passwd are -rw-r–r–, which means that only the owner (root) can write to the file, and everyone can read it. The /etc/shadow file contains information about the system's users' passwords. It is owned by user root and group shadow, and has 640 permissions. The password is stored as a long string of characters, which is a combination of the hashing algorithm, optional salt applied, and the hashed password itself. Other users are not allowed to read the file directly, to prevent them from gathering hashed passwords of others. Therefore, the most correct permissions for /etc/shadow are -r—, which means that only the owner (root) can read the file, and no one else can read or write to it. Reference:

Question: 71

A French user has installed the French language pack, but currencies are still being displayed with a leading '$' sign in his spreadsheets. What must be done to fix this?

A. Alter the locale.
B. Set the timezone correctly.
C. Edit /etc/currency.
D. Reinstall the French language pack.

Answer: A

Explanation:

The locale is a set of environmental variables that defines the language, country, and character encoding settings for the applications and shell session on a Linux system. The locale affects things

such as the time/date format, the first day of the week, numbers, currency and many other values formatted in accordance with the language or region/country you set on a Linux system1. The currency sign (¤) is a character used to denote an unspecified currency2. To display the correct currency symbol for a specific region, the locale must be set accordingly. For example, to display the euro symbol (€) for France, the locale can be set to fr_FR.UTF-81. Setting the timezone correctly, editing /etc/currency, or reinstalling the French language pack will not affect the currency symbol displayed in the spreadsheets. Reference: 1: How to Change or Set System Locales in Linux - Tecmint 2: decimal.ToString ("C") produces ¤ currency symbol on Linux - Stack Overflow

Question: 72

Which of the following can the chage command NOT change?

A. The number of days since January 1, 1970 after which the user's account will no longer be accessible.
B. The number of days since January 1, 1970 after which the password can change.
C. The number of days since January 1, 1970 since the password was last changed.
D. The maximum number of days during which a password is valid.
E. The number of days of inactivity after a password has expired before the account is locked.

Answer: E

Explanation:

The chage command can change the following parameters related to user password expiry and aging:
The last password change date (-d or --lastday option)
The password expiry date (-E or --expiredate option)
The minimum number of days between password changes (-m or --mindays option)
The maximum number of days during which a password is valid (-M or --maxdays option)
The number of days of warning before password expires (-W or --warndays option)
The chage command cannot change the number of days of inactivity after a password has expired before the account is locked. This parameter is controlled by the -I or --inactive option of the usermod command, which modifies the user account information. The chage command only displays the current value of this parameter, but does not allow changing it. Reference: chage command in Linux with examples - GeeksforGeeks
10 chage command examples in Linux [Cheat Sheet] - GoLinuxCloud
How to Use the Chage Command in Linux – TecAdmin
How to Manage User Password Expiration and Aging in Linux - Tecmint

Question: 73

What command will display the group names and GIDs to which a user belongs? (Provide only the command name with or without path information)

Answer: id, /usr/bin/id

Explanation:

The id command will display the user ID (uid), the primary group ID (gid), and the supplementary groups (groups) of a user.
The output will show the names and the numerical IDs of the groups. For example: id linuxize
The command will show the user ID (uid), the user's primary group (gid), and the user's secondary groups (groups)
uid=1001(linuxize) gid=1001(linuxize) groups=1001(linuxize),27(sudo)
To print only the names instead of the numbers use the -n option.
id -nG linuxize
The command will show only the names of the groups linuxize sudo
The id command is part of the GNU coreutils package and is available on all Linux systems. The full path of the command is /usr/bin/id. Reference: id(1) - Linux manual page

How to List Groups in Linux | Linuxize

Question: 74

Of the ways listed, which is the best method to temporarily suspend a user's ability to interactively login?

A. Use passwd -d username to give the user an empty password.
B. Use chage to expire the user account.
C. Change the user's password.
D. Add the command exit to the user's .login file.

Answer: B

Explanation:
The chage command can be used to change the expiration date of a user account. By setting the expiration date to a past date, the user account will be disabled and the user will not be able to login interactively. This is a temporary method, as the expiration date can be changed back to a future date or removed to re-enable the user account. The other options are either permanent, insecure, or ineffective. Option A is insecure, as it allows anyone to login as the user without a password. Option C is permanent, as it changes the user's password without saving the original one. Option D is ineffective, as it only affects the user's .login file, which is used by the csh and tcsh shells, and not by other shells such as bash or zsh. Therefore, option B is the best method to temporarily suspend a user's ability to interactively login.
Reference: https://linuxconfig.org/disabling-user-logins-to-linux- system
https://askubuntu.com/questions/282806/how-to-enable-or-disable-a-user

Question: 75

What is the conventional purpose of Linux UIDs that are lower than 100?

A. They are reserved for super user accounts.
B. They are reserved for the system admin accounts.
C. They are reserved for system accounts.
D. They are unused, aside from 0, because they are targets of exploits.
E. They are used to match with GIDs in grouping users.

Answer: C

Explanation:
Linux UIDs (user identifiers) are numbers that are used to identify users and groups on a Linux system. Each user and group has a unique UID and GID (group identifier) respectively. The UID 0 is always reserved for the root or superuser account, which has full privileges to access and modify the system. The UIDs lower than 100 (or 1000 on some modern systems) are typically reserved for system accounts, which are used by various services and daemons that run on the system. These accounts are not meant for human users, but for specific purposes such as managing files, processes, network, security, etc. For example, some common system accounts are bin, daemon, mail, sshd, etc. The UIDs higher than 100 (or 1000) are usually allocated for regular user accounts, which have limited privileges and can be created and deleted by the system administrator. The system accounts are defined in the /etc/passwd file, which contains the username, UID, GID, home directory, shell, and other information for each account12345. Reference: 1: Linux User Management - Tecmint 2: What are the well-known UIDs? - Stack Overflow 3: user ID less than 1000 on CentOS 7 - Unix

& Linux Stack Exchange 4: Recommended GID for users group in Linux (100 or 1000)? - Unix & Linux Stack Exchange 5: What is the conventional purpose of Linux UIDs that are lower than 100? - VCE Guide

Question: 76

How is the file format of /etc/crontab different from a normal crontab file? (Select TWO correct answers)

A. The /etc/crontab file can specify a year field.
B. A normal crontab file must be installed with the crontab command.
C. A normal crontab file allows for environment variable substitution.
D. The /etc/crontab file has a user field for commands.

Answer: B, D

Explanation:
The /etc/crontab file is the system-wide crontab file that can be edited only by root. It has a different format from the normal crontab files that can be edited by individual users using the crontab command. The differences are:
The /etc/crontab file can specify a year field as the sixth field in a cron entry. This allows for scheduling jobs that run only in specific years. The normal crontab files do not have a year field and assume the current year for all entries.
The /etc/crontab file has a user field as the seventh field in a cron entry. This allows for running commands as different users from the crontab owner (root). The normal crontab files do not have a user field and run commands as the crontab owner.
The /etc/crontab file does not need to be installed with the crontab command. It is read by the cron daemon automatically. The normal crontab files need to be installed with the crontab command to be recognized by the cron daemon.
The /etc/crontab file and the normal crontab files both allow for environment variable substitution. However, the /etc/crontab file sets some default environment variables such as SHELL, PATH, MAILTO, and HOME, which can be overridden by entries in the file. The normal crontab files inherit the environment variables from the cron daemon, which are usually minimal.

Reference:
crontab(5) - Linux manual page
Linux Crontab Format
How to schedule a task using Linux crontab (/etc/crontab) file
/etc/crontab - Linux Bash Shell Scripting Tutorial Wiki

Question: 77

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 is similar to the at command, except that it executes commands when the system load levels permit; in other words, when the load average drops below 1.5, or the value specified in the invocation of atd1. The at command allows us to schedule jobs using any of two commands: at and batch. While at runs commands at our specified time, batch runs commands when our system's load average is below 0.82. Both commands read commands from standard input or a specified file, and both commands send the output of the commands to the user by mail1. Therefore, the main difference between them is the time of execution: at runs at a fixed time, while batch runs when the system is idle. Reference: 1: Linux At, Batch, Atq, Atrm Command Help and Examples - Computer Hope 2: The "at" Command in Linux | Baeldung on Linux

Question: 78

The correct crontab entry to execute the script chklog three times per month between 3 p.m. and 5 p.m.:

A. * 3,4,5 1 * * chklog
B. 3 3,4,5 1 * * chklog
C. 3 15,16,17 * * * chklog D. 0 15,16,17 1 * * chklog E. * 15,16,17 1 * * chklog

Answer: C

Explanation:

The correct crontab entry to execute the script chklog three times per month between 3 p.m. and 5 p.m. is:
3 15,16,17 * * * chklog

The crontab entry has five fields that specify the time and frequency of the job, followed by the command or script to be executed. The fields are:

Minute: the minute of the hour when the job should run, from 0 to 59
Hour: the hour of the day when the job should run, from 0 to 23 (in 24-hour format)
Day of month: the day of the month when the job should run, from 1 to 31
Month: the month of the year when the job should run, from 1 to 12
Day of week: the day of the week when the job should run, from 0 to 6 (where 0 and 7 are Sunday) The asterisk (*) means any value, and the comma (,) means a list of values. Therefore, the crontab entry above means:
Run the job at the 3rd minute of the hour
Run the job at the 15th, 16th, and 17th hour of the day (which are 3 p.m., 4 p.m., and 5 p.m.)
Run the job on any day of the month
Run the job on any month of the year
Run the job on any day of the week
This will execute the script chklog three times per day, every day of the month, and every month of the year, which is equivalent to three times per month.

The other options are incorrect because:

A . This will run the job at any minute of the hour, but only at the 3rd, 4th, and 5th hour of the day (which are 3 a.m., 4 a.m., and 5 a.m.), and only on the 1st day of the month.

B . This will run the job at the 3rd minute of the hour, but only at the 3rd, 4th, and 5th hour of the day (which are 3 a.m., 4 a.m., and 5 a.m.), and only on the 1st day of the month.

D . This will run the job at the 0th minute of the hour (which is the top of the hour), but only at the 15th, 16th, and 17th hour of the day (which are 3 p.m., 4 p.m., and 5 p.m.), and only on the 1st day of the month.

E . This will run the job at any minute of the hour, but only at the 15th, 16th, and 17th hour of the day (which are 3 p.m., 4 p.m., and 5 p.m.), and only on the 1st day of the month.
Reference:
Crontab Explained in Linux [With Examples]
'crontab' in Linux with Examples - GeeksforGeeks
Linux Crontab Command Help and Examples - Computer Hope
Crontab in Linux with 20 Useful Examples to Schedule Jobs - TecAdmin
Linux crontab tutorial with Examples - Linux Tutorials - Learn Linux ...

Question: 79

Fill in Blanks

The command is used to add a group to the system.

Answer: groupadd, /usr/sbin/groupadd

Explanation:

The groupadd command creates a new group using the options specified on the command line and the default values from the /etc/login.defs file. It adds an entry for the new group to the /etc/group and /etc/gshadow files. Only the root user or a user with sudo privileges can create new groups using this command. The general syntax for the groupadd command is as follows: groupadd [OPTIONS] GROUPNAME
Some of the common options for the groupadd command are:

- g, --gid GID: Specify the numeric group ID for the new group. If not given, the system will assign the next available GID from the range of group IDs specified in the login.defs file.
- r, --system: Create a system group with a GID chosen from the range of system group IDs specified in the login.defs file. System groups are usually used for some special system operation purposes, like creating backups or doing system maintenance.
- f, --force: Suppress the error message if the group already exists and exit successfully. This option is useful for scripts that need to ensure the existence of a group.
- K, --key KEY=VALUE: Override the default values from the /etc/login.defs file. The valid keys are GROUP_MIN_ID, GROUP_MAX_ID, SYS_GROUP_MIN_ID, SYS_GROUP_MAX_ID, and GID_INCREMENT.

Reference: https://www.makeuseof.com/linux-file-ownership-groups-guide/ https://linuxize.com/post/how-to-create-groups-in-linux/ https://linuxhandbook.com/groupadd-command/

Question: 80

Which command will set the local machine's timezone to UTC?

A. cat UTC > /etc/timezone
B. ln -s /usr/share/zoneinfo/UTC /etc/localtime
C. date --timezone=UTC
D. mv /usr/timezone/UTC /etc

Answer: B

Explanation:

The command ln -s /usr/share/zoneinfo/UTC /etc/localtime will create a symbolic link from the file /etc/localtime to the file /usr/share/zoneinfo/UTC, which contains the binary time zone data for the UTC timezone. This will set the system's timezone to UTC, which is the Coordinated Universal Time, the primary time standard by which the world regulates clocks and time1. The /etc/localtime file is used by various system programs and libraries to determine the local time according to the configured timezone2. The /usr/share/zoneinfo directory contains the time zone information files for different regions and cities around the world3. The other commands are either invalid or will not change the system's timezone permanently. The command cat UTC > /etc/timezone will overwrite the /etc/timezone file with the string "UTC", which is not a valid timezone identifier. The /etc/timezone file is a plain text file that contains the name of the timezone, such as "America/New_York" or "Europe/Paris"4. The command date --timezone=UTC will display the current date and time in UTC, but will not change the system's timezone setting. The command mv /usr/timezone/UTC /etc will move the file /usr/timezone/UTC to the /etc directory, but this file does not exist by default and has no effect on the system's timezone configuration. Reference: 1: Coordinated Universal Time - Wikipedia 2: localtime(5) - Linux manual page 3: tz database - Wikipedia 4: How to Change or Set System Locales in Linux - Tecmint : date(1) - Linux manual page : How do I change my timezone to UTC/GMT? - Ask Ubuntu

Question: 81

Which command should be added to /etc/bash_profile to change the language of messages from an internationalised program to Portuguese (pt)? (Select TWO correct answers)

A. export LANGUAGE="pt"
B. export MESSAGE="pt"
C. export LANG="pt"
D. export LC_MESSAGES="pt"
E. export ALL_MESSAGES="pt"

Answer: C, D

Explanation:

The commands that should be added to /etc/bash_profile to change the language of messages from an internationalised program to Portuguese (pt) are:

export LANG="pt"
export LC_MESSAGES="pt"

The LANG and LC_MESSAGES environment variables are used to control the language of messages from an internationalised program. The LANG variable sets the default locale for all categories, such as collation, currency, date and time formats, etc. The LC_MESSAGES variable sets the locale for the language of messages, overriding the LANG variable for this category. Therefore, to change the language of messages to Portuguese, both variables should be set to "pt" in /etc/bash_profile, which is a script that is executed when a user logs in. This will affect the current user and any subsequent login sessions.

Reference:

Locale Environment Variables in Linux – Baeldung on Linux
Environment Variables - The Open Group

[LPI Linux Essentials - 1.4 Localization and Internationalization]

Question: 82

A user was not given permission to use the CRON scheduling system. What file needs to be modified to provide that access? (Please specify the full path to the file)

Answer: /etc/cron.allow

Explanation:

The /etc/cron.d/cron.allow file is a text file that contains the names of the users who are allowed to use the crontab command to create and manage their own cron jobs12. If this file exists, only the users listed in it can use the crontab command, and all other users are denied access12. If this file does not exist, the /etc/cron.d/cron.deny file is checked to see which users are not allowed to use the crontab command12. If neither file exists, only the root user can use the crontab command12. To modify the /etc/cron.d/cron.allow file, the root user can use any text editor to add or remove the names of the users who will be allowed to use the crontab command1234. For example, to allow the user frank to use the crontab command, the root user can append the name frank to the /etc/cron.d/cron.allow file1234. The root user must always be included in this file, otherwise the superuser access to the crontab command will be denied4.

Reference: 1: Linux At, Batch, Atq, Atrm Command Help and Examples - Computer Hope 2: 107.2 Lesson 1 - Linux Professional Institute Certification Programs 3: /etc/cron.allow - Linux Bash Shell Scripting Tutorial Wiki - nixCraft 4: Controlling Access to crontab (System Administration Guide ... - Oracle

Question: 83

Which commands can you use to change a user's account aging information? (Choose THREE correct answers.)

A. usermod
B. passwd
C. chattr
D. chage
E. chsh

Answer: A, B, D

Explanation:

The commands that can be used to change a user's account aging information are: usermod: this command can modify various user account properties, including the password expiration date, the account expiration date, the minimum and maximum password age, the password warning period, and the password inactivity period. To use this command, you need to specify the option and the value for the property you want to change, followed by the username. For example, to set the password expiration date for the user test to February 11, 2022, you can run: usermod -e 2022-02-11 test

To view the current account aging information for a user, you can use the -l option with the usermod command. For example, to view the information for the user test, you can run:

usermod -l test

passwd: this command can change the password of a user account, as well as some password aging options. To use this command, you need to specify the username and the option for the property you want to change. For example, to change the password of the user test, you can run: passwd test

To set the maximum password age for the user test to 90 days, you can run: passwd -x 90 test

To view the current password aging information for a user, you can use the -S option with the passwd command. For example, to view the information for the user test, you can run: passwd -S test

chage: this command can change the user password expiry and aging information, such as the password expiration date, the account expiration date, the minimum and maximum password age, the password warning period, and the password inactivity period. To use this command, you need to specify the option and the value for the property you want to change, followed by the username. For example, to set the account expiration date for the user test to February 11, 2022, you can run: chage -E 2022-02-11 test

To view the current account aging information for a user, you can use the -l option with the chage command. For example, to view the information for the user test, you can run: chage -l test

The other options are incorrect because:

chattr: this command can change the file attributes on a Linux file system, such as making a file immutable, append-only, or undeletable. It has nothing to do with user account aging information. chsh: this command can change the login shell of a user account, such as bash, zsh, or ksh. It has nothing to do with user account aging information.

Reference:

How to Manage User Password Expiration and Aging in Linux - Tecmint

Use the Chage Command in Linux

How to set user password expirations on Linux | Enable Sysadmin

How to change password and account expiry options on Linux using chage - Linux Tutorials - Learn Linux Configuration

3 ways to change user password expiration date in Linux - howtouselinux

Question: 84

Why is /etc/shadow not world readable if the passwords are stored in an encrypted fashion?

A. The encrypted passwords are still subject to brute force attacks.
B. This is just for historical reasons.
C. There is other information in the file that needs to be kept secret.
D. The passwords can be decrypted by anyone with root access.

Answer: A

Explanation:

The /etc/shadow file is not world readable because the encrypted passwords stored in it are still vulnerable to ofline brute force attacks. A brute force attack is a method of trying every possible password until finding the correct one. With modern hardware and software, millions of passwords can be tried per second. If the /etc/shadow file was world readable, anyone who logged in to the system, even as a guest, could copy the file and attempt to crack the passwords without leaving any trace. By making the file readable only by the root user, the system prevents unauthorized access to the password hashes and reduces the risk of password compromise. The other options are incorrect because they do not explain the reason for the file permissions. Option B is false, as the /etc/shadow file was created to address the security issues of the /etc/passwd file, which used to store the passwords in

a world readable file. Option C is partially true, as the /etc/shadow file does contain other information related to password expiration and account locking, but this is not the main reason for making the file not world readable. Option D is irrelevant, as the passwords cannot be decrypted by anyone, even with root access, as the encryption is one-way and irreversible. Reference: https://www.computernetworkingnotes.com/linux-tutorials/etc-shadow-file-in-linux-explained-with-examples.html
https://kerneltalks.com/user-management/understanding-etc-shadow-file/

Question: 85

Of the ways listed, which is the best way to temporarily suspend a single user's ability to interactively login?

A. Add the user name to /etc/nologin.
B. Change the user's password.
C. Change the user name in /etc/passwd.
D. Use change to expire the user account.
E. Place the command logout in the user's profile.

Answer: D

Explanation:
The best way to temporarily suspend a single user's ability to interactively login is to use the chage command to expire the user account. The chage command can modify the expiration date of a user account, which is stored in the /etc/shadow file. By setting the expiration date to a past date, the user account will be locked and the user will not be able to login. This method is temporary because the expiration date can be changed again to a future date or removed to unlock the user account. For example, to expire the user account linuxconfig, we can use the following command:
chage -E 0 linuxconfig
This will set the expiration date to January 1, 1970, which is the epoch date. To check the expiration date of a user account, we can use the -l option:
chage -l linuxconfig
Last password change : Aug 24, 2021
Password expires : never
Password inactive : never
Account expires : Jan 01, 1970
Minimum number of days between password change 0
Maximum number of days between password change 99999
Number of days of warning before password expires 7
To remove the expiration date of a user account, we can use the -E option with an empty argument: # chage -E "" linuxconfig
The other options are either invalid or not recommended. Adding the user name to /etc/nologin will not work, because /etc/nologin is a file that contains a message to be displayed to users who try to login when the system is down for maintenance. Changing the user's password is not a good idea, because it will affect the user's authentication and may cause security issues. Changing the user name in /etc/passwd will also affect the user's authentication and may cause inconsistencies with other files and services. Placing the command logout in the user's profile will not prevent the user from logging in, but only log them out immediately after login, which is not very elegant or secure.
Reference: 1: How to disable user login with Linux nologin - LinuxConfig.org 2: Disable a user's login without disabling the account - Unix & Linux Stack Exchange 3: How to Block or Disable Normal User Logins in Linux? - GeeksforGeeks 4: How

to Disable User Logins on Linux | Baeldung on Linux 5: How to Disable a User in Linux - Linux Nightly 6: How to deactivate or disable a user account in Ubuntu 20.04 LTS - Vitux 7: chage(1) - Linux manual page

Question: 86

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: AC

Explanation:

A . It is true that every user may have their own crontab. This allows individual users to schedule tasks to be run at specific times without requiring administrative privileges1. C. The cron daemon does indeed reload crontab files automatically when necessary. This means that after editing a crontab file, there is no need to manually restart the cron service for the changes to take effect1. Reference:

107.2 Lesson 1 - Linux Professional Institute Certification Programs, which covers the automation of system administration tasks by scheduling jobs with cron.

The LPIC2 Exam Prep, which provides additional information on the LPIC-2 objectives and topics related to crontab and cron daemon.

Question: 87

Which crontab entry could be used to set the system time at regular intervals?

A. 1 0 * * * date $d $t $24
B. 1 0 * * * ntpdate ntp1.digex.net
C. 1 0 * * * date ntp1.digex.net
D. 1 0 * * * runcron date ntp1.digex.net
E. 1 0 * * * settime $d $t $24

Answer: B

Explanation:

The crontab entry that could be used to set the system time at regular intervals is the one that uses the ntpdate command to synchronize the system clock with a Network Time Protocol (NTP) server. The ntpdate command takes one or more NTP server names or IP addresses as arguments and adjusts the system clock accordingly12. The crontab entry B specifies that the ntpdate command should be executed at the first minute of the zeroth hour (i.e., 00:01) of every day of every month of every weekday, using the NTP server ntp1.digex.net34. This will ensure that the system time is updated

daily with a reliable source.
The other crontab entries are either invalid or ineffective for setting the system time at regular intervals. The date command can be used to display or set the system date and time, but it requires a specific format for the argument, not an NTP server name5. The runcron and settime commands are not standard Linux commands and their functionality is unknown. The $d, $t, and $24 variables are also undefined and meaningless in this context.
Reference: 1: Linux At, Batch, Atq, Atrm Command Help and Examples - Computer Hope 2: How to set a cron job to run at a exact time? - Stack Overflow 3: 107.2 Lesson 1 - Linux Professional Institute Certification Programs 4: How to setup a crontab to execute at specific time - Stack Overflow 5: Writing a specific format of time in a text file every minute using ... - Ask Ubuntu

Question: 88

Which of the following commands can be used to convert text files in one character encoding to another character encoding?

A. cat
B. convert
C. dd
D. iconv
E. utf2utf

Answer: D

Explanation:
The command that can be used to convert text files in one character encoding to another character encoding is:
iconv: this command can convert text files from one form of encoding to another, such as UTF-8, ISO- 8859-1, ASCII, etc. To use this command, you need to specify the input encoding, the output encoding, and the file name. For example, to convert a file named input.txt from ISO-8859-1 to UTF- 8, you can run:

iconv -f ISO-8859-1 -t UTF-8 input.txt
The output will be printed to the standard output, which can be redirected to another file or piped to another command. You can also use the -o option to specify the output file name. For example, to convert the same file and save the output to output.txt, you can run: iconv -f ISO-8859-1 -t UTF-8 -o output.txt input.txt
To list all the supported encodings, you can use the -l option. For example, to see all the encodings that start with UTF, you can run: iconv -l | grep UTF
The iconv command is part of the GNU libc package and is available on most Linux systems. The full path of the command is /usr/bin/iconv.
The other options are incorrect because:
cat: this command can concatenate and print files to the standard output, but it does not perform any encoding conversion. It can be used to display the contents of a text file, but it will not change the encoding of the file.
convert: this command can convert image files from one format to another, such as PNG, JPEG, GIF, etc. It is part of the ImageMagick suite of tools and is not related to text encoding conversion. dd: this command can copy and convert data from one source to another, such as files, devices, or pipes. It can perform some conversions, such as changing the case of letters, swapping bytes, or converting between ASCII and EBCDIC, but it does not support common text encodings such as UTF-8 or ISO-8859-1.

utf2utf: this is not a valid command on Linux. There is no such tool that can convert between different UTF encodings.

Reference:

How to Convert Files to UTF-8 Encoding in Linux - Tecmint

Best way to convert text files between character sets? - Stack Overflow

how to change encoding of a text file without openning the file in shell program - Stack Overflow HowTo: Check and Change File Encoding In Linux - ShellHacks

How to change character encoding of a text file on Linux - Xmodulo

Topic 4, Essential System Services

Question: 89

Which option in the /etc/ntp.conf file specifies an external NTP source to be queried for time information? (Specify ONLY the option without any values or parameters.)

Answer: server

Explanation:

The server option is used to configure a persistent association with a remote server or peer. It takes an argument that is either a host name or a numeric IP address of the NTP server. The ntpd daemon will periodically send NTP packets to the specified server and adjust the local clock according to the received responses. Multiple server options can be used to specify more than one NTP source. For example, the following lines in the /etc/ntp.conf file configure four external NTP sources: server 0.asia.pool.ntp.org server 0.oceania.pool.ntp.org server 0.europe.pool.ntp.org server 0.north-america.pool.ntp.org

Reference:

https://docs.ntpsec.org/latest/ntp_conf.html

https://vceguide.com/which-option-in-the-etc-ntp-conf-file-specifies-an-external-ntp-source-to-be- queried-for-time-information-2/

https://vceguide.com/which-option-in-the-etcntp-conf-file-specifies-an-external-ntp-source-to-be- queried-for-time-information/

Question: 90

After configuring printing on a Linux server, the administrator sends a test file to one of the printers and it fails to print. What command can be used to display the status of the printer's queue? (Specify ONLY the command without any path or parameters.)

Answer: lpq, /usr/bin/lpq, lpstat, /usr/bin/lpstat

Explanation:

The command lpq can be used to display the status of the printer's queue on a Linux server. The lpq command is part of the cups-bsd package, which provides the Berkeley commands for CUPS (Common UNIX Printing System), the standard printing system for Linux. The lpq command shows the status of a specified printer or the default printer if none is specified. It also lists the jobs that are queued for printing, along with their job IDs, owners, sizes, and names12. For example, to display the status of the printer lp1, we can use the following command: $ lpq -P lp1 lp1 is

ready Rank Owner Job File(s) Total Size
active user1 123 test.txt 1024 bytes
1st user2 124 report.pdf 2048 bytes
The output shows that the printer lp1 is ready, and that there are two jobs in the queue, one of which is active and the other is waiting. The output also shows the owners, job IDs, file names, and sizes of the jobs. To display the status of all printers, we can use the -a option: $ lpq -a lp1 is ready
Rank Owner Job File(s) Total Size
active user1 123 test.txt 1024 bytes
1st user2 124 report.pdf 2048 bytes
lp2 is ready no entries
The output shows that there are two printers, lp1 and lp2, and that lp2 has no entries in the queue. To display more information about the jobs, such as the priority, submission time, and status, we can use the -l option: $ lpq -l -P lp1 lp1 is ready
Rank Owner Job File(s) Total Size
active user1 123 test.txt 1024 bytes
priority 50 Apr 27 10:00 processing since Apr 27 10:01
1st user2 124 report.pdf 2048 bytes
priority 50 Apr 27 10:05 waiting for lp1
The output shows that the jobs have the same priority, and that the first job is processing while the second job is waiting. The lpq command can be useful for troubleshooting printing problems, such as checking if the printer is ready, if there are any stuck or failed jobs, or if there are any conflicts or delays in the queue34. Reference: 1: lpq(1) - Linux manual page 2: How to Use the lp Command in Linux to Print Files From Terminal - Make Tech Easier 3: Linux sysadmin printing reference guide - PenguinTutor 4: How to manage print jobs on Linux - Network World

Question: 91

Which of the following tasks can be accomplished using the command date? (Choose TWO correct answers.)

A. Synchronize the hardware and system clocks.
B. Output date and time in different formats.
C. Set the system clock.
D. Set the hardware clock.
E. Update the time via NTP.

Answer: B

Explanation:
The date command is used to display and set the system date and time. It can also be used to print the time in different formats and calculate future and past dates. The date command has the following syntax: date [option]... [+format] The format controls begin with the % symbol and are substituted by their current values. For example, to display the current year, month, and day, we can use the following command: date +"Year: %Y, Month: %m, Day: %d" To set the system clock manually, we can use the --set or -s option followed by the date and time string. For example, to set the date and time to 5:30 PM, May 13, 2010, we can use the following command:
date --set="20100513 05:30"
The other tasks are not possible with the date command because:
Synchronize the hardware and system clocks: The date command cannot synchronize the hardware and system clocks. To do this, we need to use the hwclock command, which can read or set the hardware clock, and also synchronize it

with the system clock.
Set the hardware clock: The date command cannot set the hardware clock. To do this, we need to use the hwclock command with the --systohc or -w option, which will copy the system time to the hardware clock.
Update the time via NTP: The date command cannot update the time via NTP (Network Time Protocol). To do this, we need to use the ntpdate command, which will query an NTP server and set the system clock accordingly.
Reference:
Date Command in Linux: How to Set, Change, Format and Display Date date command in Linux with examples - GeeksforGeeks
Date Command in Linux | Linuxize

Question: 92

Which of the following are syslog facilities? (Choose TWO correct answers.)

A. local5
B. mail
C. advanced
D. postmaster
E. remote

Answer: A

Explanation:
https://learning.lpi.org/en/learning-materials/102-500/108/108.2/
The syslog facilities are predefined categories of messages that can be used to classify the source and type of the log events12. The syslog facilities are defined by the syslog protocol and are standardized across different implementations of syslog12. The syslog facilities are: auth: Security and authorization messages, such as login failures or sudo usage12.
authpriv: Same as auth, but used for private security messages that should not be available to all users12.
cron: Messages from the cron daemon, such as scheduled jobs or errors12.
daemon: Messages from system daemons, such as sshd or ntpd12.
kern: Messages from the kernel, such as boot messages or hardware errors12.
lpr: Messages from the line printer subsystem, such as print jobs or errors12.
mail: Messages from the mail subsystem, such as sendmail or postfix12.
news: Messages from the network news subsystem, such as news servers or clients12.
syslog: Messages generated internally by the syslog daemon, such as configuration errors or restarts12.
user: Messages from user-level processes, such as applications or scripts12.
uucp: Messages from the Unix-to-Unix copy subsystem, such as file transfers or errors12.
local0 to local7: Custom facilities that are not used by any system processes and can be assigned to user applications or scripts123.
Therefore, the correct answers are A. local7 and B. mail, as they are both valid syslog facilities. The other options are not syslog facilities and are either made up (C. advanced and E. remote) or refer to a specific process rather than a category of messages (D. postmaster).
Reference: 1: Prepare for LPIC-1 exam 2 - topic 108.2: System logging - IBM Developer
Tutorial 2: 108.2 System logging - Linux Professional Institute Certification Programs 3: What is the local6 (and all other local#) facilities in syslog?

Question: 93

What is the purpose of the command mailq?

A. It fetches new emails from a remote server using POP3 or IMAP.
B. It is a multi-user mailing list manager.
C. It is a proprietary tool contained only in the qmail MTA.
D. It queries the mail queue of the local MTA.
E. It is a command-line based tool for reading and writing emails.

Answer: D

Explanation:

The mailq command is a widely used tool for checking the email queue in Linux. It provides a summary of all the messages in the queue, including information such as message IDs, sender addresses, recipient addresses, and delivery status1. The mail queue is a collection of messages that are waiting to be delivered by the local Mail Transfer Agent (MTA), such as sendmail, postfix, or exim1. The mailq command is the same as the sendmail -bp command that also prints the mail queue2. The mailq command can also accept various options to filter or modify the output, such as -v for verbose mode, -Ac for mail submission queue, or -q for processing the queue3. The mailq command is part of the LPI's multi-level Linux professional certification program, and it is covered in the topic 108.3 Mail Transfer Agent (MTA) basics of the exam 102 objectives4. Reference: 4: https://www.lpi.org/our-certifications/exam-102-objectives/ 3: https://www.thegeekdiary.com/mailq-command-examples-in-linux/ 2: https://sites.ualberta.ca/dept/chemeng/AIX- 43/share/man/info/C/a_doc_lib/cmds/aixcmds3/mailq.htm 1: https://www.emaillistvalidation.com/blog/check-email-queue-in-linux-a-comprehensive-guide-to- managing-message-queues/

Question: 94

Which file inside the CUPS configuration directory contains the definition of the printers?

A. cups-devices.conf
B. snmp.conf
C. printcap.conf
D. printers.conf
E. cupsd.conf

Answer: D

Explanation:

The printers.conf file inside the CUPS configuration directory contains the definition of the printers. It is a text file that lists the names, locations, descriptions, and options for each printer queue. Each printer queue has a corresponding <Printer> or <DefaultPrinter> section in the file. The file is normally located in the /etc/cups directory and is automatically updated by the cupsd (8) daemon when printers are added or modified. The file can also be edited manually, but the changes will not

take effect until the cupsd daemon is restarted or the command cupsctl --reload-config is issued. Reference: cups-files.conf - file and directory configuration file for cups [printers.conf - printer configuration file for cups]

Question: 95

What is true regarding the command sendmail?

A. With any MTA, the sendmail command must be run periodically by the cron daemon.
B. All MTAs, including Postfix and Exim, provide a sendmail command.
C. The sendmail command prints the MTAs queue history of which mails have been sent successfully.
D. It is only available when the sendmail MTA is installed.

Answer: B

Explanation:

The sendmail command is a generic interface to various mail transfer agents (MTAs), such as Sendmail, Postfix, Exim, Qmail, etc. The sendmail command is used to send emails from the command line or from other programs that need to deliver emails. The sendmail command accepts various flags and parameters to specify the sender, recipient, subject, body, and attachments of the email. The sendmail command also reads the standard input for the email content if no file is specified. The sendmail command is part of the sendmail package, which is the original and most widely used MTA for Unix-like systems. However, other MTAs, such as Postfix and Exim, also provide a sendmail command for compatibility reasons. The sendmail command provided by these MTAs may have slightly different syntax and options, but they all support the basic functionality of sending emails. Therefore, the statement that all MTAs, including Postfix and Exim, provide a sendmail command is true123.

The other statements are false. The sendmail command does not need to be run periodically by the cron daemon, as it is not a daemon itself, but a command-line tool. The sendmail command does not print the MTA's queue history, but rather sends the email to the MTA for delivery. The sendmail command is not only available when the sendmail MTA is installed, but also when other MTAs that provide a sendmail command are installed. Reference: 1: Linux Sendmail Command Help and Examples - Computer Hope 2: Send Email in Linux from Command Line | DigitalOcean 3: 5 Ways To Send Email from Linux Command Line - TecAdmin

Question: 96

After adding a new email alias to the configuration, which command must be run in order to ensure the MTA knows about it? (Specify the command without any path but including all required parameters.)

Answer: newaliases, sendmail -bi

Explanation:

The command that must be run in order to ensure the MTA knows about the new email alias is: newaliases
This command updates the MTA's aliases database and makes the changes effective. It is equivalent to the commands

sendmail -bi or sendmail -I12. The newaliases command should be run after making modifications to the /etc/aliases file, which contains the email aliases for the system3.

Question: 97

Why is the correct configuration of a system's time zone important?

A. Because the conversion of Unix timestamps to local time relies on the time zone configuration. B. Because the time zone is saved as part of the modification times of files and cannot be changed after a file is created.
C. Because the environment variables LANG and LC_MESSAGES are, by default, set according to the time zone.
D. Because NTP chooses servers nearby based on the configured time zone.

Answer: A

Explanation:
The correct configuration of a system's time zone is important because it affects how the system displays and interprets the local time from the Unix timestamps. A Unix timestamp is a number that represents the number of seconds that have elapsed since January 1, 1970 (UTC)1. Unix timestamps are independent of time zones and are the same for all systems1. However, when a system needs to display or interpret the local time from a Unix timestamp, it needs to know the offset from UTC, which is determined by the time zone configuration23. If the time zone configuration is incorrect, the system may display or interpret the local time incorrectly, which can cause problems with scheduling tasks, logs, and other applications45.
For example, suppose a system has a Unix timestamp of 1638374400, which corresponds to December 1, 2021, 12:00:00 UTC6. If the system's time zone is configured correctly as UTC, it will display the local time as December 1, 2021, 12:00:00. However, if the system's time zone is configured incorrectly as EST (Eastern Standard Time), which is 5 hours behind UTC, it will display the local time as December 1, 2021, 07:00:00, which is 5 hours earlier than the actual local time6. This can lead to confusion and errors for the system and the user.
Therefore, the correct answer is A. Because the conversion of Unix timestamps to local time relies on the time zone configuration.
Reference: 1: Unix time - Wikipedia 2: How to Set or Change the Time Zone in Linux – TecAdmin 3: Set the date, time, and timezone on a Linux server 4: Configure the time zone (TZ) on Linux systems - Linux Audit 5: Setting the timezone under Linux - Learn Linux Configuration 6: Epoch Converter - Unix Timestamp Converter

Question: 98

Which of the following parameters are used for journalctl to limit the time frame of the output?
(Choose TWO correct answers.)

A. --from=
B. --since=
C. --until=
D. --upto=
E. --date=

Answer: B, C

Explanation:

The journalctl command is a tool for viewing and filtering the systemd journal logs. It accepts various parameters to control the output format, the source of the logs, and the filtering criteria. Two of the parameters that are used to limit the time frame of the output are --since= and --until=. These parameters take a date and time value in the format of "YYYY-MM-DD hh:mm:ss" or a relative value such as "-1h" for one hour ago. For example, the command journalctl --since="2023-11-22 23:00:00" --until="2023-11-23 00:00:00" will show the logs from 11:00 PM to 12:00 AM on November 22, 2023. The --since= and --until= parameters are part of the LPI's multi-level Linux professional certification program, and they are covered in the topic 106.1 System logging of the exam 102 objectives1. Reference: 1: https://www.lpi.org/our-certifications/exam-102-objectives/

Question: 99

Which of the following are commonly used Mail Transfer Agent (MTA) applications? (Choose THREE correct answers.)

A. Postfix
B. Procmail
C. Sendmail
D. Exim
E. SMTPd

Answer: A, C, D

Explanation:

Postfix, Sendmail, and Exim are three of the most commonly used Mail Transfer Agent (MTA) applications on Linux systems. An MTA is a software that transfers and routes electronic mail messages from one computer to another using the Simple Mail Transfer Protocol (SMTP). An MTA receives messages from another MTA or from a Mail User Agent (MUA), which is a computer application that end users use to access or send emails. An MTA can also query the MX records of the recipient's domain to find the destination mail server and forward the message accordingly. An MTA can also perform other functions such as filtering, encryption, authentication, and bounce handling.

Postfix is a cross-platform, popular MTA that was designed and developed by Wietse Zweitze Venema for his mail server while working at the IBM research department. It was primarily developed as an alternative to well-known and popular Sendmail MTA. Postfix runs on Linux, Mac OSX, Solaris, and several other Unix-like operating systems. It borrows a lot of Sendmail properties on the outside, but it has a totally and comprehensively distinct internal operation. Additionally, it bids to be fast in performance with easy configurations and secure operation mechanism1.

Sendmail, now known as Proofpoint (after Proofpoint, Inc acquired Sendmail, Inc), is by far the most popular and one of the oldest MTA on the Linux server platform. Sendmail has a lot of limitations though, in comparison to modern MTAs. Because of its complicated configuration steps and demands, and weak security mechanisms, many new MTAs have come up as alternatives to Sendmail, but importantly, it offers everything to do with mail on a network1.

Exim is a free MTA developed for Unix-like operating systems such as Linux, Mac OSX, Solaris, and many more. Exim offers a great level of flexibility in routing mail on a network, with outstanding mechanisms and facilities for incoming mail monitoring. Its notable features include among others: no support for POP and IMAP protocols, supports protocols such as RFC 2821 SMTP and RFC 2033 LMTP email message transport, configurations include access control lists, content

scanning, encryption, routing controls among others1.

Procmail is not an MTA, but a mail processing utility that can be used to filter, sort, and deliver incoming mail. It can be invoked by an MTA or run as a standalone program. Procmail can process mail based on various criteria such as sender, subject, header, body, size, date, and more. It can also execute external programs, forward mail to another address, or write mail to a file.

SMTPd is not an MTA, but a generic name for a daemon (a background process) that implements the SMTP protocol. A daemon is a program that runs continuously and performs certain tasks at predefined times or in response to certain events. An SMTP daemon listens for incoming SMTP connections from other MTAs or MUAs and handles the mail transfer accordingly. SMTPd can also refer to a specific SMTP daemon that is part of the OpenSMTPD project, which is a free implementation of the SMTP protocol for Unix systems. Reference:

7 Best Mail Transfer Agents (MTA's) for Linux

Mail Transfer Agent (MTA) Explained | Mailtrap Blog

What is a Message Transfer Agent (MTA)? - Definition from Techopedia

Mail Transfer Agent (MTA) – Glossary of Email Terms | Mailgun

[Procmail - Wikipedia]

[SMTP daemon - Wikipedia]

Question: 100

Which of the following is observed and corrected by a NTP client?

A. The skew in time between the system clock and the hardware clock.
B. The skew in time between the system clock and the reference clock.
C. Changes in the time zone of the current computer's location.
D. Adjustments needed to support Daylight Saving Time.

Answer: B

Explanation:

The Network Time Protocol (NTP) is a protocol that enables the accurate synchronization of time and date information across networked computer systems. NTP uses a hierarchical system of time servers, where each server has a stratum level that indicates its distance from the primary reference source. The primary reference source is usually an atomic clock or a GPS receiver, which provides the Coordinated Universal Time (UTC). The NTP clients are the computer systems that want to synchronize their system clocks with the UTC. The system clock is a software clock that runs in the kernel and keeps track of the current time and date. The system clock can be influenced by various factors, such as the hardware clock, the CPU frequency, the temperature, the load, and the network latency. These factors can cause the system clock to drift or skew from the UTC, resulting in inaccurate timekeeping. A NTP client observes and corrects the skew in time between the system clock and the reference clock, which is the clock of the NTP server that the client is connected to. The NTP client periodically sends requests to the NTP server and receives the server's time stamps. The NTP client then calculates the offset and the round-trip delay between its system clock and the reference clock, and adjusts its system clock accordingly. The NTP client can also use multiple NTP servers and apply algorithms to select the best one and filter out outliers. The NTP client can also discipline the system clock by using a feedback loop that controls the clock frequency and reduces the clock drift. By using NTP, the system clock can achieve a high accuracy and precision, usually within a few milliseconds or microseconds of the UTC123.

The other options are not correct. The skew in time between the system clock and the hardware clock is not observed and corrected by a NTP client, but by a separate utility called hwclock, which can read and set the hardware clock. The hardware clock is a battery-powered device that keeps time even when the system is powered off. The hardware clock is

usually less accurate than the system clock, and can be synchronized with the system clock at boot or shutdown time. The changes in the time zone of the current computer's location are not observed and corrected by a NTP client, but by a configuration tool called timedatectl, which can set the system time zone and other parameters. The time zone is a geographical region that has a uniform standard time and date. The time zone does not affect the system clock, which always keeps the UTC, but only the display of the local time and date for the user. The adjustments needed to support Daylight Saving Time (DST) are not observed and corrected by a NTP client, but by the system's time zone database, which contains the rules and transitions for DST. DST is a practice of advancing the clocks by one hour during summer months to make better use of daylight. DST is not observed in all regions and countries, and can vary in start and end dates. The system's time zone database is updated regularly to reflect the changes in DST rules, and can be applied to the system clock to calculate the correct local time and date. Reference: 1: Network Time Protocol - Wikipedia 2: How NTP Works - NTP Pool Project 3: How To Set Up Time Synchronization on Ubuntu 20.042 : hwclock(8) - Linux manual page : timedatectl(1) - Linux manual page : Daylight saving time - Wikipedia

Question: 101

Which command is used to sync the hardware clock to the system clock? (Specify ONLY the command without any path or parameters.)

Answer: hwclock, /sbin/hwclock, /usr/sbin/hwclock

Explanation:

The command that is used to sync the hardware clock to the system clock is: hwclock --systohc
This command copies the current system time to the hardware clock, which runs even when the system is shut down. It is equivalent to the command hwclock -w12. The hardware clock is also called the BIOS clock or the RTC (Real Time Clock)3.

Question: 102

Which command, available with all MTAs, is used to list the contents of the MTA's mail queue? (Specify ONLY the command without any path or parameters.)

Answer: mailq, /usr/bin/mailq, sendmail -bp, /usr/sbin/sendmail -bp, /usr/lib/sendmail -bp,

Explanation:sendmail, /usr/sbin/sendmail, /usr/lib/sendmail

The command that is used to list the contents ofthe MTA's mail queue is mailq12. This command is available with all MTAs, such as sendmail, postfix, exim, etc12. The mailq command prints the mail queue, which is the list of messages that are waiting to be sent12. The output ofthe mailq command shows the queue ID, size, time, sender, and recipient of each message12. The mailq command can also take various options to modify the output, such as -v for verbose mode, -Ac for mail submission queue, -qL for lost items, and -qQ for quarantined items3.
Reference: 1: mailq Command in Linux with Examples - GeeksforGeeks 2: mailq Command Examples in Linux – The

Geek Diary 3: linux - How to see entire sendmail queue? - Server Fault

Question: 103

Please specify the top directory containing the configuration files for the CUPS printing system. (Specify the full path to the directory.)

Answer: /etc/cups, /etc/cups/

Explanation:

The top directory containing the configuration files for the CUPS printing system is /etc/cups. This directory stores various files that control the behavior and functionality of the CUPS scheduler, cupsd (8), such as cups-files.conf (5), cupsd.conf (5), mime.convs (5), mime.types (5), printers.conf (5), and subscriptions.conf (5). The /etc/cups directory also contains subdirectories for classes, interfaces, ppd, and ssl, which store information about printer classes, device interfaces, printer drivers, and encryption certificates, respectively12. The /etc/cups directory is part of the LPI's multi-level Linux professional certification program, and it is covered in the topic 105.5 Print service of the exam 102 objectives3. Reference: 1: cups-files (5) - Linux Manuals 2: cupsd.conf - server configuration file for cups 3: Exam 102 Objectives

Question: 104

Which of the following is a legacy program provided by CUPS for sending files to the printer queues on the command line?

A. lpd B. lpp C. lpq D. lpr

Answer: D

Explanation:

The lpr command is a legacy program provided by CUPS for sending files to the printer queues on the command line. It is one of the Berkeley (lpr) printing commands that CUPS supports for compatibility with other Unix-like systems. The lpr command accepts one or more filenames as arguments and sends them to the default or specified printer. It also supports several options to control the printing process, such as the number of copies, the page size, the orientation, and the priority. The lpr command is equivalent to the lp command, which is one of the System V (lp) printing commands that CUPS also supports. However, the lp command has more options and features than the lpr command, and is recommended for use with CUPS. Reference: Command-Line Printing and Options - CUPS Command-Line Printer Administration - CUPS
Linux cups tutorial for beginners - Linux Tutorials - Learn Linux ...
CUPS Command-Line Utilities - Configuring and Managing ... - Oracle

Question: 105

What entry can be added to the syslog.conf file to have all syslog messages generated by a system displayed on console 12?

A. *.* /dev/tty12
B. /var/log/messages | /dev/tty12
C. | /dev/tty12
D. syslog tty12
E. mail.* /dev/tty12

Answer: A

Explanation:
The entry that can be added to the syslog.conf file to have all syslog messages generated by a system displayed on console 12 is A. . /dev/tty12. This entry consists of a selector field and an action field, separated by a space or a tab. The selector field specifies the pattern of facilities and priorities that

match the action. The action field specifies the destination where the matching messages are sent. In this case, the selector field is ., which means all facilities and all priorities. The action field is /dev/tty12, which is the device file for the console 12. This means that any syslog message generated by the system will be displayed on the console 12, regardless of its facility or priority. This can be useful for debugging or monitoring purposes, but it can also be very noisy and distracting, as it will show all kinds of messages, including debug, info, notice, warning, err, crit, alert, and emerg12.
The other options are not correct. Option B. /var/log/messages | /dev/tty12 is invalid, as it uses a pipe (|) character in the selector field, which is not allowed. The pipe character can only be used in the action field to indicate that the matching messages are piped to an external program1. Option C. | /dev/tty12 is also invalid, as it has an empty selector field, which is not allowed. The selector field must specify at least one facility and one priority1. Option D. syslog tty12 is also invalid, as it has a missing period (.) between the facility and the priority in the selector field, and a missing slash (/) before the device file in the action field. The correct syntax for this option would be syslog.* /dev/tty12, which would display only the messages with the syslog facility and any priority on the console 121. Option E. mail.* /dev/tty12 is valid, but it would not display all syslog messages generated by a system, but only the messages with the mail facility and any priority on the console 12. This would exclude the messages from other facilities, such as auth, cron, daemon, kern, user, etc1. Reference: 1: syslog.conf (5) - Linux man page 2: Beginner's Guide to Syslogs in Linux [Real World Examples]

Question: 106

What is true about the ntpdate command?

A. It is the primary management command for the NTP time server.
B. It updates the local system's date (i.e. day, month and year) but not the time (i.e. hours, minutes, seconds).
C. It queries one or more NTP time servers and adjusts the system time accordingly.
D. It sends the local system time to one or many remote NTP time servers for redistribution.
E. It can be used by any user to set the user clock independently of the system clock.

Answer: C

Explanation:

The ntpdate command is a tool used to synchronize the system date and time with the NTP (Network Time Protocol)

server(s) specified as arguments. It can be run manually as necessary to set the system clock, or it can be run from a cron script to periodically update the system clock. The ntpdate command has the following syntax: ntpdate [options] server [server ...]

The ntpdate command obtains a number of samples from each server and applies a subset of the NTP clock filter and selection algorithms to select the best one. It then adjusts the system clock either by stepping it (if the offset is larger than 0.5 seconds) or by slewing it (if the offset is smaller than 0.5 seconds). The ntpdate command can also be used to query the date and time from a server without setting the system clock by using the -q option.

The other statements are false because:

It is not the primary management command for the NTP time server. The primary management command for the NTP time server is ntpd, which is a daemon that runs continuously and disciplines the system clock using sophisticated algorithms.

It updates both the local system's date and time, not just the date. The ntpdate command sets the system date and time according to the configured timezone information.

It does not send the local system time to any remote NTP time servers. The ntpdate command only queries the time from the servers and does not transmit any time information to them.

It cannot be used by any user to set the user clock independently of the system clock. The ntpdate command must be run as root on the local host and it affects the system clock for all users.

Reference:

Linux ntpdate Command Tutorial – LinuxTect

ntpdate - set the date and time via NTP

How to Use NTPDATE to Sync Time in Ubuntu Linux? – TheITBros

Question: 107

What is true regarding the file ~/.forward?

A. As it is owned by the MTA and not writable by the user, it must be edited using the editaliases command.
B. After editing ~/.forward the user must run newaliases to make the mail server aware of the changes.
C. Using ~/.forward, root may configure any email address whereas all other users may configure only their own addresses.
D. When configured correctly, ~/.forward can be used to forward each incoming mail to more than one other recipient.

Answer: D

Explanation:

The file ~/.forward is a text file that contains one or more email addresses to which the incoming mail for the user will be forwarded123. The file is owned by the user and can be edited with any text editor123. The file does not require any special syntax or commands, just a list of email addresses separated by commas or newlines123. For example, if the user wants to forward their mail to alice@example.com and bob@example.com, they can create a ~/.forward file with the following content: alice@example.com, bob@example.com

The MTA will read the ~/.forward file and send a copy of each incoming mail to the specified addresses123. Therefore, the correct answer is D. When configured correctly, ~/.forward can be used to forward each incoming mail to more than one other recipient.

The other options are false regarding the file ~/.forward. The file is not owned by the MTA and does not need to be

edited with the editaliases command, which is used to edit the system-wide aliases file, not the user-specific ~/.forward file4. The user does not need to run newaliases to make the MTA aware of the changes, as the MTA will check the ~/.forward file every time a mail is delivered to the user123. The newaliases command is used to rebuild the system-wide aliases database, not the userspecific ~/.forward file4. The file ~/.forward does not have any restrictions on the email addresses
that can be used for forwarding, as long as they are valid and reachable123. The root user can also use the ~/.forward file to forward their mail, but it is not recommended for security reasons. Reference: 1: LPIC 102 – Configure e-mail aliases and forwarding on Linux using MTA -
TechViewLeo 2: topic 108.3: Mail transfer agent (MTA) basics - IBM Developer 3: 108.3 Mail Transfer Agent (MTA) basics - Linux Professional Institute ... 4: Linux At, Batch, Atq, Atrm Command Help and Examples - Computer Hope : [How to forward root's email to another email address - nixCraft]

Question: 108

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

A. rotatelog
B. striplog
C. syslogd --rotate
D. logrotate
E. logger

Answer: D

Explanation:
The logrotate command is a tool for rotating, compressing, and mailing system logs. It is designed to ease the administration of systems that generate large numbers of log files. It allows automatic rotation, compression, removal, and mailing of log files. Each log file may be handled daily, weekly, monthly, or when it grows too large. Normally, logrotate is run as a daily cron job1. The logrotate command reads the configuration files specified on the command line or in the /etc/logrotate.conf and /etc/logrotate.d directories. These configuration files can set global options and specify log files to rotate and how to handle them. For example, the compress option enables compression of old log files, the mail option sends the log files to a specified email address before being rotated, and the rotate option sets the number of log files to keep12. The logrotate command is part of the LPI's multi-level Linux professional certification program, and it is covered in the topic 106.1 System logging of the exam 102 objectives3. Reference: 1: logrotate(8) - Linux man page 2: logrotate command in Linux with examples - Linux command line tutorial 3: Exam 102 Objectives

Question: 109

To exclude all log messages of a given logging facility, you should use a logging priority of

Answer: none

Explanation:

To exclude all log messages of a given logging facility, you should use a logging priority of none. This means that no messages from that facility will be logged, regardless of their severity level. For example, if you want to exclude all

messages from the local0 facility, you can use local0.none in your syslog configuration file. This will prevent any messages from local0 from being written to any log file or destination that matches that selector12. The logging priority of none is part of the LPI's multi-

level Linux professional certification program, and it is covered in the topic 106.1 System logging of the exam 102 objectives3. Reference: 1: logging - exclude syslog facility from all others - Server Fault 2: rsyslog.conf(5) - Linux manual page - man7.org 3: Exam 102 Objectives

Question: 110

What command can be used to generate syslog entries of any facility and priority? (supply just the command name without a path)

Answer: logger

Explanation:

The logger command can be used to generate syslog entries of any facility and priority. It is a shell command interface to the syslog system log module. It allows users to write messages to the system log from the command line or from a script. The logger command supports several options to specify the facility, priority, tag, message, and other attributes of the log entry. For example, the following command generates a log entry with the facility user and the priority info: logger -p user.info "This is a test message"

The facility and priority can be any of the values defined in the syslog protocol, such as kern, mail, auth, local0, etc. for the facility, and emerg, alert, crit, err, warn, notice, info, debug, etc. for the priority. The default facility is user and the default priority is notice. The logger command can also read messages from standard input or from a file. For more information, see the logger man page or the logger - Linux man page online. Reference:

Syslogs in Linux: Understanding Facilities and Levels

What are Syslog Facilities and Levels? - Trend Micro syslog-ng Open Source Edition 3.30 - Administration Guide Syslog Logging Guide: Advanced Concepts - CrowdStrike

Question: 111

You need to pause the CUPS printer HPLaserjet4, and you want to cancel all print jobs with a message, "hello". Which command will do this?

A. cupsreject -c -r hello HPLaserjet4
B. cupsreject -p -m hello HPLaserjet4
C. cupsdisable -c -r hello HPLaserjet4
D. cupsdisable -p -m hello HPLaserjet4

Answer: C

Explanation:

The command cupsdisable -c -r hello HPLaserjet4 will pause the CUPS printer HPLaserjet4 and cancel all print jobs with a message, "hello". The cupsdisable command is used to stop printers and classes, while the cupsenable command is used to start them. The -c option cancels all jobs on the named destination, and the -r option sets the message associated with the stopped state. The message will be displayed to the users who try to print to the paused printer. For example, the

output of lpstat -p HPLaserjet4 after running the command will show:

printer HPLaserjet4 disabled since Wed 23 Jun 2023 11:54:03 AM UTC - hello

The other options are not correct. The cupsreject command is used to reject or accept jobs for a printer or class, not to pause or resume them. The -p and -m options are not valid for either cupsdisable or cupsreject. The correct syntax for cupsreject is:

cupsreject [-E] [-U username] [-h server [:port]] [-r reason] destination (s)

The -E option forces encryption of the connection to the server, the -U option uses the specified username when connecting to the server, the -h option uses the specified server and port, and the -r option sets the message associated with the rejecting state123. Reference: 1: cupsdisable(8) - Linux manual page 2: cupsreject(8) - Linux manual page 3: CUPS Administration - Page: 1.4 - Seite 3 » Raspberry Pi Geek

Question: 112

What is pool.ntp.org?

A. A deprecated feature for maintaining system time in the Linux kernel
B. A website which provides binary and source packages for the OpenNTPD project
C. A virtual cluster of various timeservers
D. A community website used to discuss the localization of Linux

Answer: C

Explanation:

C . pool.ntp.org is indeed a virtual cluster of various timeservers. It provides a reliable and easy-to- use NTP (Network Time Protocol) service for millions of clients worldwide. The pool.ntp.org project allows systems to synchronize their clocks with internet time servers, which are part of a large virtual cluster1.

Reference:

pool.ntp.org: the internet cluster of ntp servers, which explains the purpose and functioning of the pool.ntp.org project.

How do I setup NTP to use the pool?, which provides instructions on how to use pool.ntp.org for time synchronization.

NTP pool - Wikipedia, which offers additional information about the NTP pool and its role in time synchronization across the internet.

Question: 113

On a dual boot system, every time the system is booted back into Linux the time has been set backward by one day. Which of the following commands will correct the problem?

A. date -d '+ 1 day'
B. hwclock --systohc --localtime
C. ntpdate pool.ntp.org
D. time hwclock

Answer: B

Explanation:

The command that will correct the problem of the time being set backward by one day on a dual boot system is hwclock --systohc --localtime. This command will set the hardware clock (RTC) to the current system time and use the local time standard instead of UTC12. This will prevent the time inconsistency issue that occurs when dual booting Linux and Windows, as Windows assumes that the hardware clock is using local time while Linux assumes that it is using UTC34. By using the same time standard for both operating systems, the time will be displayed correctly on both Linux and Windows.

The other commands are either invalid or ineffective for solving the problem. The date -d '+ 1 day' command will display the date and time one day ahead of the current system time, but it will not change the system time or the hardware clock5. The ntpdate pool.ntp.org command will synchronize the system time with an NTP server, but it will not affect the hardware clock or the time standard6. The time hwclock command will measure the time taken by the hwclock command, which will display the hardware clock time, but it will not change anything7.

Reference: 1: hwclock(8) - Linux man page 2: How to Fix Windows and Linux Showing Different Times When Dual Booting 3: Dual-Booting Linux Messed Up Windows Time? Here's How to Fix It 4: Wrong Time in Windows 10 After Dual Boot With Linux 5: date(1) - Linux man page 6: ntpdate(8) - Linux man page 7: time(1) - Linux man page

Question: 114

Which file, when using Sendmail or a similar MTA system, will allow a user to redirect all their mail to another address and is configurable by the user themselves?

A. /etc/alias
B. ~/.alias
C. /etc/mail/forwarders
D. ~/.forward
E. ~/.vacation

Answer: D

Explanation:

The ~/.forward file is a file that users can create in their home directories to redirect mail or send mail using sendmail or a similar MTA system. The file contains a list of recipient addresses, which can be email addresses, file names, program names, or :include: files. The file must be owned by the user and have the read permission bit set for the owner. The file cannot be a symbolic link or have more than one hard link. The file is processed by sendmail when a recipient address selects a delivery agent with the F=w flag set. If the file contains a backslash, further processing is disabled and the message is delivered to the user's mail-spooling directory. If the file does not exist or cannot be read, it is silently ignored. The ~/.forward file is different from the /etc/aliases file, which is a system-wide file that maps aliases to one or more recipient addresses. The /etc/aliases file is maintained by the system administrator and requires running the newaliases command after any changes. The ~/.alias file is not a valid file for sendmail or similar MTA systems. The /etc/mail/forwarders file is not a standard file for sendmail or similar MTA systems. The ~/.vacation file is a file that contains a vacation message that is sent to the sender when the user is away. The ~/.vacation file is used in conjunction with the vacation program, which can be invoked from the ~/.forward file. Reference:

Question: 115

What command should be used to print a listing of email in the system's mail queue?

A. lpq
B. mailq
C. mlq
D. sendmail –l

Answer: B

Explanation:
The mailq command prints the list of messages that are in the mail queue. The mail queue is where outgoing mail is stored until a receiving server connection is available. The mailq command is the same as the sendmail -bp command, which also prints the mail queue. The mailq command is part of the topic 108.3: Mail transfer agent (MTA) basics, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: https://learning.lpi.org/en/learning-materials/102-500/ 2: https://www.lpi.org/our-certifications/exam-102-objectives/

Question: 116

What entry can you add to syslog.conf file to have all syslog messages generated by your system go to virtual console 12?

A. *.* /dev/tty12
B. /var/log/messages | /dev/tty12
C. | /dev/tty12
D. syslog tty12
E. mail.* /dev/tty12

Answer: A

Explanation:

The syslog.conf file is the main configuration file for the syslogd daemon, which logs system

messages on Linux systems. This file specifies rules for logging, using a selector field and an action field. The selector field consists of a facility and a priority, separated by a period. The facility indicates the subsystem that produced the message, such as mail, auth, or kern. The priority indicates the severity of the message, such as debug, info, or emerg. An asterisk (*) stands for all facilities or all priorities, depending on where it is used. The action field specifies where the message should be logged, such as a file, a user, or a device.

To have all syslog messages generated by the system go to virtual console 12, which is represented by the device file /dev/tty12, the following entry can be added to the syslog.conf file: . /dev/tty12
This means that any facility and any priority (.) should be logged to the device /dev/tty12. This will redirect all the messages that would normally go to /var/log/messages to the console 12. To see the messages, the user can press

Ctrl-Alt-F12 to switch to that console.

Reference:

syslog.conf (5) - Linux man page

Beginner's Guide to Syslogs in Linux [Real World Examples]

Configuration Formats — rsyslog 8.33-20180109-54df0f2 documentation

Topic 5, Networking Fundamentals

Question: 117

What is the command to delete the default gateway from the system IP routing table? (Choose TWO correct answers.)

A. route del default
B. ifconfig unset default
C. netstat -r default
D. ip route del default
E. sysctl ipv4.default_gw=0

Answer: A, D

Explanation:

The command to delete the default gateway from the system IP routing table is either route del default or ip route del default. Both commands will remove the default route that matches the specified parameters. The route command is the older and more widely supported tool, while the ip command is the newer and more powerful tool that can manipulate various aspects of the network configuration. The other options are either invalid or do not affect the default gateway.

Reference:

1: How to Remove Default Gateways via ip | Baeldung on Linux

2: How to remove all default gateways - Unix & Linux Stack Exchange

4: How to Add or Change the Default Gateway in Linux: 9 Steps - wikiHow

Question: 118

What is the purpose of the nsswitch.conf file?

A. It is used to configure where the C library looks for system information such as host names and user passwords.
B. It is used to configure network protocol port numbers such as for HTTP or SMTP.
C. It is used to configure LDAP authentication services for the local system.
D. It is used to configure which network services will be turned on during the next system boot.

Answer: A

Explanation:

The nsswitch.conf file is a configuration file that determines the sources and the order of the sources that are queried for

various system databases, such as user information, group information, host names, network services, and more. The C library uses this file to look up various system information when a program or a command requests it. For example, when a user logs in, the C library will use the nsswitch.conf file to determine where to find the user's password, whether it is in the local /etc/passwd file, or in a remote LDAP server, or both. The nsswitch.conf file allows the system administrator to configure the system databases in a flexible and modular way. Reference: [LPI Linux Essentials - Topic 106: The Linux Operating System]
[LPI Linux Administrator - Exam 102 Objectives - Topic 110: Security]
[Linux man page for nsswitch.conf]

Question: 119

With IPv6, how many bits have been used for the interface identifier of an unicast address? (Specify the number using digits only.)

Answer: 64

Explanation:

With IPv6, the interface identifier of an unicast address is typically a 64-bit value that is used to identify a host's network interface. The interface identifier can be derived from the MAC address of the network card, or it can be randomly generated or manually configured. The interface identifier is the rightmost 64 bits of the most commonly encountered address types, such as global unicast (2000::/3) and link-local (fe80::/10). The interface identifier is different from the network prefix, which is the leftmost bits of the address that indicate the network or subnet to which the host belongs. The network prefix can vary in length, depending on the address type and the subnetting scheme. The network prefix and the interface identifier are separated by a double colon (::) in the IPv6 address notation. For example, in the address 2001:db8:1234:5678:abcd:ef12:3456:7890, the network prefix is 2001:db8:1234:5678 and the interface identifier is abcd:ef12:3456:7890. Reference: https://study-ccna.com/ipv6-interface-identifier/https://networklessons.com/ipv6/ipv6-eui-64-explained

Question: 120

Which of the following commands will help identify a broken router between the local and the remote machine?

A. ps
B. netstat
C. nslookup
D. ifconfig
E. traceroute

Answer: E

Explanation:

The traceroute command will help identify a broken router between the local and the remote machine. The traceroute command sends a series of packets with increasing time-to-live (TTL) values to a destination and displays the routers that the packets pass through along the way. If a router is broken or unreachable, the traceroute command will show a *

symbol or a timeout message. The traceroute command is part of the topic 109.1: Fundamentals of internet protocols, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: https://learning.lpi.org/en/learning-materials/102-500/ 2: https://www.lpi.org/our- certifications/exam-102-objectives/

Question: 121

Which of the following details is NOT provided in any output from the netstat utility?

A. broadcast services
B. interface statistics
C. masquerading connections
D. network connections
E. routing tables

Answer: A

Explanation:

The netstat utility is a command-line tool that displays network connections, routing tables, interface statistics, masquerade connections, and multicast memberships. It can be used with various options to filter and customize the output. However, it does not provide any information about broadcast services, which are a type of network communication that sends data to all devices on a network segment. Broadcast services are usually handled by other tools, such as ping, traceroute, or arp. Reference: netstat | Microsoft Learn
28 Netstat Commands {A Comprehensive List With Examples} - phoenixNAP

Question: 122

Which of the following commands can be used to display the local routing table? (Choose TWO correct answers.)

A. ifconfig
B. dig
C. netstat
D. route
E. trackroute

Answer: C, D

Explanation:
The commands that can be used to display the local routing table are netstat and route. Both commands can show the kernel routing tables, which contain information about the network destinations and the gateways to reach them. The netstat command can be used with the -r option to display the routing table, and the -n option to show numeric addresses only. The route command can also be used with the -n option to display the routing table without resolving

names. However, both netstat and route are considered obsolete and have been replaced by the ip route command, which is the current recommended way of printing the routing table in Linux12. Reference: 1: Understanding Routing Table - nixCraft
2: How To Display Routing Table In Linux - RootUsers
3: linux networking - What is the local routing table used for? - Server Fault

Question: 123

Which of the following is true about IPv6?

A. With IPv6, the TCP port numbers of most services have changed.
B. IPv6 no longer supports broadcast addresses.
C. IPv4 addresses can be used without any change with IPv6.
D. IPv6 no longer supports multicast addresses.
E. For IPv6, UDP and TCP have been replaced by the Rapid Transmission Protocol RTP.

Answer: B

Explanation:
Broadcast addresses are used to send a message to all devices on a network segment. IPv4 supports broadcast addresses, but IPv6 does not. Instead, IPv6 uses multicast addresses, which are used to send a message to a group of devices that have joined a multicast group. Multicast addresses are more efficient and flexible than broadcast addresses, as they allow the sender to specify the recipients more precisely and avoid unnecessary network traffic. IPv6 also supports anycast
addresses, which are used to send a message to the nearest device that provides a specific service.
Anycast addresses are useful for load balancing and redundancy. Reference:
[LPI Linux Administrator - Exam 102 Objectives - Topic 109: Networking Fundamentals]
IPv6 - Features - Online Tutorials Library
IPv6 - Wikipedia

Question: 124

Which command is used to set the hostname of the local system? (Specify ONLY the command without any path or parameters.)

Answer: hostname

Explanation:

The hostname command is used to set the hostname of the local system. The hostname command can take a single argument, which is the new hostname to be assigned to the system. For example, to set the hostname to linux, one can run: hostname linux
The hostname command can also be used without any arguments to display the current hostname of the system. For example, to show the current hostname, one can run: hostname
The hostname command only changes the hostname temporarily, meaning that the original hostname will be restored after a reboot. To change the hostname permanently, one has to edit the configuration files that store the hostname information, such as /etc/hostname, /etc/hosts, /etc/sysconfig/network, etc. The exact files and commands may vary

depending on the Linux distribution and the system initialization process. For more details, please refer to the web search results1 or the question answering results2. Reference:

Question: 125

Which parameter must be passed to ifconfig to activate a previously inactive network interface? (Specify the parameter only without any command, path or additional options)

Answer: up

Explanation:

The parameter that must be passed to ifconfig to activate a previously inactive network interface is up. The up parameter tells the kernel to activate the network interface and allow it to send and receive packets. The opposite of up is down, which deactivates the network interface. The up parameter is part of the topic 109.2: Basic network configuration, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: https://learning.lpi.org/en/learning-materials/102-500/ 2: https://www.lpi.org/our-certifications/exam-102-objectives/

Question: 126

What is true regarding a default route?

A. The default route is always used first. When the default route is not available more specific routes are tried.
B. When a default route is set, all other routes are disabled until the default route is deleted.
C. The default route is only used if there is not a more specific route to a destination host or network.
D. Without a default route, no network communication even in directly attached networks is possible.

Answer: C

Explanation:

A default route is a special type of route that specifies where to send packets when there is no explicit route for the destination in the routing table. A default route is usually configured on a router or a gateway that connects to another network, such as the internet. A default route is often represented by the destination 0.0.0.0/0, which means any IP address.

A default route is not always used first. It is only used as a last resort, when there is no more specific route for the destination. For example, if a host wants to send a packet to 192.168.1.10, and the routing table contains the following entries:

Destination Gateway Genmask Flags Metric Ref Use Iface 192.168.1.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0 0.0.0.0 192.168.1.1 0.0.0.0 UG 0 0 0 eth0

The host will use the first entry, which is more specific, and send the packet directly to 192.168.1.10 via eth0 interface. The second entry, which is the default route, will not be used in this case. However, if the host wants to send a packet to 8.8.8.8, which is not in the same network, the host will use the default route and send the packet to 192.168.1.1, which is the gateway to the internet. Setting a default route does not disable other routes. It only adds an entry to the routing table that can be used when no other route matches the destination. Other routes are still valid and can be used if they are more specific.

Without a default route, network communication in directly attached networks is still possible, as long as there are

routes for those networks in the routing table. However, network communication to other networks that are not directly connected will not be possible, unless there are specific routes for those networks in the routing table.

Reference:

How to Set the Default Gateway in Linux - How-To Geek Linux setup default gateway with route command - nixCraft How to set a default route permanently in Linux - Xmodulo

Question: 127

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 valid lines in the file /etc/hosts are A and C. The format of the /etc/hosts file is as follows12: IP_address canonical_hostname [aliases...]

where IP_address is the IPv4 or IPv6 address of the host, canonical_hostname is the official name of the host, and aliases are optional alternative names for the host. Each field is separated by whitespace (spaces or tabs). The # character indicates the beginning of a comment, and the rest of the line is ignored.

The lines B, D, and E are invalid because they do not follow the format of the /etc/hosts file. Line B has the hostname and aliases before the IP address, which is incorrect. Line D has multiple IP addresses and hostnames separated by commas, which is also incorrect. Line E has two IP addresses for the same host, which is not supported by the /etc/hosts file. If a host has more than one IP address, it should have a separate line for each address3.

Reference:

1: hosts(5) - Linux manual page - man7.org
2: Format of /etc/hosts on Linux (different from Windows?)
3: hosts File Format for TCP/IP - IBM

Question: 128

Which of the following keywords can be used in the file /etc/nsswitch.conf to specify a source for host name lookups? (Choose TWO correct answers.)

A. resolve
B. dns
C. remote
D. files
E. hosts

Answer: B, D

Explanation:
: The keywords dns and files can be used in the /etc/nsswitch.conf file to specify a source for host name lookups. The keyword dns means that the system will use the Domain Name System (DNS) to resolve host names to IP addresses. The keyword files means that the system will use the local /etc/hosts file to resolve host names to IP addresses. The order of the keywords on the line determines the order in which the sources will be queried. For example, the following line in /etc/nsswitch.conf: hosts: files dns

means that the system will first check the /etc/hosts file for a matching host name, and if not found, it will query the DNS servers configured in /etc/resolv.conf. The other keywords in the question are not valid for the hosts database. The keyword resolv is used for the services database, which contains network service names and port numbers. The keyword remote is not a standard keyword, but it may

be used by some applications to implement their own name service providers. The keyword hosts is the name of the database itself, not a source for it. Reference:
[LPI Linux Administrator - Exam 102 Objectives - Topic 109: Networking Fundamentals]
nsswitch.conf(5) - Linux manual page
What is the /etc/nsswitch.conf file in Linux – TecAdmin

Question: 129

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:

Network interfaces may become active or inactive, the routing table may change, and IP addresses may change. Comprehensive The ifconfig command is a network management tool that is used to configure and view the status of the network interfaces in Linux operating systems1. With ifconfig, you can assign IP addresses, enable or disable interfaces, manage ARP cache, routes, and more1. Some of the possible consequences of using the ifconfig command are:
Network interfaces may become active or inactive. The ifconfig command can take an interface name as an argument and display the configuration information for that interface. For example, to view the configuration of the eth0 interface, one can run: ifconfig eth0
The output shows whether the interface is UP or DOWN, meaning active or inactive. The ifconfig command can also activate or deactivate an interface by using the up or down options. For example, to deactivate the eth0 interface, one can run: sudo ifconfig eth0 down
To activate the eth0 interface, one can run: sudo ifconfig eth0 up
The routing table may change. The ifconfig command can assign IP addresses, netmask, and broadcast address to a

network interface. For example, to assign an IP address of 192.168.1.10, a netmask of 255.255.255.0, and a broadcast address of 192.168.1.255 to the eth0 interface, one can run:

sudo ifconfig eth0 192.168.1.10 netmask 255.255.255.0 broadcast 192.168.1.255

These parameters affect the routing table, which is a list of rules that determine where the packets are sent. The routing table can be viewed by using the route command. For example, to view the routing table, one can run: route -n

The output shows the destination, gateway, netmask, flags, metric, reference, use, and interface for each route. The ifconfig command can also delete an IP address from an interface, which may remove the corresponding route from the routing table. For example, to delete the IP address of 192.168.1.10 from the eth0 interface, one can run: sudo ifconfig eth0 0

IP addresses may change. The ifconfig command can assign or delete IP addresses to a network interface, as explained above. The IP address is a unique identifier that allows the network interface to communicate with other devices on the network. The IP address can be viewed by using the ifconfig command without any options or with the interface name only. For example, to view the IP address of the eth0 interface, one can run: ifconfig eth0

The output shows the inet addr, which is the IP address of the interface. The ifconfig command can also create an alias for the network interface, which is a virtual interface that shares the same physical interface but has a different IP address. For example, to create an alias for the eth0 interface with an IP address of 192.168.1.11, one can run:

sudo ifconfig eth0:0 192.168.1.11

The alias can be viewed by using the ifconfig command with the alias name. For example, to view the configuration of the eth0:0 alias, one can run: ifconfig eth0:0

The output shows the inet addr, which is the IP address of the alias.

The other options are not correct because:

New name servers may be added to the resolver configuration. The ifconfig command does not affect the resolver configuration, which is a file that contains the names and addresses of the name servers that resolve domain names to IP addresses. The resolver configuration is stored in the /etc/resolv.conf file and can be viewed or edited by using a text editor. For example, to view the resolver configuration, one can run: cat /etc/resolv.conf

The output shows the nameserver entries, which are the IP addresses of the name servers. The ifconfig command does not add or remove name servers from this file.

The system's host name may change. The ifconfig command does not affect the host name, which is a name that identifies the system on the network. The host name is stored in the /etc/hostname file and can be viewed or edited by using a text editor. For example, to view the host name, one can run: cat /etc/hostname

The output shows the host name of the system. The ifconfig command does not change the host name of the system.

Reference:

https://linuxize.com/post/ifconfig-command/

https://www.ibm.com/docs/en/aix/7.2?topic=i-ifconfig-command

Question: 130

What is true regarding TCP port 23?

A. Port 23 is the well known port for the telnet service which is a plain text protocol that should no longer be used.
B. Port 23 is the well known port for the SSH service which provides secure logins.
C. Port 23 is the well known port for the rlogin service which is SSL secured by default.
D. Port 23 is the well known port for the system login services which are encrypted when the user runs the starttls commandin his login shell.

Answer: A

Explanation:
Port 23 is the well known port for the telnet service, which is a remote connection tool similar to SSH, but without the security of SSH. It uses a client/server model - a telnet client connects to a telnet server using TCP port 23. Due to lack of security, usage of telnet is discouraged in most situations, and firewalls routinely block port 23 to prevent incoming telnet connections. Port 23 is part of the topic 109.1: Fundamentals of internet protocols, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: https://learning.lpi.org/en/learning-materials/102- 500/ 2: https://www.lpi.org/our-certifications/exam-102-objectives/

Question: 131

How many IP-addresses can be used for unique hosts inside the IPv4 subnet 192.168.2.128/28? (Specify the number only without any additional information.)

Answer: 14

Explanation:

To find the number of IP-addresses that can be used for unique hosts inside an IPv4 subnet, we need to calculate the number of bits that are used for the host part of the IP address. The host part is the part that is not used for the network prefix, which is indicated by the slash notation (/) followed by a number. The number after the slash represents the number of bits that are used for the network prefix, out of the total 32 bits of an IPv4 address. The remaining bits are used for the host part. For example, in the subnet 192.168.2.128/28, the number 28 means that the first 28 bits are used for the network prefix, and the last 4 bits are used for the host part. The number of IP-addresses that can be used for unique hosts is equal to 2^n - 2, where n is the number of bits in the host part. The -2 is because the first and the last IP addresses in a subnet are reserved for the network address and the broadcast address, respectively, and cannot be assigned to hosts. Therefore, in the subnet 192.168.2.128/28, the number of IP-addresses that can be used for unique hosts is 2^4 - 2, which is 14.
Reference:
IPv4 - Subnetting - Online Tutorials Library
IP Subnet Calculator

Question: 132

What is the lowest numbered unprivileged TCP port? (Specify the number in digits only.)

Answer: 1024

Explanation:

The lowest numbered unprivileged TCP port is 1024. A port number is a 16-bit unsigned integer, thus ranging from 0 to 65535. The port numbers in the range from 0 to 1023 are the well-known ports or system ports. They are used by system processes that provide widely used types of network services.
On Unix-like operating systems, a process must execute with superuser privileges to be able to bind a network socket to an IP address using one of the well-known ports. Therefore, the lowest numbered port that can be used by a normal user without root access is 1024, which is the first unprivileged port123 Reference:

1: How to bind to port number less than 1024 with non root access?
2: lowest numbered unprivileged TCP port - Bing

3: List of TCP and UDP port numbers - Wikipedia
4: Privileged Ports - World Wide Web Consortium (W3C)
5: What is the lowest TCP port number? – TeachersCollegesj

Question: 133

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 statement hosts: files dns is valid in the file /etc/nsswitch.conf. It means that the system will use the local /etc/hosts file and the Domain Name System (DNS) to resolve host names to IP addresses. The order of the sources on the line determines the order in which they will be queried. In this case, the system will first check the /etc/hosts file for a matching host name, and if not found, it will query the DNS servers configured in /etc/resolv.conf. The other statements in the question are not valid in the /etc/nsswitch.conf file. The statement multi on is not a valid keyword or source for any database. The statement 192.168.168.4 dns-server is not a valid syntax for specifying a source or an action. The statement include /etc/nsswitch.d/ is not a valid way to include another file or directory in the /etc/nsswitch.conf file. Reference: [LPI Linux Administrator - Exam 102 Objectives - Topic 109: Networking Fundamentals] nsswitch.conf(5) - Linux manual page
What is the /etc/nsswitch.conf file in Linux – TecAdmin

Question: 134

Which command, depending on its options, can display the open network connections, the routing tables, as well as network interface statistics. (Specify ONLY the command without any path or parameters.)

Answer: netstat, /bin/netstat, ss, /usr/bin/ss

Explanation:

The netstat command, meaning network statistics, is a command-line utility in the Linux system to display network configuration and activity, including network connections, routing tables, interface statistics, masquerade connections, and multicast memberships1. The netstat command can display different types of network data depending on the command line option selected. Some of the common options are:

- a: This option displays active TCP connections, TCP connections with the listening state, as well as UDP ports that are being listened to.
- r: This option displays the routing table information, which is a list of rules that determine where the packets are sent.
- i: This option displays the network interface information, such as the name, MTU, RX-OK, TX-OK, etc.
- s: This option displays the network statistics by protocol, such as TCP, UDP, ICMP, IP, etc.

For example, to display the open network connections, one can run: netstat -a
To display the routing table, one can run: netstat -r
To display the network interface statistics, one can run: netstat -i
To display the network statistics by protocol, one can run: netstat -s
For more details and examples, please refer to the web search results1 or the question answering results2.
Reference:
https://netref.soe.ucsc.edu/node/7
https://bing.com/search?q=command+to+display+network+connections%2c+routing+tables%2c+and+interface+statistics

Question: 135

Which port is the default server port for the HTTPS protocol? (Specify the port number using digits.)

Answer: 443

Explanation:

The port number 443 is the default server port for the HTTPS protocol, which is a secure version of HTTP that uses SSL/TLS certificates to encrypt the data transmission between web servers and browsers. The port number 443 is recognized by the Internet Engineering Task Force (IETF) as the standard port for HTTPS connections1. The port number 443 is part of the topic 109.1: Fundamentals of internet protocols, which is one of the objectives of the LPI Linux Administrator - 102 exam23. Reference: 1: HTTPS Port: What It Is, How to Use It, and More (2023) - Hostinger 2: LPI Linux Administrator - 102 (LPIC-1) 3: Exam 102 Objectives

Question: 136

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

A. 127.0.0.0/8
B. 10.0.0.0/8
C. 169.255.0.0/16
D. 172.16.0.0/12
E. 192.168.0.0/16

Answer: B, D, E

Explanation:
According to the RFC 19181, the Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IPv4 address space for private internets:
1.1. .0.0 - 10.255.255.255 (10/8 prefix)
172.16. 0.0 - 172.31.255.255 (172.16/12 prefix)
192.168. 0.0 - 192.168.255.255 (192.168/16 prefix)
These address blocks are not globally routable and are intended for use within private networks, such as home, office, or campus networks. They can be assigned to any device that does not need to communicate directly with the public internet, or that can use network address translation (NAT) to do so. Private addresses allow for more efficient use of the limited IPv4 address space and reduce the need for public addresses.

The other options are not reserved for private use by IANA. Option A, 127.0.0.0/8, is reserved for loopback addresses, which are used to refer to the local host. Option C, 169.255.0.0/16, is a typo and should be 169.254.0.0/16, which is reserved for link-local addresses, which are used for automatic address configuration on a local network segment. Option F, 224.0.0.0/4, is reserved for multicast addresses, which are used for one-to-many communication.

Reference:

RFC 1918: Address Allocation for Private Internets - RFC Editor

IANA IPv4 Special-Purpose Address Registry

Private network - Wikipedia

Question: 137

Which of the following tools used for DNS debugging, reports not only the response from the name server but also details about the query?

A. dnsq
B. dig
C. hostname
D. dnslookup
E. zoneinfo

Answer: B

Explanation:

The tool that reports not only the response from the name server but also details about the query is dig. Dig stands for domain information groper and it is a command-line tool that can query DNS servers for various types of records. Dig can also provide additional information such as the query time, the server address, the query options, and the response code.

Dig is a powerful and flexible

tool that can be used for DNS troubleshooting and testing123 Reference:

1: How to use the dig command - Linux.com
2: dig(1) - Linux manual page - man7.org
3: Top 6 Tools for DNS Troubleshooting | Total Uptime®

Question: 138

What of the following can be done by the command ifconfig? (Choose TWO correct answers.)

A. Set a network interface active or inactive.
B. Specify the kernel module to be used with a network interface.
C. Allow regular users to change the network configuration of a network interface.
D. Change the netmask used on a network interface.
E. Specify which network services are available on a network interface.

Answer: A, D

Explanation:

The command ifconfig can be used to set a network interface active or inactive by using the up or down options. For example, the following command will activate the eth0 interface: sudo ifconfig eth0 up
The command ifconfig can also be used to change the netmask used on a network interface by specifying the netmask option followed by the desired netmask value. For example, the following command will change the netmask of the eth0 interface to 255.255.255.0: sudo ifconfig eth0 netmask 255.255.255.0
The other options in the question are not possible with the ifconfig command. The command ifconfig cannot specify the kernel module to be used with a network interface. This is done by the modprobe command or the /etc/modules file. The command ifconfig cannot allow regular users to change the network configuration of a network interface. This is controlled by the sudoers file or the polkit framework. The command ifconfig cannot specify which network services are available on a network interface. This is done by the firewall rules or the /etc/services file. Reference: [LPI Linux Administrator - Exam 102 Objectives - Topic 109: Networking Fundamentals]
Linux ifconfig Command | Linuxize
15 Useful "ifconfig" Commands to Configure Network in Linux - Tecmint ifconfig command in Linux with Examples - GeeksforGeeks

Question: 139

Which of the following programs can be used to determine the routing path to a given destination?

A. dig
B. netstat
C. ping
D. route
E. traceroute

Answer: E

Explanation:

The traceroute program can be used to determine the routing path to a given destination by sending packets with incrementing TTL values and recording the source of the ICMP time exceeded messages. This way, it can show the intermediate hops and the round-trip times for each packet. The other programs have different purposes: dig is used to query DNS servers, netstat is used to display network connections and statistics, ping is used to test the reachability of a host by sending ICMP echo requests and measuring the response time, and route is used to manipulate the routing table.

Reference:

LPI 102-500 Exam Objectives, Topic 110: Network Fundamentals, Weight: 4, 110.3 Basic network troubleshooting
LPI 102-500 Study Guide, Chapter 10: Network Fundamentals, Section 10.3: Basic Network Troubleshooting, Page 125-126

Question: 140

Given the following routing table:

```
Kernel IP routing table
    Destination     Gateway          Genmask          Flags Metric  Ref     UseIface
    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 directly transmitted on the device wlan0.
E. It would be passed to the router 192.168.1.1 on eth0.

Answer: E

Explanation:

The routing table shows how the kernel will route packets to different destinations based on the destination IP address, the gateway, the netmask, the flags, the metric, and the interface. The kernel will try to find the most specific route that matches the destination IP address, which means the route with the longest netmask. If there are multiple routes with the same netmask, the kernel will use the route with the lowest metric. If there is no matching route, the kernel will use the default

route, which is the route with the destination 0.0.0.0.
In this case, the destination IP address is 192.168.2.150, which belongs to the network 192.168.2.0/24. The routing table has a specific route for this network, which is the second entry. The gateway for this route is 0.0.0.0, which means that the packet will be directly transmitted on the interface eth0, without passing through any router. The netmask for this route is 255.255.255.0, which means that the network has 256 possible hosts. The flags for this route are U, which means that the route is up, and G, which means that the route is to a gateway. The metric for this route is 0, which means that it has the highest priority. Therefore, the kernel will use this route to handle the outgoing packet to the destination 192.168.2.150.

Reference:

How To Display Routing Table In Linux - RootUsers
route command in Linux with Examples - GeeksforGeeks Understand the basics of Linux routing | TechRepublic

Question: 141

Which of the following is a valid IPv6 address?

A. 2001:db8:3241::1
B. 2001::db8:4581::1
C. 2001:db8:0g41::1
D. 2001%db8%9990%%1

E. 2001.db8.819f..1

Answer: A

Explanation:

A valid IPv6 address is represented as a set of 16-bit hexadecimals separated by colons. The address is divided into eight groups, and each 16-bit group is represented by four hexadecimal numbers. A valid IPv6 address is in the form "x1:x2:x3:x4:x5:x6:x7:x8" where each xi is a hexadecimal string which may contain digits, lower-case English letter ('a' to 'f') and upper-case English letters ('A' to 'F'). Leading zeros are allowed in xi. The longest sequence of consecutive all-zero fields is replaced with two colons (::).1

Option A is the only one that follows these rules. Option B has two consecutive colons twice, which is not allowed. Option C has an invalid hexadecimal character 'g'. Option D uses percentage signs instead of colons, which is not a valid separator. Option E uses dots instead of colons, and has two consecutive dots, which are both invalid.

Reference: 1: IPv4 and IPv6 address formats - IBM

Question: 142

Which of the following keywords can be used in the file /etc/resolv.conf? (Choose TWO correct answers.)

A. substitute
B. nameserver
C. search
D. lookup
E. method

Answer: B, C

Explanation:

The file /etc/resolv.conf is the configuration file for the DNS resolver, which translates domain names to IP addresses by querying the DNS servers. The file supports several keywords that provide various types of resolver information. Two of the keywords that can be used in /etc/resolv.conf are: nameserver: This keyword specifies the IP address of the DNS server that the resolver can query against. Up to three nameservers can be configured, and the resolver will try them in order until one responds or all fail.

search: This keyword specifies a list of search domains that the resolver will append to the domain name when performing a query. For example, if the search list is example.com example.net, and the resolver queries for host, it will try host.example.com and host.example.net in order. The search list can have up to six domains, with a maximum of 256 characters in total.

The other keywords in the question are not valid for /etc/resolv.conf. The file does not support any keywords for substitution, lookup, or method. However, there are other keywords that can be used, such as:

domain: This keyword specifies the local domain name of the system. It is mutually exclusive with the search keyword, and only one instance of either can be used.

options: This keyword specifies various options that modify the behavior of the resolver. For example, the option rotate can be used to rotate the nameservers in a round-robin fashion, instead of trying them in order. Multiple options can be specified, separated by spaces.

Reference:
3: The /etc/resolv.conf File | Baeldung on Linux
1: /etc/resolv.conf - QNX
4: Chapter 33. Manually configuring the /etc/resolv.conf file

Question: 143

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:
The route command displays the kernel's routing table, which contains information about how packets are routed to different destinations. By default, route tries to resolve the IP addresses of the routers and destinations to their hostnames using DNS. If DNS is not working properly, this can cause a delay in displaying the routing table as route waits for the DNS queries to time out. To avoid this, route can be used with the -n option, which prevents DNS lookups and displays only numeric addresses. Reference:
LPI 102-500 Exam Objectives, Topic 109: Network Fundamentals, 109.3 Basic network troubleshooting
LPI 102-500 Study Guide, Chapter 9: Network Troubleshooting, Section 9.2: Troubleshooting Routing Problems

Question: 144

Which keyword must be listed in the hosts option of the Name Service Switch configuration file in order to make host lookups consult the /etc/hosts file?

Answer: files

Explanation:

The keyword files must be listed in the hosts option of the Name Service Switch configuration file in order to make host lookups consult the /etc/hosts file. The files service specifies that the local files, such as /etc/hosts, should be used as a source of information. The order of the services on the line determines the order in which those services will be queried, in turn, until a result is found. For example, if the hosts option is set to: hosts: files dns
then the /etc/hosts file will be searched first, and if no match is found, the DNS server will be queried next. If the hosts option is set to: hosts: dns files
then the DNS server will be queried first, and if no match is found, the /etc/hosts file will be searched next. Reference:
LPI 102-500 Exam Objectives, Topic 110: Network Fundamentals, Weight: 4, 110.3 Basic network troubleshooting
LPI 102-500 Study Guide, Chapter 10: Network Fundamentals, Section 10.3: Basic Network Troubleshooting, Page

125-126
nsswitch.conf: Name Service Switch configuration file

Topic 6, Security

Question: 145

In an xinetd configuration file, which attribute specifies the network address that will be used to provide the service?

Answer: bind, interface

Explanation:

The bind attribute in an xinetd configuration file specifies the network address that will be used to provide the service. It can be either an IP address or a hostname. If the bind attribute is not specified, xinetd will listen on all available addresses on the system. The bind attribute can be used to restrict the service to a specific interface or network. For example, bind = 192.168.1.100 will only allow the service to be accessed from the 192.168.1.0/24 network. The bind attribute can also be used to provide different configurations for the same service on different addresses. For example, one can have two telnet configuration files, one with bind = 192.168.1.100 and another with bind = 192.168.2.100, to offer different access rules or options for the telnet service on each address.

Reference:

xinetd - Wikipedia

17.4. xinetd Configuration Files - Red Hat Customer Portal

How to configure xinetd ? - Red Hat Customer Portal

Question: 146

What argument to the -type option of find will match files that are symbolic links? (Specify only the argument and no other options or words.)

Answer: l

Explanation:

The -type option of the find command allows you to specify the type of file you want to search for. The argument l (lowercase L) will match files that are symbolic links, which are special files that point to another file or directory. Symbolic links are also known as soft links or symlinks. For example, the command find /home -type l will find all the symbolic links in the /home directory and its subdirectories.

Reference: 1: How do I find all of the symlinks in a directory tree? - Stack Overflow 2: find(1) - Linux manual page - man7.org 3: Symbolic Links (GNU Findutils 4.9.0) 4: Find All Symbolic Links in Linux - Linux Handbook 5: Find all symbolic links with the find command

Question: 147

With X11 forwarding in ssh, what environment variable is automatically set in the remote host shell that is not set when X11 forwarding is not enabled? (Specify only the environment variable without any additional commands or values.)

Answer: DISPLAY, $DISPLAY

Explanation:

With X11 forwarding in ssh, the environment variable that is automatically set in the remote host shell is DISPLAY. This variable specifies the name of the X display to which X11 clients should connect. When X11 forwarding is enabled, the ssh server sets the DISPLAY variable to a value like localhost:10.0, which means that the X11 clients will connect to a proxy X11 display on the remote host. The proxy display will then forward the X11 protocol over ssh to the X server on the local host.

This way, the X11 clients can display their graphical output on the local host, even though they are running on the remote host. If X11 forwarding is not enabled, the DISPLAY variable is not set by the ssh server, and the X11 clients will not be able to connect to any X display unless the user manually sets the DISPLAY variable to a valid value. However, this may not work if the X server on the local host does not allow remote connections or if there are firewall rules that block the X11 traffic. Reference:

3: Built-in SSH X11 forwarding in PowerShell or Windows Command Prompt - X410

4: Understanding X11 Forwarding through SSH - start to finish steps

1: Why use ssh X11 forwarding with LSF; How to use ssh X11 forwarding - IBM

Question: 148

The presence of what file will temporarily prevent all users except root from logging into the system? (Specify the full name of the file, including path.)

Answer: /etc/nologin

Explanation:

The /etc/nologin file is used to prevent all users except root from logging into the system. This file is usually created by the system administrator when the system is going down for maintenance or reboot. The file can contain a message that is displayed to the users who try to log in, explaining the reason for the system shutdown. The file is automatically removed by the system when it boots up again. Reference:

LPI 102-500 Exam Objectives, Topic 104: Administrative Tasks, 104.5 Manage user accounts

LPI 102-500 Study Guide, Chapter 4: User and Group Management, Section 4.3: Preventing Users from Logging In

Question: 149

Which configuration file would be edited to change the default options for outbound SSH sessions?

A. /etc/ssh/sshd_config
B. /etc/ssh/ssh
C. /etc/ssh/client
D. /etc/ssh/ssh_config
E. /etc/ssh/ssh_client

Answer: D

Explanation:

The /etc/ssh/ssh_config file is the global configuration file for the OpenSSH client. It contains the default values for the options that apply to all outbound SSH sessions initiated from the system. The options in this file can be overridden by the user's configuration file (~/.ssh/config) or by commandline arguments. The /etc/ssh/sshd_config file is the configuration file for the OpenSSH server, and it does not affect outbound SSH sessions. The other options are not valid configuration files for OpenSSH. Reference:

LPIC-1 Exam 102 Objectives, Topic 110: Security, 110.3 Perform security administration tasks, Key Knowledge Areas: Basic client-side DNS configuration, Configure SSH and remote X
[OpenSSH manual page], FILES section, /etc/ssh/ssh_config description

Question: 150

Which of the following programs uses the hosts.allow file to perform its main task of checking for access control restrictions to system services?

A. tcpd
B. inetd
C. fingerd
D. mountd
E. xinetd

Answer: A

Explanation:
The tcpd program is a wrapper for network services that use the TCP protocol. It intercepts incoming connection requests and checks them against the rules specified in the /etc/hosts.allow and /etc/hosts.deny files. If the connection is allowed, tcpd executes the actual service program and passes the connection to it. If the connection is denied, tcpd logs the attempt and sends an error message to the client. The tcpd program can be used to enhance the security and control of network access to various services, such as SSH, FTP, Telnet, etc.
The other programs listed are not directly related to the hosts.allow file, although they may be affected by it if they are wrapped by tcpd. The inetd and xinetd programs are super-servers that listen for incoming connections and launch the appropriate service program. The fingerd program is a service that provides information about users on a remote system. The mountd program is a service that handles NFS mount requests from clients. Reference:
tcpd(8) - Linux man page
Control server access using hosts.allow and hosts.deny files hosts.allow format and example on Linux

Question: 151

Which command is used to set restrictions on the size of a core file that is created for a user when a program crashes?

A. core
B. edquota
C. ulimit

D. quota

Answer: C

Explanation:

The ulimit command is used to set or display the limitations on the system resources available to the current shell and its descendants. One of the resources that can be controlled by ulimit is the maximum size of a core file that is created when a program crashes. A core file is a snapshot of the memory and registers of a process at the time of termination, which can be used for debugging purposes. By default, the core file size limit is zero, which means no core file will be generated. To change the core file size limit, the option -c can be used with ulimit, followed by a number that represents the maximum number of blocks (usually 512 bytes) that can be written to a core file. For example, the command ulimit -c 1000 will set the core file size limit to 512000 bytes. To remove the core file size limit, the option -c can be used with ulimit, followed by unlimited. For example, the command ulimit -c unlimited will allow core files of any size to be created. Reference:
LPIC-1 Exam 102 Objectives, Topic 103: Linux Installation and Package Management, Subtopic 103.3: Manage shared libraries, Weight: 1, Key Knowledge Areas: Identify the location and purpose of important file and directories as defined in the FHS, Objective: Use the ulimit command to set or display limitations on the system resources available to the current shell and its descendants.
LPIC-1 Exam 102 Learning Materials, Topic 103: Linux Installation and Package Management, Subtopic 103.3: Manage shared libraries, Section 103.3.2: ulimit, Page 14-15.

Question: 152

When trying to unmount a device it is reported as being busy. Which of the following commands could be used to determine which process is causing this?

A. debug
B. lsof
C. nessus
D. strace
E. traceroute

Answer: B

Explanation:

The lsof command stands for list open files, and it can be used to show which processes have opened files on a device or mount point. This can help to identify which process is causing a device to be busy and prevent it from being unmounted. The syntax of the lsof command is: lsof [options] [file|directory|device]
For example, to list the processes that have opened files on the /dev/sda1 device, the command would be:
lsof /dev/sda1
The output of the lsof command will show the process ID (PID), the user name, the command name, the file descriptor, the file type, the device number, the file size, the node number, and the file name

for each open file. The file descriptor column can indicate the mode of access, such as r for read, w for write, u for read and write, and - for unknown.
The other options in the question are not relevant for this task. The debug command is used to examine and modify the memory of a running process. The nessus command is used to launch the Nessus vulnerability scanner. The strace command is used to trace system calls and signals of a process. The traceroute command is used to display the route and measure the transit delays of packets across a network.
Reference:
LPI 102-500 Exam Objectives, Topic 104.3: Manage file permissions and ownership
LPI 102-500 Study Guide, Chapter 4: Devices, Linux Filesystems, Filesystem Hierarchy Standard, Section 4.3: Mounting and Unmounting Filesystems lsof man page

Question: 153

Which configuration file would be edited to change default options for the OpenSSH server?

A. /etc/ssh/sshd_config
B. /etc/ssh/ssh
C. /etc/ssh/server
D. /etc/ssh/ssh_config
E. /etc/ssh/ssh_server

Answer: A

Explanation:

The configuration file for the OpenSSH server is called sshd_config. It is typically located in /etc/ssh on most *NIX systems, but is /etc/sshd_config in the case of MacOS X and perhaps other systems. OpenSSH has two different sets of configuration files: one for client programs (ssh, scp, and sftp) and one for the server daemon (sshd). System-wide SSH configuration information is stored in the /etc/ssh/ directory1. Reference: 1: Where is the configuration file for OpenSSH server?

Question: 154

Which of the following find commands will print out a list of files owned by root and with the SUID bit set in /usr?

A. find /usr -uid 0 -perm +4000
B. find -user root +mode +s /usr
C. find -type suid -username root -d /usr
D. find /usr -ls *s* -u root
E. find /usr -suid -perm +4000

Answer: A

Explanation:
This command will find all the files in the /usr directory that have the user ID (UID) of 0, which is the root user, and have

the permission of 4000, which is the SUID bit. The SUID bit allows the file to be executed with the privileges of the file owner, regardless of who runs it. The -uid option tests for a specific UID, and the -perm option tests for a specific permission. The + sign before the permission means that at least those bits are set; the - sign means that exactly those bits are set. The other options are either invalid or do not match the criteria. Reference:

LPIC-1 Exam 102 Objectives, Topic 104: Devices, Linux Filesystems, Filesystem Hierarchy Standard, 104.4 Find system files and place files in the correct location, Key Knowledge Areas: Search for files by type, size, or time

find manual page, -uid and -perm options description

Find Command in Linux with Practical Examples, Example 8: Find Files with SUID and SGID Permissions

Question: 155

Which directory holds the files that configure the xinetd service when using several configuration files instead of an integrated configuration file? (Specify the full path to the directory.)

Answer:
/etc/xinetd.d/,
/etc/xinetd.d

Explanation:

The /etc/xinetd.d/ directory holds the files that configure the xinetd service when using several configuration files instead of an integrated configuration file. Each file in this directory corresponds to a specific service that is managed by xinetd, such as telnet, ftp, ssh, etc. The name of the file matches the name of the service. The files in this directory contain service-specific options that override or supplement the global options defined in the /etc/xinetd.conf file. The files are read only when the xinetd service is started, so any changes require a restart of the service. The /etc/xinetd.d/ directory allows for a modular and flexible configuration of the xinetd service, as well as easier management and maintenance of the individual service files. Reference:

How to configure xinetd ? - Red Hat Customer Portal Understanding /etc/xinetd.d directory under Linux xinetd - Wikipedia

Question: 156

Which file lists which users can execute commands using sudo? (Specify the full name of the file, including path.)

Answer: /etc/sudoers

Explanation:

The /etc/sudoers file lists which users can execute commands using sudo, as well as which commands they can run, on which hosts, and as which users. The /etc/sudoers file is the main configuration file for the sudo command, which allows users to run commands as another user, usually the superuser or root. The /etc/sudoers file has a specific syntax and should be edited only with the visudo command, which checks the file for errors and locks it to prevent concurrent edits. The /etc/sudoers file contains entries that follow the format: user host = (runas) command where user is the name of the user who can run sudo, host is the name of the host where the user can run sudo, runas is the name of the user as whom the command will be executed, and command is the name of the command or a list of commands that the user can run with sudo. For example, the entry: alice ALL = (root) /bin/ls, /usr/bin/whoami means that the user alice can run sudo on any host, and can execute the commands /bin/ls and /usr/bin/whoami as the root user. The /etc/sudoers file also supports aliases, variables, wildcards, and other features that make it more flexible and powerful. For more details, see the sudoers manual page.

Reference:
LPIC-1 Exam 102 Objectives, Topic 110: Security, Subtopic 110.2: Use sudo to manage access to the root account, Weight: 2, Key Knowledge Areas: Configure sudo and sudoers. Use sudo to execute commands as another user.
LPIC-1 Exam 102 Learning Materials, Topic 110: Security, Subtopic 110.2: Use sudo to manage access to the root account, Section 110.2.1: sudo and sudoers, Page 3-5.

Question: 157

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 path.)https://lh3.googleusercontent.com/-5cd-cImKnbk/AAAAAAAAAAI/AAAAAAAAADM/-SXesH19Ido/s46-c-k-no/photo.jpg

Answer: /etc/hosts.allow

Explanation:

The /etc/hosts.allow file contains a set of rules that specify which services and hosts are allowed to connect to the server by going through a TCP Wrapper program such as tcpd. TCP Wrappers are a security mechanism that can filter incoming requests based on the source address, destination address, and service name. TCP Wrappers can also perform logging, redirection, and execution of commands based on the rules.
The /etc/hosts.allow file has the following format: service_list : host_list [: option_list]
The service_list is a comma-separated list of service names, such as sshd, telnet, or ftp. The host_list is a comma-separated list of host names, IP addresses, or network masks that are allowed to access the services. The option_list is an optional list of keywords that can modify the behavior of the rule, such as twist, spawn, deny, or allow.
For example, the following rule in /etc/hosts.allow allows ssh access from any host in the 192.168.1.0/24 network, and logs the connection attempt:
sshd : 192.168.1.0/255.255.255.0 : spawn /bin/echo %a from %h attempted to access %d >> /var/log/sshd.log
The /etc/hosts.allow file is processed before the /etc/hosts.deny file, which contains the rules for denying access to the server. If a request matches a rule in /etc/hosts.allow, it is granted access and the processing stops. If it does not match any rule in /etc/hosts.allow, it is checked against the rules in /etc/hosts.deny. If it matches a rule in /etc/hosts.deny, it is denied access and the processing stops. If it does not match any rule in either file, it is granted access by default.
Reference:
LPI 102-500 Exam Objectives, Topic 110.3: Implement host security
LPI 102-500 Study Guide, Chapter 10: Securing Your System, Section 10.3: TCP Wrappers hosts.allow man page

Question: 158

Which of the following commands preloads and manages keys that are used for automatic authentication while logging in to other machines using SSH?

A. sshd
B. ssh-agent
C. ssh-keygen
D. ssh-add

Answer: B

Explanation:

The ssh-agent command is a program that runs in the background and acts as a key manager for SSH. It can store multiple private keys in memory and provide them to SSH clients when needed. This way, the user does not have to enter the passphrase for each key every time they log in to another machine using SSH. The ssh-agent can also forward the authentication request to another agent running on the original machine, allowing the user to hop between different machines without reentering the passphrase1.
To use ssh-agent, the user needs to start it and add the private keys to it using the ssh-add command. The ssh-add command can also list, delete, and lock the keys stored in the agent. The user can then use the SSH_AUTH_SOCK environment variable to connect to the agent and use the keys for authentication2.
The other commands are not related to the ssh-agent. The sshd command is the SSH server daemon that listens for incoming connections and handles the authentication and encryption. The ssh-keygen command is a tool for generating, managing, and converting SSH keys. The ssh command is the SSH client that initiates the connection to the remote machine3. Reference: 1: SSH Essentials: Working with SSH Servers, Clients, and Keys. 2: [ssh-agent(1) - Linux manual page]. 3: SSH command usage, options, and configuration in Linux/Unix.

Question: 159

What is a purpose of an SSH host key?

A. It must be sent by any SSH client in addition to a user key in order to identify the client's host.
B. It provides the server's identity information to connecting SSH clients.
C. It is the root key by which all user SSH keys must be signed.
D. It authenticates any user that logs into a remote machine from the key's host.
E. It is used by system services like cron, syslog or a backup job to automatically connect to remote hosts.

Answer: B

Explanation:

An SSH host key is a cryptographic key used for authenticating computers in the SSH protocol. Host keys are key pairs, typically using the RSA, DSA, or ECDSA algorithms. Public host keys are stored on and/or distributed to SSH clients, and private keys are stored on SSH servers. Each host (i.e., computer) should have a unique host key. Host keys are used for authentication towards the connecting client, analogous to user SSH keys. Host keys are generated using asymmetric encryption algorithms like RSA, DSA, or ECDSA algorithms12. When a client connects to the host, the host sends its public host key to the client, and the client verifies that the host key matches the one stored in its known hosts file. If the host key is unknown or has changed, the client will display a warning and prompt the user to accept or reject the host key. This is to prevent man-in-the-middle attacks, where an attacker intercepts the connection and pretends to be the legitimate host. The other options are either incorrect or irrelevant to the purpose of an SSH host key. Reference:
What is an SSH Host Key & How are They Configured?, What are SSH Host Keys? section
SSH Host Key Management Demystified, What are SSH host keys? section
What is SSH host key - omnisecu.com, first paragraph

Question: 160

Which of the following commands can be used to limit the amount of memory a user may use?

A. umask
B. usermod
C. ulimit
D. passwd
E. chage

Answer: C

Explanation:
The ulimit command can be used to limit the amount of memory a user may use. The ulimit command is a shell builtin that allows the user to view or modify the resource limits imposed by the operating system. The resource limits can affect the maximum size of files, the maximum number of processes, the maximum amount of CPU time, and the maximum amount of virtual memory a user

can access.
To limit the amount of memory a user may use, the -v option can be used with the ulimit command. The -v option sets the maximum amount of virtual memory available to the current shell and its children in kilobytes. For example, the command ulimit -v 1000000 would limit the virtual memory to 1 GB. The -m option can also be used to set the maximum resident set size, which is the amount of physical memory used by a process, but this option is not supported by all systems.
The ulimit command can be used interactively in a shell session, or it can be placed in a shell initialization file, such as .bashrc or .profile, to apply the limits to all future shell sessions. The ulimit command can also be used in conjunction with the /etc/security/limits.conf file, which allows the system administrator to set global or per-user resource limits for all users and processes. The /etc/security/limits.conf file can specify hard and soft limits for each resource, as well as the scope of the limit, such as user, group, or domain.
The other commands listed are not related to limiting the amount of memory a user may use. The umask command sets the default file permissions for newly created files and directories. The usermod command modifies the user account information, such as the home directory, the login shell, or the password expiration date. The passwd command changes the user password. The chage command changes the password aging information, such as the minimum and maximum number of days between password changes, or the number of days before the password expires. Reference: ulimit(1) - Linux manual page
How to limit user environment with ulimit Linux command Control server access using hosts.allow and hosts.deny files

Question: 161

On a Linux system with shadow passwords enabled, which file in the file system contains the password hashes of all local users? (Specify the full name of the file, including path.)

Answer: /etc/shadow

Explanation:

On a Linux system with shadow passwords enabled, the file that contains the password hashes of all local users is /etc/shadow. This file is a replacement for the password field in /etc/passwd, which is a world-readable file that contains basic information about users. The /etc/shadow file is not readable by regular users, and it stores the encrypted passwords (or hashes) of each user, along with other information such as password expiration dates, minimum and maximum password ages, and password warning periods. The /etc/shadow file has nine colon-delimited fields for each user: Username: The name used when the user logs into the system.

Password: The encrypted password of the user, or a special character that indicates the password status. For example, an asterisk (*) means the account is locked, and an exclamation mark (!) means the password is expired.

Last Password Change: The date of the last password change, expressed as the number of days since January 1, 1970.

Minimum Password Age: The minimum number of days required between password changes. A zero means the password can be changed anytime.

Maximum Password Age: The maximum number of days the password is valid. After this number of days, the password must be changed. A zero means the password never expires.

Password Warning Period: The number of days before the password expires that the user will be warned. A zero means no warning is given.

Password Inactivity Period: The number of days after the password expires that the account will be disabled. A negative value means the account is never disabled.

Account Expiration Date: The date when the account will be disabled, expressed as the number of days since January 1, 1970. A zero means the account never expires.

Reserved Field: A field for future use.

The /etc/shadow file can be modified by using the commands passwd and chage, which are used to change the password and the password aging information of a user, respectively. The /etc/shadow file should not be edited directly, but always through the tools provided by the distribution. For more details, see the shadow manual page.

Reference:

LPIC-1 Exam 102 Objectives, Topic 110: Security, Subtopic 110.2: Use sudo to manage access to the root account, Weight: 2, Key Knowledge Areas: Configure sudo and sudoers. Use sudo to execute commands as another user.

LPIC-1 Exam 102 Learning Materials, Topic 110: Security, Subtopic 110.2: Use sudo to manage access to the root account, Section 110.2.1: sudo and sudoers, Page 3-5.

Question: 162

Which of the following commands connects to the remote host example.com which has OpenSSH listening on TCP port 2222? (Choose TWO correct answers.)

A. ssh --port 2222 example.com
B. ssh -p 2222 example.com
C. ssh -o Port=2222 example.com
D. ssh -o GatewayPort=2222 example.com
E. ssh example.com:2222

Answer: B, C

Explanation:

T The ssh command is used to connect to a remote host using the Secure Shell (SSH) protocol, which provides encrypted and authenticated communication. The ssh command has the following syntax: ssh [options] [user@]hostname

[command]

The options can modify the behavior of the ssh command, such as specifying the port number, the identity file, the cipher, the compression, and the timeout. The user@hostname specifies the username and the hostname of the remote host to connect to. The command is an optional argument that specifies the command to execute on the remote host.

To connect to the remote host example.com which has OpenSSH listening on TCP port 2222, two possible options are:

B . ssh -p 2222 example.com: This option uses the -p flag to specify the port number of the remote host. The -p flag is a shortcut for the Port option, which can also be used with the -o flag.

C . ssh -o Port=2222 example.com: This option uses the -o flag to specify a configuration option for the ssh command. The -o flag can be followed by any option that is valid in the ssh_config file, such as Port, IdentityFile, Cipher, Compression, and ConnectTimeout. The Port option sets the port number of the remote host.

The other options in the question are not correct for this task. The --port option is not a valid option for the ssh command. The GatewayPort option is used to specify whether remote hosts are allowed to connect to local forwarded ports. The example.com:2222 syntax is not valid for the ssh command. Reference:

LPI 102-500 Exam Objectives, Topic 110.1: Perform security administration tasks

LPI 102-500 Study Guide, Chapter 10: Securing Your System, Section 10.1: Configuring SSH

ssh man page

opic 7, Misc Questions New

Question: 163

Which command included in NetworkManager is a curses application which provides easy acces to the NetworkManager on the command line? (Specify only the command without any path or parameters.)

Answer: nmtui

Explanation:

The command nmtui is a curses application that provides easy access to the NetworkManager on the command line. It is included in the networkmanager package, along with nmcli, which is another command line interface for NetworkManager. nmtui allows the user to view, edit, activate and deactivate network connections, as well as set the system hostname. It has a simple and userfriendly interface that can be navigated with the keyboard or mouse12. Reference: 1: Wireless Network Manager command line ncurses GUI. 2: NetworkManager - ArchWiki.

Question: 164

What is true about the Hop Limit field in the IPv6 header?

A. The field is not changed during the transport of a package.
B. The field is transmitted within a hop-by-hop extension header.
C. Each router forwarding the packet increases the field's value.
D. Each router forwarding the packet decreases the field's value.
E. For multicast packages, the field's value is always 1.

Answer: D

Explanation:
The Hop Limit field in the IPv6 header is similar to the Time to Live (TTL) field in the IPv4 header. It specifies the maximum number of hops (routers) that a packet can traverse before reaching its destination. Each router that receives the packet decrements the Hop Limit field by one and forwards the packet. If the Hop Limit field reaches zero, the packet is discarded and an ICMPv6 error message is sent back to the source. This mechanism prevents packets from looping indefinitely in the network12 Reference: 1: IPv6 packet - Wikipedia 2: IP Time to Live (TTL) and Hop Limit Basics - Packet Pushers

Question: 165

Which of the following nmcli subcommands exist? (Choose two.)

A. nmcli ethernet
B. nmcli device
C. nmcli wifi
D. nmcli address
E. nmcli connection

Answer: B,E

Explanation:
The nmcli command is a command-line interface for NetworkManager, which is a tool for configuring and managing network settings on Linux systems. The nmcli command consists of different subcommands that correspond to different aspects of network configuration and management. The subcommands are: nmcli general: shows status and permissions of NetworkManager, as well as system hostname and logging level and domains.
nmcli connection: enables you to create, modify, activate, deactivate, delete, and show network connections.
nmcli device: enables you to show, modify, and control network devices, such as interfaces, bonds, teams, bridges, etc.
nmcli monitor: monitors activity of NetworkManager and watches for changes in the state of connectivity and devices.
nmcli networking: enables or disables overall networking.
nmcli radio: enables or disables radio transmitters for Wi-Fi, Bluetooth, and WWAN devices. nmcli agent: registers as a secret agent that provides and caches network credentials.
The other options listed are not valid nmcli subcommands. There is no nmcli ethernet, nmcli wifi, or nmcli address subcommand. However, nmcli device and nmcli connection can be used to configure and manage Ethernet and Wi-Fi connections and addresses. Reference: NetworkManager configuration and usage | SLE Micro 5.3 nmcli: NetworkManager Reference Manual - GNOME nmcli: command not found – The Geek Diary

Question: 166

Which of the following changes may occur as a consequence of using the command ip? (Choose three.)

A. Network interfaces may become active or inactive.
B. New name servers may be added to the resolver configuration.
C. The system's host name may change.
D. IP addresses may change.
E. The routing table may change.

Answer: A,D,E

Explanation:
The ip command is a versatile tool that can be used to configure and manage various aspects of the network interfaces, such as IP addresses, routes, tunnels, and more. Depending on the options and arguments used, the ip command can cause different changes to the network configuration. Some of the possible changes are:
Network interfaces may become active or inactive. The ip command can be used to bring up or down a network interface, which means to activate or deactivate its connection to the network. For example, the command ip link set eth0 up will bring up the interface eth0, while the command ip link set eth0 down will bring it down. This can affect the network connectivity and performance of the system.
IP addresses may change. The ip command can be used to assign or remove IP addresses to a network interface, which are the numerical identifiers that allow the system to communicate with other hosts in the network. For example, the command ip addr add 192.168.1.100/24 dev eth0 will assign the IP address 192.168.1.100 with a subnet mask of 255.255.255.0 to the interface eth0, while the command ip addr del 192.168.1.100/24 dev eth0 will remove it. This can affect the network reachability and routing of the system.
The routing table may change. The ip command can be used to add or delete routes to the routing table, which is a data structure that stores the information about how to reach different network destinations. For example, the command ip route add 10.0.0.0/8 via 192.168.1.1 dev eth0 will add a route to the network 10.0.0.0/8 through the gateway 192.168.1.1 using the interface eth0, while the command ip route del 10.0.0.0/8 via 192.168.1.1 dev eth0 will delete it. This can affect the network traffic and efficiency of the system.
The ip command does not affect the following settings:
New name servers may be added to the resolver configuration. The resolver configuration is a file that specifies the name servers that the system uses to resolve domain names to IP addresses. The resolver configuration file is usually /etc/resolv.conf, and it is not modified by the ip command. To add or remove name servers, the file has to be edited manually or by another tool, such as resolvconf or NetworkManager.
The system's host name may change. The host name is a human-readable name that identifies the system in the network. The host name is usually stored in the file /etc/hostname, and it is not changed by the ip command. To change the host name, the file has to be edited manually or by another tool, such as hostnamectl or nmtui.
Reference:
LPIC-1 Exam 102 Objectives, Topic 109: Networking Fundamentals, Subtopic 109.2: Persistent network configuration, Weight: 2, Key Knowledge Areas: Query and modify the behavior of network interfaces. Objective: Use the ip command to configure and modify the behavior of network interfaces.
LPIC-1 Exam 102 Learning Materials, Topic 109: Networking Fundamentals, Subtopic 109.2: Persistent network configuration, Section 109.2.2: ip, Page 17-19.

Question: 167

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: C

Explanation:
The IPv4 subnet 192.168.2.128/26 is a Class C network with a subnet mask of 255.255.255.192. This means that the network has 26 bits for the network prefix and 6 bits for the host part. To calculate the number of IP addresses that can be used for unique hosts inside the subnet, we can use the formula:
2^(number of host bits) - 2
The -2 is because the first and the last IP addresses in the subnet are reserved for the network address and the broadcast address, respectively. Therefore, the number of IP addresses for unique hosts in the subnet is:
2^6 - 2 = 64 - 2 = 62
However, this is not the correct answer, because the question asks for the number of IP addresses inside the subnet 192.168.2.128/26, not the entire network 192.168.2.0/26. A subnet is a smaller division of a network that can have its own range of IP addresses. The subnet 192.168.2.128/26 has a network address of 192.168.2.128 and a broadcast address of 192.168.2.191. Therefore, the IP addresses that can be used for unique hosts inside the subnet are:
192.168.2.129 - 192.168.2.190
This is a range of 62 IP addresses, but we have to subtract 2 more, because the question specifies that the IP addresses 192.168.2.130 and 192.168.2.140 are already in use by other hosts. Therefore, the final answer is:
62 - 2 = 60
Reference:
LPI 102-500 Exam Objectives, Topic 105.1: Customize and use the shell environment
LPI 102-500 Study Guide, Chapter 5: Customizing Shell Environments, Section 5.1: Working with the Shell
IPv4 Subnet Calculator

Question: 168

Which of the following commands configure network interfaces based on the system's existing distribution-specific configuration files? (Choose two.)

A. ifconf B. ifdown C. ifpause D. ifstart
E. ifup

Answer: B,E

Explanation:
The commands ifdown and ifup are used to configure network interfaces based on the system's existing distribution-specific configuration files. These files are typically located in /etc/network/interfaces or /etc/sysconfig/network-scripts, depending on the Linux distribution. The ifdown command shuts down a network interface, while the ifup command brings up a network interface. These commands can be used to apply changes made to the configuration files without rebooting the system12.
The other commands are not related to network interface configuration. The ifconf command does not exist in Linux. The

ifpause and ifstart commands are not standard Linux commands, but they may be aliases or scripts defined by some users or distributions. Reference: 1: NetworkConfigurationCommandLine - Community Help Wiki. 2: [How to Configure Network Static IP Address on RHEL/CentOS 8/7/6].

Question: 169

Which of the following statements is true if the UID of a regular user is identical to the GID of a group?

A. UID have precedence over GIDs, therefore the user is available while the group doesn't.
B. The user as well as the group are not available to avoid ambiguity due to the ID conflict.
C. UIDs and GIDs are independent of each other, therefore the user as well as the group are still available.
D. The user is the only member of the group, even if the group configuration contains other members.
E. GIDs have precedence over UIDs, therefore the group is available while the user isn't.

Answer: C

Explanation:
UIDs and GIDs are two different types of identifiers for users and groups in Linux. They are not related to each other, and they do not affect each other's availability or functionality. A user can have the same UID as another user's GID, or vice versa, without any problem. The only restriction is that UIDs and GIDs must be unique within their own domain, i.e., no two users can have the same UID, and no two groups can have the same GID. Having the same UID as a GID does not imply any special

relationship between the user and the group, nor does it grant any extra permissions or access rights. The user and the group are still treated as separate entities by the system123 Reference: 1: Linux sysadmin basics: User account management with UIDs and GIDs 2: How to (Correctly) Change the UID and GID of a user/group in Linux 3: Linux File Permission: uid vs gid - CBT Nuggets

Question: 170

Which of the following information is stored in /etc/shadow for each user?

A. The timestamp of the user's last login
B. The user's private SSH keys
C. The hashed password of the user
D. The numerical user ID (UID)
E. The path to the user's home directory

Answer: C

Explanation:
The /etc/shadow file is a text file that stores encrypted passwords, along with user name, password expiration values,

and last password change date. The credential information in the shadow file is encrypted using a one-way hash function to disable decryption. The /etc/shadow file contains one entry per line for each user listed in /etc/passwd file. Each line of the /etc/shadow file contains nine comma-separated fields, and the second field is the encrypted password of the user. The password field uses the $type$salt$hashed format, where $type is the method of cryptographic hash algorithm, salt is a random string, and hashed is the result of applying the hash function to the user's password and the salt. The /etc/shadow file is only readable by the root user, and it is used to enhance the security and control of user passwords.

The other information listed are not stored in /etc/shadow file, but in /etc/passwd file. The /etc/passwd file is a text file that contains basic information about each user account on the system. Each line of the /etc/passwd file contains seven colon-separated fields, and they are: Username: The name of the user account.

Password: An x character indicates that the encrypted password is stored in /etc/shadow file.

User ID (UID): The numerical identifier of the user account.

Group ID (GID): The numerical identifier of the primary group of the user account.

User ID Info: The comment field that can store additional information about the user, such as full name, phone number, etc.

Home Directory: The absolute path to the user's home directory, where the user's personal files and settings are stored.

Shell: The absolute path to the user's default login shell, which is the program that runs when the user logs in to the system.

Reference:

Understanding the /etc/shadow File | Linuxize

Understanding /etc/shadow file format on Linux - nixCraft

/etc/shadow file format | Linux# [/etc/passwd file format | Linux#]

Question: 171

Which of the following commands shows all active systemd timers?

A. systemctl-timer show
B. timectl list
C. systemctl –t
D. systemctl list-timers
E. timeq

Answer: D

Explanation:

The command systemctl list-timers shows all active systemd timers, which are units that can be used to schedule the execution of other units at specific times or after certain intervals. The output of the command includes the following columns: NEXT: The next time the timer will trigger.

LEFT: The time left until the next trigger.

LAST: The last time the timer triggered.

PASSED: The time passed since the last trigger.

UNIT: The name of the timer unit.

ACTIVATES: The name of the unit that is activated by the timer.

For example, the following output shows two active timers: apt-daily.timer and apt-daily- upgrade.timer, which are used

to perform automatic updates on Debian-based systems. NEXT LEFT LAST PASSED UNIT ACTIVATES Mon 2021-11-15 06:00:00 UTC 9h left Sun 2021-11-14 06:00:01 UTC 20h ago apt-daily.timer apt-daily.service Mon 2021-11-15 06:23:51 UTC 9h left Sun 2021-11-14 06:23:51 UTC 20h ago apt-daily-upgrade.timer apt-daily-upgrade.service 2 timers listed. The other commands in the options are either invalid or unrelated to systemd timers: systemctl-timer show is not a valid command. To show the details of a specific timer unit, the command systemctl show unit.timer can be used, where unit is the name of the unit that is activated by the timer.

timectl list is not a valid command. To list the available time zones, the command timedatectl listtimezones can be used. To list the current time and date settings, the command timedatectl can be used without any arguments.

systemctl -t is not a complete command. To list all units of a specific type, the command systemctl -t type can be used, where type is the name of the unit type, such as service, timer, socket, etc. timeq is not a valid command. It may be confused with the time command, which measures the time taken by a command or program to execute.

Reference:

LPIC-1 Exam 102 Objectives, Topic 107: Administrative Tasks, Subtopic 107.2: Automate system administration tasks by scheduling jobs, Weight: 4, Key Knowledge Areas: Use cron and systemd timers to run jobs at regular intervals and to use anacron to manage system cron jobs. Objective: Use systemd timers to run jobs at regular intervals and to use anacron to manage system cron jobs.

LPIC-1 Exam 102 Learning Materials, Topic 107: Administrative Tasks, Subtopic 107.2: Automate system administration tasks by scheduling jobs, Section 107.2.3: systemd timers, Page 21-22.

Question: 172

Which of the following tasks can the date command accomplish? (Choose two.)

A. Set the system's date and time.
B. Set the system's date but not the time.
C. Calculate the time span between two dates.
D. Print a calendar for a month or a year.
E. Display time in a specific format.

Answer: A,E

Explanation:

The date command is used to display or set the system's date and time. The date command has the following syntax:

date [options] [+format] [time]

The options can modify the behavior of the date command, such as setting the time zone, printing the date in RFC 3339 format, or updating the hardware clock. The +format argument can specify the output format of the date command, using various conversion specifiers that represent different components of the date and time, such as %Y for the year, %m for the month, %d for the day, %H for the hour, %M for the minute, and %S for the second. The time argument can set the system's date and time, using the format MMDDhhmm[[CC]YY][.ss], where MM is the month, DD is the day, hh is the hour, mm is the minute, CC is the century, YY is the year, and ss is the second.

Therefore, the date command can accomplish the following tasks:

A . Set the system's date and time. For example, to set the system's date and time to November 8, 2023, 18:30:00, the command would be: date 110818302023.00

E . Display time in a specific format. For example, to display the current date and time in the format YYYY-MM-DD HH:MM:SS, the command would be: date +%Y-%m-%d %H:%M:%S

The other options in the question are not correct for this task. The date command cannot set the system's date but not the time, as the time argument requires both the date and the time components. The date command cannot calculate the time span between two dates, as it can only display or set the current date and time. The date command cannot print a calendar for a month or a year, as that is the function of the cal command.

Reference:

LPI 102-500 Exam Objectives, Topic 105.1: Customize and use the shell environment

LPI 102-500 Study Guide, Chapter 5: Customizing Shell Environments, Section 5.1: Working with the Shell

date man page

Question: 173

Which file, if present, must contain all users that are allowed to use the cron scheduling system? (Specify the full name of the file, including path.)

Answer: crontab

Explanation:

The file /etc/cron.allow, if present, must contain all users that are allowed to use the cron scheduling system. This file is used to restrict the access to cron for security reasons. Only users listed in this file can create and edit their own crontab files using the crontab command. If the file does not exist, all users can use cron, unless the file /etc/cron.deny exists, which lists the users that are not allowed to use cron12. Reference: 1: Cron and Crontab usage and examples. 2: How to use cron in Linux.

Question: 174

What can be specified with useradd? (Choose two.)

A. Commands the user can run using sudo.
B. The absolute path to the user's home directory.
C. Which printers are available for the new user.
D. The SSH keys used to login to the new account.
E. The numeric user ID (UID) of the user.

Answer: B,E

Explanation:

The useradd command is used to create new user accounts in Linux. It has many options that can be specified to customize the user creation process. Two of these options are:

-d, --home HOME_DIR: This option allows the user to specify the absolute path to the user's home directory. The default is to append the username to the base directory specified by the HOME variable in /etc/default/useradd, or /home by default. The directory does not have to exist but will not be created if it is missing12

-u, --uid UID: This option allows the user to specify the numeric user ID (UID) of the user. The UID must be unique and not already in use by another user. The default is to use the next available UID from the range specified by the UID_MIN and UID_MAX variables in /etc/login.defs13

Reference: 1: useradd(8) - Linux man page 2: How to Create Users in Linux (useradd Command) | Linuxize 3: Linux Useradd Command Help and Examples - Computer Hope

Question: 175

Which of the following statements is true regarding systemd timer units?

A. Timer units can only be defined within a service unit's file.
B. The command executed by the timer is specified in the timer unit's [Cmd] section.
C. A dedicated system service, systemd-cron, handles the execution of timer units.
D. Timer units only exist in the system scope and are not available for users.
E. Each systemd timer unit controls a specific systemd service unit.

Answer: E

Explanation:
systemd timer units are a type of systemd unit files that define when and how to activate another systemd unit, usually a service unit. Each timer unit has a corresponding service unit that it controls, and by default, the name of the timer unit matches the name of the service unit, except for the suffix. For example, a timer unit named foo.timer activates and manages a service unit named foo.service. The service unit defines what to run when the timer elapses, and the timer unit defines when and how to run the service unit. The timer unit can specify different types of timers, such as calendar-based or monotonic timers, and various options to control the frequency, accuracy, and persistence of the timer. The timer unit can also override the service unit to activate by using the Unit= option in the [Timer] section of the timer unit file.
The other statements listed are false regarding systemd timer units. Timer units are not defined within a service unit's file, but in a separate file with the .timer extension. The command executed by the timer is not specified in the timer unit's file, but in the service unit's file that the timer controls. There is no dedicated system service named systemd-cron that handles the execution of timer units, as timer units are managed by systemd itself. Timer units can exist in both the system scope and the user scope, and users can create and manage their own timer units in their home directories.
Reference:
systemd/Timers - ArchWiki
systemd.timer - freedesktop.org
Working with systemd Timers - SUSE Documentation

Question: 176

Which of the following fields are available in the standard format of both the global /etc/crontab file as well as in user-specific crontab files? (Choose two.)

A. Year
B. Minute
C. Username
D. Effective group ID
E. Command

Answer: B,E

Explanation:

The standard format of both the global /etc/crontab file and user-specific crontab files consists of six fields separated by spaces or tabs. The first five fields indicate when to execute the command that is specified in the sixth field. The fields are:

Minute: The minute of the hour (0-59) when the command should run.

Hour: The hour of the day (0-23) when the command should run.

Day of month: The day of the month (1-31) when the command should run.

Month: The month of the year (1-12 or Jan-Dec) when the command should run.

Day of week: The day of the week (0-7 or Sun-Sat, with 0 or 7 representing Sunday) when the command should run.

Command: The command or script to execute.

For example, the following entry in a crontab file will run the command /usr/bin/backup.sh every day at 2:30 AM:

30 2 * * * /usr/bin/backup.sh

The global /etc/crontab file has an additional field between the fifth and sixth fields, which is: Username: The name of the user who will execute the command.

For example, the following entry in the /etc/crontab file will run the command /usr/bin/apt update as the root user every hour:

0 * * * * root /usr/bin/apt update

The other fields in the options are not part of the standard format of crontab files:

Year: This field is not supported by the standard cron daemon, but it may be available in some implementations, such as the Vixie cron. It would specify the year (1970-2099) when the command should run, and it would be placed after the month field.

Effective group ID: This field is not supported by any cron implementation, and it would not make sense to specify the group ID of the user who will execute the command, since it can be derived from the user ID.

Reference:

LPIC-1 Exam 102 Objectives, Topic 107: Administrative Tasks, Subtopic 107.2: Automate system administration tasks by scheduling jobs, Weight: 4, Key Knowledge Areas: Use cron and systemd timers to run jobs at regular intervals and to use anacron to manage system cron jobs. Objective: Use cron to run jobs at regular intervals.

LPIC-1 Exam 102 Learning Materials, Topic 107: Administrative Tasks, Subtopic 107.2: Automate system administration tasks by scheduling jobs, Section 107.2.1: cron, Page 18-20.

Question: 177

Which of the following commands should be executed when starting a login shell in order to change the language of messages for an internationalized program to Portuguese (pt)?

A. export LANGUAGE="pt"
B. export LC_MESSAGES="pt"
C. export UI_MESSAGES="pt"
D. export MESSAGE="pt"
E. export ALL_MESSAGES="pt"

Answer: B

Explanation:

The LC_MESSAGES environment variable is used to specify the language of messages for an internationalized program. Internationalization is the process of designing and developing a program that can adapt to different languages, cultures, and regions without requiring modifications. Localization is the process of translating and customizing a program for a specific language, culture, or region.

The LC_MESSAGES environment variable is one of the several locale categories that can affect the behavior of a program. A locale is a set of parameters that defines the user's language, country, and any special variant preferences that the user wants to see in their user interface. The locale categories are:

LC_CTYPE: Character classification and case conversion.

LC_NUMERIC: Numeric, monetary, and time formats.

LC_TIME: Date and time formats.

LC_COLLATE: Collation order.

LC_MONETARY: Monetary formats.

LC_MESSAGES: Formats of informative and diagnostic messages and interactive responses. LC_PAPER: Paper size.

LC_NAME: Name formats.

LC_ADDRESS: Address formats and location information.

LC_TELEPHONE: Telephone number formats.

LC_MEASUREMENT: Measurement units (metric or other).

LC_IDENTIFICATION: Metadata about the locale information.

The locale categories can be set individually by using the export command, such as: export LC_MESSAGES="pt" This will set the language of messages for an internationalized program to Portuguese (pt) for the current login shell and any child processes. Alternatively, the locale categories can be set collectively by using the LANG or LC_ALL environment variables, such as: export LANG="pt" This will set the default locale for all the categories to Portuguese (pt) for the current login shell and any child processes, unless overridden by another LC_* variable.

The other options in the question are not correct for this task. The LANGUAGE environment variable is used to specify a priority list of languages for programs using the GNU gettext library. The UI_MESSAGES, MESSAGE, and ALL_MESSAGES environment variables are not valid locale categories. Reference:

LPI 102-500 Exam Objectives, Topic 105.3: Localization and internationalization

LPI 102-500 Study Guide, Chapter 5: Customizing Shell Environments, Section 5.3: Localization and Internationalization

Locale man page

Question: 178

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.)

Answer: journalctl

Explanation:

The command journalctl is included in systemd and supports selecting messages from the systemd journal by criteria such as time or unit name. The systemd journal is a binary log file that stores system and service messages. The journalctl command can be used to view, filter, export, and manipulate the journal entries. For example, to show all messages from a specific unit, such as sshd.service, the command would be: journalctl -u sshd.service

To show all messages from a specific time range, such as yesterday, the command would be: journalctl --

since=yesterday

The journalctl command has many options and arguments that can be used to customize the output and perform various operations on the journal. For more information, see the man page of journalctl or the official documentation1.
Reference: [LPI 102-500 Exam Objectives], Topic 106.2: System logging, Weight: 3. [systemd-journald.service(8) — systemd — Debian unstable — Debian Manpages], Section NAME.

Question: 179

Which of the following statements about sytemd-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 a system service that collects and stores logging data from various sources, such as kernel, user-mode programs, and services1. It creates and maintains structured, indexed journals that include metadata and binary data where necessary1. The journal format is secure and unfakeable1. systemd-journald is not incompatible with syslog and can coexist with it. It can forward log messages to a syslog daemon for further processing, filtering, or storage2. This can be enabled by setting the ForwardToSyslog option to yes in the /etc/systemd/journald.conf file2. systemd-journald does not only process messages of systemd, but also messages of any other tools that use the standard logging interfaces, such as syslog(3), sd_journal_print(3), or systemd-cat(1)1. systemd- journald also supports syslog facilities, such as kern, user, and auth, which are used to specify the type of program that is logging the message3. These facilities can be used to filter the journal entries by using the -p or --priority option of the journalctl command4. For example, to show only kernel messages, we can use journalctl -p kern4. Reference: systemd-journald.service
Introduction to the Systemd journal
systemd/Journal
journalctl: Query the systemd Journal

Question: 180

Which option in the chrony configuration file changes the initial interval of polls to a NTP server in order to speed up the initial synchronization?

A. iburst
B. quickstart
C. fast
D. D. fsync
E. flood

Answer: A

Explanation:

The option in the chrony configuration file that changes the initial interval of polls to a NTP server in order to speed up the initial synchronization is iburst. The iburst option allows chronyd to send four requests to the server at intervals of 2 seconds or less, instead of the interval specified by the minpoll option, which is usually 64 seconds. This way, chronyd can make the first update of the clock shortly after start1. The iburst option is recommended for all servers, especially if the network connectivity is not reliable1.

The other options are not valid or do not have the same effect as iburst. The quickstart option does not exist in the chrony configuration file. The fast option is used to specify a fast initial correction of the system clock, but it does not change the polling interval1. The fsync option is used to enable or disable synchronization of the system clock to the real-time clock (RTC) every 11 minutes1. The flood option is used to enable a mode of operation where chronyd sends a burst of requests to the server at a high rate, which can be useful for testing or initial synchronization of a very inaccurate clock1. Reference:

LPI Linux Essentials: 1.4. Using sudo

LPI Linux Administrator: 102.5. Use Debian package management

LPI Linux Engineer: 201.1. Measure and Troubleshoot Resource Usage

LPI Linux Professional Certification Program 1

Question: 181

What is the top-level directory which contains the configuration files for CUPS? (Specify the full path to the directory.)

Answer: /etc/cups/
cups-files.conf

Explanation:

The top-level directory which contains the configuration files for CUPS is /etc/cups. CUPS stands for

Common UNIX Printing System, which is the printer and print job manager for Linux. The /etc/cups directory contains several configuration files related to CUPS, such as cupsd.conf, which is the main configuration file for the cupsd print server daemon, and printers.conf, which contains the definition of the printers. The /etc/cups directory is part of the topic 108.4: Manage printers and printing, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: LPI Linux Administrator - 102 (LPIC-1) 2: Exam 102 Objectives

Question: 182

Which of the following commands lists all queued print jobs?

A. lpd
B. lpr
C. lp
D. lsq
E. lpq

Answer: E

Explanation:

The lpq command, meaning list print queue, is a command-line utility in the Linux system to display the status of the print queue for a specified printer or class1. The lpq command can take various options and arguments to filter and format the output. By default, the lpq command shows the print queue for the default printer or class, which is determined by the PRINTER or LPDEST environment variables, or the /etc/printcap file. For example, to display the print queue for the default printer, one can run: lpq

The output shows the printer name, status, rank, owner, job number, file size, and file name for each print job. For example, the output may look like:

Printer: lp@localhost 'HP LaserJet 1020' Queue: no printable jobs in queue Server: no server active Status: printer idle. enabled since Tue 09 Nov 2021 10:00:00 AM EST Rank Owner Job File(s) Total Size 1st alice 123 report.pdf 1024 bytes 2nd bob 124 memo.docx 2048 bytes 3rd charlie 125 presentation.pptx 4096 bytes

This shows that there are three print jobs in the queue for the printer lp@localhost, which is an HP LaserJet 1020. The first job belongs to alice, the second to bob, and the third to charlie. The lpq command can also display the print queue for a specific printer or class by using the -P option, followed by the printer or class name. For example, to display the print queue for the printer lp1, one can run: lpq -P lp1

The output shows the print queue for the printer lp1, which may be different from the default printer. The lpq command can also display the print queue for all printers or classes by using the -a option. For example, to display the print queue for all printers or classes, one can run: lpq -a

The output shows the print queue for each printer or class, separated by a blank line. The lpq command is different from the following commands:

lpd: This is not a command, but a daemon that provides print spooling and network printing services for the Linux system2. The lpd daemon is started by the system initialization scripts and runs in the background. It does not display the print queue for any printer or class.

lpr: This is a command that submits print jobs to the print queue for a specified printer or class3. The lpr command can take various options and arguments to specify the print options and the files to be printed. For example, to print the file report.pdf to the default printer, one can run:

lpr report.pdf

The lpr command does not display the print queue for any printer or class.

lp: This is a command that submits print jobs to the print queue for a specified printer or class, similar to the lpr command4. The lp command can take various options and arguments to specify the print options and the files to be printed. For example, to print the file report.pdf to the printer lp1, one can run:

lp -d lp1 report.pdf

The lp command does not display the print queue for any printer or class.

lsq: This is not a valid command in the Linux system. It may be a typo or a misspelling of the lpq command. It does not display the print queue for any printer or class. Reference: 1

Question: 183

Which of the following entries in /etc/syslog.conf writes all mail related events to the file /var/log/maillog and sends all critical events to the remote server logger.example.com?

A. mail.*/var/log/maillogmail,crit@logger.example.org
B. mail.*/var/log/maillogmail.critsyslog://logger.example.org

C. mail/var/log/maillogmail.crit@logger.example.org
D. mail.*/var/log/maillogmail.crit@logger.example.org
E. mail */var/log/maillogmail crit@logger.example.org

Answer: D

Explanation:
The /etc/syslog.conf file is used to configure the syslog daemon, which handles the logging of system messages. The file consists of lines that have the following format: selector action
The selector specifies the type and priority of the messages to be logged, and the action specifies what to do with the messages. The selector has two parts, separated by a dot: the facility and the priority. The facility indicates the source of the message, such as mail, auth, kern, etc. The priority indicates the severity of the message, such as emerg, alert, crit, err, etc. A priority can also be preceded by an equal sign (=) to match only that priority, or a minus sign (-) to match all priorities except that one.
The action can be one of the following:
A filename, starting with a slash (/), indicating the file to write the messages to.
A hostname, preceded by an at sign (@), indicating the remote host to send the messages to via UDP. A username, indicating the user to send the messages to via wall.
An asterisk (*), indicating all logged-in users.
A pipe symbol (|), followed by a command, indicating the program to pipe the messages to.
In this question, the correct entry is D. mail.*/var/log/maillogmail.crit@logger.example.org. This entry means:
Log all mail related messages (mail.*) to the file /var/log/maillog.
Log all critical mail messages (mail.crit) to the remote host logger.example.org.
The other options are incorrect because:
Option A is missing a dot between mail and crit, and uses the wrong domain name (example.org instead of example.com).
Option B uses an invalid action (syslog://logger.example.org) that is not supported by syslog.conf.
Option C is missing a dot between mail and *.
Option E is missing dots between mail and * and between mail and crit.
Reference: [LPI 102-500 Exam Objectives], Topic 106.1: Maintain system time, Weight: 3. [LPI Linux Essentials Study Guide], Chapter 9: Administrative Tasks, Section 9.4: System Logging.

Question: 184

Which of the following protocols is related to the term open relay?

A. SMTP
B. POP3
C. NTP
D. IMAP
E. LDAP

Answer: A

Explanation:

SMTP stands for Simple Mail Transfer Protocol, which is a standard for sending and receiving email messages over the Internet. An open relay is a SMTP server that is configured to allow anyone on the Internet to send email through it, without verifying the identity or the origin of the sender. This can be exploited by spammers and hackers to send large volumes of unwanted or malicious email, while hiding their true identity and location. Open relays are considered a security risk and a source of spam, and are often blocked or blacklisted by other email servers. To prevent open relays, SMTP servers should be configured to only accept email from authorized users or domains, and to reject or filter email from unknown or suspicious sources. Reference:

Open mail relay - Wikipedia

What is Open Relay? - Definition from Techopedia

What is an open relay and how do I close one? - Validity Help Center

LPI Linux Essentials: 1.5 Security and File Permissions: 1.5.3 Network Security

LPIC-1: System Administrator: 102.5 Implement basic network security: 102.5.1 TCP Wrappers

Question: 185

Which of the following commands displays all environment and shell variables?

A. getargs
B. lsenv
C. ls
D. env
E. lsshell

Answer: D

Explanation:

The command that displays all environment and shell variables is env. The env command prints a list of the current environment variables, which are variables that are defined for the current shell and are inherited by any child shells or processes. The env command can also be used to run another program in a custom environment without modifying the current one1. The env command is part of the GNU coreutils package and is available on most Linux systems2.

The other commands are not valid or do not display all environment and shell variables. The getargs command does not exist in Linux. The lsenv command is a utility to list information about IBM Power Systems firmware, not environment variables. The ls command lists the files and directories in the current working directory, not environment variables. The lsshell command is a utility to list the available shells on a system, not environment variables.

Reference:

LPI Linux Essentials: 1.4. Using sudo

LPI Linux Administrator: 102.5. Use Debian package management

LPI Linux Engineer: 201.1. Measure and Troubleshoot Resource Usage

LPI Linux Professional Certification Program

https://linuxconfig.org/how-to-set-and-list-environment-variables-on-linux

https://www.digitalocean.com/community/tutorials/how-to-read-and-set-environmental-and-shell- variables-on-linux

https://unix.stackexchange.com/questions/176001/how-can-i-list-all-shell-variables

Question: 186

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 comparison operators for test that work on elements in the file system are -d and -f. The -d operator tests if a given file name refers to a directory, and returns true if it does. The -f operator tests if a given file name refers to a regular file, and returns true if it does. These operators are part of the topic 105.3: Perform basic file management, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: LPI Linux Administrator - 102 (LPIC-1) 2: Exam 102 Objectives

Question: 187

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 echo command is a built-in Linux feature that prints out arguments as the standard output1. The echo command can take various options and arguments to display different types of information. One of the arguments that can be used with the echo command is $$, which represents the process ID (PID) of the current shell2. A process ID is a unique number that identifies a running process in the system. The current shell is the shell that is executing the echo command. For example, if you are using the Bash shell and run the following command: echo $$
The output will show the PID of the Bash shell, such as: 1234
The echo commandcanbeusefultocheckwhichshellyouareusing,ortofindoutthePIDofthecurrentshellfordebuggingormonitoringpurposes.Theecho command is different from the following commands:
echo $!: This command displays the PID of the last command executed in the background2. A background command is a

command that runs without blocking the shell, allowing you to continue using the shell while the command executes. For example, if you run the following command: sleep 10 &

This command will put the sleep command, which pauses the execution for 10 seconds, in the background. The output will show the PID of the sleep command, such as: 1 2345

If you then run the following command:

echo $!

The output will show the same PID of the sleep command, such as: 2345

echo $?: This command displays the exit status of the last command executed2. The exit status is a number that indicates whether the command was successful or not. A zero exit status means the command was successful, while a non-zero exit status means the command failed or encountered an error. For example, if you run the following command: ls /home

This command will list the contents of the /home directory. If the command succeeds, the output will show the files and directories in the /home directory, such as: alice bob charlie

If you then run the following command: echo $?

The output will show the exit status of the ls command, which is zero, meaning the command was successful:

0

echo $0: This command displays the name of the current shell or script2. The name of the current shell is the name of the executable file that runs the shell, such as bash, zsh, ksh, etc. The name of the current script is the name of the file that contains the script, such as script.sh, script.py, etc. For example, if you are using the Bash shell and run the following command: echo $0

The output will show the name of the current shell, such as: bash

Reference:

2

Question: 188

What output is produced by the following command sequence? echo '1 2 3 4 5 6' | while read a b c; do echo result $c $b $a;

done

A. result: 6 5 4
B. result: 1 2 3 4 5 6
C. result: 3 4 5 6 2 1
D. result: 6 5 4 3 2 1
E. result: 3 2 1

Answer: A

Explanation:

The command sequence uses a while loop to read the input from the echo command and assign the values to the variables a, b, and c. The read command reads one line of input at a time and splits it into words according to the IFS variable, which is a space by default. The first word is assigned to the first variable, the second word to the second variable, and so on. If there are more words than variables, the remaining words are assigned to the last variable. In this case, the input line has six words, so the read command assigns 1

to a, 2 to b, and 3 4 5 6 to c. Then, the echo command prints the result with the variables in reverse order, i.e., $c $b $a. Therefore, the output is result: 6 5 4 3 2 1. However, the answer choices only show the first three words of the output, so the correct answer is A. result: 6 5 4. Reference: [LPI 102-500 Exam Objectives], Topic 105.3: Customize and use the shell environment, Weight: 4. [LPI Linux Essentials Study Guide], Chapter 7: The Linux Operating System, Section 7.2: Shell Scripting.

Question: 189

What is the purpose of TCP wrapper?

A. Manage and adjust bandwidth used by TCP services.
B. Bind a network service to a TCP port.
C. Encapsulate TCP messages in IP packets.
D. Add SSL support to plain text TCP services.
E. Limit access to a network service.

Answer: E

Explanation:

TCP wrapper is a security tool that allows you to restrict the access to a network service based on the source IP address or hostname of the client. TCP wrapper works by intercepting the incoming connection requests to a service and checking them against a set of rules defined in the /etc/hosts.allow and /etc/hosts.deny files. If the client is allowed, the connection is passed to the service. If the client is denied, the connection is rejected and an error message is logged. Reference: LPI Linux Essentials: 1.5 Security and File Permissions: 1.5.3 Network Security
LPIC-1: System Administrator: 102.5 Implement basic network security: 102.5.1 TCP Wrappers

Question: 190

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 sudo configuration file (/etc/sudoers) defines the rules for granting privileges to users or groups to execute commands as another user, usually the superuser or root. The format of the sudo configuration file is as follows:

user_list host_list=effective_user_list tag_list command_list

The user_list specifies the users who can run the commands, the host_list specifies the hosts where the commands can be run, the effective_user_list specifies the user as whom the commands can be run, the tag_list specifies some options for the commands, and the command_list specifies the commands that can be run.

In this case, the user_list is jane, the host_list is ANY (meaning any host), the effective_user_list is not specified (meaning root by default), the tag_list is NOPASSWD or PASSWD (meaning whether a password is required or not), and the command_list is /bin/kill, /bin/id, or /sbin/fdisk.

Therefore, the sudo configuration file allows jane to run /bin/kill, /bin/id, or /sbin/fdisk as root on any host, but with different password requirements. Specifically:

Jane can run /bin/kill without specifying a password, because the tag_list is NOPASSWD for this command. This means that jane can execute sudo /bin/kill and the command will run as root without asking for any password. This makes option D true.

Jane can run /bin/id without specifying a password, because the tag_list is also NOPASSWD for this command. This means that jane can execute sudo /bin/id and the command will run as root without asking for any password. This makes option E true.

Jane can run /sbin/fdisk after specifying her password, because the tag_list is PASSWD for this command. This means that jane can execute sudo /sbin/fdisk and the command will ask for jane's password before running as root. This makes option C true.

The other options are false because:

Jane cannot run /bin/id only after specifying her password, because the tag_list is NOPASSWD for this command. This makes option A false.

Jane cannot run /sbin/fdisk after specifying root's password, because the password that is required is jane's password, not root's password. This makes option B false.

Reference:

LPI Linux Essentials: 1.4. Using sudo

LPI Linux Administrator: 102.5. Use Debian package management

LPI Linux Engineer: 201.1. Measure and Troubleshoot Resource Usage

LPI Linux Professional Certification Program

Question: 191

Which configuration file contains the default options for SSH clients?

A. /etc/ssh/sshd_config
B. /etc/ssh/ssh
C. /etc/ssh/ssh_config
D. /etc/ssh/client
E. /etc/ssh/ssh_client

Answer: C

Explanation:
The configuration file that contains the default options for SSH clients is /etc/ssh/ssh_config. This file is read by the ssh program when it connects to a remote SSH server. It can contain global options that apply to all hosts, or host-specific options that only apply to certain hosts or patterns. The /etc/ssh/ssh_config file is the system-wide default SSH client configuration file. It can be overridden by a user-specific configuration file ~/.ssh/config, which is located in the user's home directory. The /etc/ssh/ssh_config file is part of the topic 110.1: Perform security administration tasks, which is one of the objectives of the LPI Linux Administrator - 102 exam12. Reference: 1: LPI Linux Administrator - 102 (LPIC-1) 2: Exam 102 Objectives

Question: 192

Which of the following commands can identify the PID od a process which opened a TCP port?

A. ptrace B. strace C. debug D. lsof
E. nessus

Answer: D

Explanation:
The lsof command, meaning list open files, is a command-line utility in the Linux system to display information about files that are opened by processes1. The lsof command can take various options and arguments to filter and format the output. One of the options that can be used to identify the PID of a process which opened a TCP port is the -i option, which selects the listing of files whose Internet address matches the specified address. The address can be specified as a port number, a host name, or a combination of both. For example, to list the processes that are listening on TCP port 80, one can run: lsof -i TCP:80
The output shows the command name, the PID, the user name, the file descriptor, the type, the device, the size/off, the node, and the name for each process. The name column shows the local and remote addresses and port numbers for the TCP connection. For example, the output may look like: COMMAND PID USER FD TYPE DEVICE SIZE/OFF NODE NAME httpd 1234 root 4u IPv4 12345 0t0 TCP *:80 (LISTEN) httpd 2345 www-data 4u IPv4 12345 0t0 TCP *:80 (LISTEN) httpd 3456 www-data 4u IPv4 23456 0t0 TCP 192.168.1.10:80->192.168.1.20:1234 (ESTABLISHED)
This shows that the httpd command, which is the Apache web server, is listening on TCP port 80 with

the PID 1234 and 2345, and has an established connection with the remote address 192.168.1.20 and port 1234 with the PID 3456. To kill the process by PID, one can use the kill command with the -SIGTERM option, which sends a termination signal to the process. For example, to kill the process with the PID 3456, one can run:
kill -SIGTERM 3456
The other options are not correct because:
ptrace: This is not a command, but a system call that allows a process to trace and control the execution of another process. It is used by debuggers and other tools that need to monitor and manipulate the behavior of other processes2. It does not display the PID of a process which opened a TCP port.
strace: This is a command that traces the system calls and signals of a process. It can be used to diagnose, debug, and monitor the interaction between a process and the kernel3. It does not display the PID of a process which opened a TCP port.

debug: This is not a command, but a general term that refers to the process of finding and fixing errors in a program or system. There are various tools and methods that can be used for debugging, such as debuggers, loggers, profilers, etc4.
It does not display the PID of a process which opened a TCP port.
nessus: This is a command that runs the Nessus vulnerability scanner, which is a tool that scans a network or a system for security flaws and potential attacks5. It does not display the PID of a process which opened a TCP port. Reference:
https://www.howtogeek.com/28609/how-can-i-tell-what-is-listening-on-a-tcpip-port-in-windows/
https://bing.com/search?q=identify+PID+of+process+that+opened+a+TCP+port

Question: 193

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 is a protocol that allows users to access graphical applications on a remote host, such as a virtual machine or a server, using a client program. SPICE also supports connecting local USB devices to remote applications, such as printers, scanners, or flash drives, using a feature called USB redirection. SPICE does not replace Xorg as the local X11 server, nor does it allow downloading and locally installing applications or uploading and running binary programs from a remote machine. These are features that are provided by other tools, such as SSH, SCP, or RDP. Reference:
Features - spice-space.org
SPICE - Wikipedia
SPICE Model ＜What is SPICE?＞ | Electronics Basics | ROHM

Question: 194

Depending on a system's configuration, which of the following files can be used to enable and disable network services running on this host?

A. /etc/profile
B. /etc/xinetd.conf
C. /etc/ports
D. /et/host.conf
E. /etc/host.conf

Answer: B

Explanation:

The file that can be used to enable and disable network services running on this host is /etc/xinetd.conf. This file is the main configuration file for the xinetd daemon, which is a superserver that can start other network services on demand. The /etc/xinetd.conf file contains global settings and a list of services that xinetd can manage. Each service has its own configuration section, which can include the disable keyword to enable or disable the service. For example, to disable the telnet service, the configuration section would look like this12: service telnet {

disable = yes socket_type = stream protocol = tcp wait = no

user = root

server = /usr/sbin/in.telnetd log_on_failure += USERID

}

The other files are not related to network services. The /etc/profile file is a global configuration file for the Bash shell, which sets environment variables and aliases for all users3. The /etc/ports file does not exist by default in Linux, and it is not a standard file for network configuration. The /et/host.conf file is a typo, and it should be /etc/host.conf, which is a file that controls the behavior of the resolver library, which is used to look up host names and IP addresses4. The /etc/host.conf file is not used to enable or disable network services, but to specify the order of host name resolution methods5. Reference:

1: How to enable or disable services with xinetd - LinuxConfig.org

2: xinetd.conf(5) - Linux manual page - man7.org

3: What is /etc/profile file in Linux? - LinuxForDevices

4: host.conf(5) - Linux manual page - man7.org

5: Linux host.conf file - Computer Notes

Question: 195

On a machine running serval X servers. how do programs identify the different instances of theX11 server?

A. By a fixed UUID that is defined in the X11 configuration file.

B. By a display name like :1.

C. By the name of the user that turns the X server like x11: bob.

D. By a device name like /dov/x11/xeorvore/1.

E. By a unique IPv6 address from the fc80: /64 subnet.

Answer: B

Explanation:

On a machine running several X servers, programs identify the different instances of the X11 server by a display name like :1. The display name consists of three parts: the hostname, the display number, and the screen number. The hostname is the name of the machine where the X server runs. The display number is a unique identifier that distinguishes different X server instances on the same machine. The screen number is used to address different physical screens that are managed by the same X server instance. The display name has the format hostname:displaynumber.screennumber. If the hostname is omitted, it means the local machine. The screen number is also optional and defaults to 0. For example, :1 means the second X server instance on the local machine, screen 0. remote:0.1 means the first X server instance on the remote machine, screen 1. The display name is part of the topic 106.1: Install and configure X11, which is one of the objectives of the LPI Linux Administrator - 102 exam12.

Reference: 1: LPI Linux Administrator - 102 (LPIC-1) 2: Exam 102 Objectives

Question: 196

Which of the following connection types, as seen in unroll connection show, may exist in Network Manager? (Choose THREE correct answers.)

A. tcp
B. Ethernet
C. wifi
D. ipv6
E. bridge

Answer: BCE

Explanation:
The connection types, as seen in nmcli connection show, are the types of network configurations that Network Manager can manage. They are not the same as the network protocols or layers, such as TCP or IPv6, but rather the logical or physical ways of connecting to a network. According to the Network Manager reference manual1, some of the possible connection types are:
wifi: This connection type is for wireless network interfaces that use the IEEE 802.11 standard. It requires a wifi device and a wifi access point to establish a connection. The connection settings include the SSID, security, password, etc.
bridge: This connection type is for creating a network bridge, which is a device that connects two or more network segments and forwards packets between them. It requires a bridge device and one or more slave devices to be attached to the bridge. The connection settings include the bridge name, MAC address, STP, etc.
vpn: This connection type is for creating a virtual private network, which is a secure tunnel between two or more network endpoints. It requires a VPN plugin and a VPN service provider to establish a connection. The connection settings include the VPN type, service name, user name, password, etc. The other options are not correct because:
tcp: This is not a connection type, but a network protocol that operates at the transport layer. It provides reliable, ordered, and error-checked delivery of data between applications. It is not a configuration option for Network Manager.
Ethernet: This is not a connection type, but a network technology that operates at the physical and data link layers. It defines the standards for wiring, signaling, and framing of data packets. It is not a configuration option for Network Manager, but rather a device type that can be used by other connection types, such as bridge or vpn.
ipv6: This is not a connection type, but a network protocol that operates at the network layer. It provides addressing and routing of data packets across networks. It is not a configuration option for Network Manager, but rather an IP configuration option that can be used by other connection types, such as wifi or vpn. Reference: https://www.networkmanager.dev/docs/api/latest/nm-settings-nmcli.html

Question: 197

How do shadow passwords improve the password security in comparison to standard non-shadow passwords?

A. Regular users do not have access to the password hashes of shadow passwords.
B. Every shadow password is valid tor 45 days and must be changed afterwards.

C. The system's host key is used to encrypt all shadow passwords.
D. Shadow passwords are always combined with a public key that has to match the user's private key.
E. Shadow passwords are stored in plain text and can be checked for weak passwords.

Answer: A

Explanation:

Question: 198

Which mechanism does ggfe use to interact with the SSH agent?

A. Connecting to pod 2222 which is used by the system-wide SSH agent.
B. Using the fixed socket .BBh-ngant/ipe.
C. Creating an alias replacing anh with calls to ssh-agent.
D. Starting ssh-agent as a child process for each sir. invocation.
E. Evaluating environment variables such as SSH_AUTH_SOCK.

Answer: E

Explanation:

Question: 199

Which directory holds configuration files for xinetd services? (Specify the full path to the directory

Answer: /etc/xinetd,d/

Explanation:

Question: 200

What command is used to add OpenSSH private keys to a running ssh-agen . instance? (Specify the command name only without any path.)

Answer: ssh-add

Explanation:

Question: 201

What information related to a user account is modified using the change command?

A. Default ownership for new files
B. Group membership
C. Set of commands available to the user
D. Password expiry information
E. Default permissions for new files

Answer: D

Explanation:

Question: 202

What is Irue regarding public and private SSH keys? (Choose TWO correct answers.)

A. For each user account, there is exactly one key pair that can be used to log into that account.
B. The private key must never be revealed to anyone.
C. Several different public keys may be generated (or the same private key.
D. To maintain the private key's confidentiality, the SSH key pair must be created by its owner.
E. To allow remote logins, the user's private key must be copied to the remote server.

Answer: B, D

Explanation:

Question: 203

Which parameter of the ssh command specifies the location of the private key used for login attempts? (Specify ONLY the option name without any values or parameters.)

Answer: ssh-keygen

Explanation:

Question: 204

After editing the TCP wrapper configuration to grant specific hosts access to a service. When do these changes become effective?

A. The new configuration becomes effective after restarting the respective service.
B. The new configuration becomes effective at the next system reboot.
C. The new configuration becomes effective when the last established connection to the service is closed.
D. The new configuration becomes effective after restarting the tcpd service.
E. The new configuration becomes effective immediately for all new connections.

Answer: E

Explanation:

Question: 205

Which of the following files is not read directly by a Bash login shell?

A. -/.bashrc

B. -/.bash_profile
C. -/.bash_login
D. -/.profile
E. /etc/profile

Answer: A

Explanation:

Question: 206

What is Irue regarding the statement beginning with #. that is found in the first line of a script? (Choose TWO correct answers.)

A. It prevents the script from being executed until the, is removed

B. It triggers the installation of the script's interpreter.
C. It specifies the path and the arguments of the interpreter used to run the script.
D. It defines the character encoding of the script.
E. It is a comment that is ignored by the script interpreter.

Answer: B, C

Explanation:

Question: 207

What output does the command seq 1 5 20 produce?

A. 1 5 10 15 B.
1
6
11
16 C.
1
2
3
4
D.
2
3
4
5
E.
5
10
15
20

Answer: B

Explanation:

Question: 208

If an alias 1s exists, which of the following commands updates the alias to point to the command 1s - 1 instead of the alias's current target?

A. Sot is='ls -1'
B. Alias 1s='s'1a -1'
C. Alias –force 1s='1s -1'
D. Alias –update is is-'1s-1'
E. Realias 1s='1s -1'

Answer: B

Explanation:

Question: 209

What is true about the file .profile in a user's home directory?

A. It must be executable.
B. It must call the binary of the login shell.
C. It must use a valid shell script syntax.
D. It must start with a shebang.
E. It must be readable for its owner only.

Answer: C

Explanation:

Question: 210

What does the term Braille Display refer to?

A. A standardized high contract graphical theme for desktop applications
B. A Linux desktop environment similar to KDE and GNOME.
C. A legacy display technology superseded by LCD.
D. A physical representation of characters using small dots.
E. A standard file format for data exchange, similar to XML

Answer: D

Explanation:

Question: 211

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:

Question: 212

Which environment variable is used by an X11 client to determine the X Server to connect to? (Specify ONLY the variable name without any preceding commands or values.)

Answer: Display

Explanation:

Question: 213

Which file is processed by newaliases? (Specify the full name of the file, including path.)

Answer: /etc/mail/alia

Explanation:

Question: 214

Where is the system journal stored?

A. /var/ jlog/ and /var/jlogd/
B. /proc/log and /proc/klog
C. /run/log/journal/or/var/log/journal/
D. /var/log/syslog.bin or /var/log/syslog-jrn
E. /etc/system/journal / or /usr/1ib/sysLend/journal/

Answer: C

Explanation:

Question: 215

Which of the following options in the chrony configuration file define remote lime sources? (Choose TWO correct answers.)

A. Source
B. Clock
C. Remote
D. Pool
E. server

Answer: D, E

Explanation:

Question: 216

Which of the following commands display a list of jobs in the print queue? (Choose TWO correct answers.)

A. Cups—list
B. Iprm -1
C. Lpstat
D. Ipr –q
E. ipq

Answer: C, E

Explanation:

Question: 217

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 Howday
C. Syslend-cst echo Howdy
D. Echo Howdy > /dev/journal
E. Journalct1 and howdy

Answer: C, E

Explanation:

Question: 218

Which of the following commands display the number of bytes transmitted and received via the etho network interface? (Choose TWO correct answer.)

A. Route –v via etho

B. Ip stats show dev etho
C. Netstat –s –I etho
D. Ifconfig etho
E. Ip –s link show etho

Answer: DE

Explanation:

Question: 219

What command enables a network interface according to distribution-specific configuration, such as /etc/network/interfaces or /etc/sysconflg/network-scripts/ifcfg-etho?(Specify ONLY the command without any path or parameters.)

Answer: UP

Explanation:

Question: 220

How does the ping command work by default?

A. It sends an ICMP Echo RequesI lo a remote host and warts to receive an ICMP Echo Response in return.
B. It sends an ARP request to a remote host and waits to receive an ARP response in return.
C. It sends a TCP SYN packet to a remote host and waits to receive an TCPACK response in return.
D. It sends a broadcast packet to all hosts on the net and waits to receive, among others, a response from the target system.
E. It sends a UDP packet to port 0 of the remote host and waits to receive a UDP error response in return.

Answer: A

Explanation:

Question: 221

Which of the following are valid host addresses for the subnet 203.0.113.64/28? (Choose TWO correct answers.)

A. 2030.113.64
B. 2030.113.78
C. 203.0.113.65
D. 203.0.113.80
E. 203.0.113.81

Answer: B, C

Explanation:

Question: 222

Which of Ihe following commands will delete Ihe defaull gateway from the system's IP touting table? (Choose TWO correct answers.)

A. Ifconfig unaet default
B. Route del default
C. Ip route del default
D. Netstat –r default
E. Sysct1 ipv1default_gw=0

Answer: B, C

Explanation:

Question: 223

What is the about the following command?
Nmcli device wifi connect WIFIo1

A. NetworkManager opens a new public hotspol with ihe SSID...
B. NetworkManager creates an unconfigured new virtual network interface named WlFIc1.
C. NetworkManager creates a new wifi connection MZTZol and activates it.
D. NetworkManager returns an error in case the connection WIFIIol does not exist.
E. NetworkManager reports an error because WiFIo1 is an invalid wifi device.

Answer: C

Explanation:

Question: 224

Which parameter is missing in the command
Ip link act dec etho
To activate the previously inactive network interface ethO? (Specrfy the parameter only without any command, path or additional options)

Answer: Up

Explanation:

Question: 225

Which of Ihe following slates can NetworkManager show regarding the system's network connectivity? (Choose TWO correct answers.)

A. Up
B. Portal
C. full
D. Login-required
E. firewalled

Answer: B, C

Explanation:

Question: 226

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: B, C

Explanation:

Question: 227

Which standardized TCP port is used by HTTPS services?

A. 25
B. 80
C. 8080
D. 443
E. 636

Answer: D

Explanation:

Question: 228

Which of Ihe following commands sets the system's time zone to the Canadian Eastern Time?

A)

lo ..‘ i. ■ -r -f / H r / nb.i HI//< ■ i ' /;'. I I.ATHTI > /«t.^/1 GCjll^. T ;■

B)

/etc/localtime

C)

$y»zl. ■< cXcck.’.- ’Canada/Easlei

D) xowroce zz :a est

E)

1 —s ' /L.s-/«r-ore/-ZQ"ei ••'c/7e"udi:/Fa«.e"n /^ .c/lCK.-i/Une

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E

Answer: E

Explanation:

Question: 229

Which of the following sections exists in a systemd timer unit?

A. [Events]
B. [Timer]
C. [Cron]
D. [Schedule]
E. [Trigger]

Answer: B

Explanation:

Question: 230

If neigher cron, allow nor cron, deny exist in /etc/, which of the following is true?

A. Without additional configuration, all users may create user specific crontabs.
B. Without additional configuration, only root may create user specific crontabs.
C. The aon 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.
E. The default settings of /etc/crond.conf define whether or not user specific crontabs are generally allowed or not.

Answer: E

Explanation:

Question: 231

How can a specific user be prevented from scheduling tasks with a i.

A. By adding the specific user to the /etc/ at,allow file.

B. By adding the specific user to the [deny] section in the /etc./atd.conf file.
C. By adding the specific user to the nojoba group.
D. By adding the specific user to the /etc/at.deny file.
E. By executing the atd –deny [user] command.

Answer: D

Explanation:

Question: 232

Which of the following getent invocations lists all existing users?

A. getent homco
B. getent uids
C. getent pmsswd
D. getent users
E. getent logins

Answer: C

Explanation:

Question: 233

Given the following user's cronlab entry:
15 14 * * 1-5 /usr/local/bin/example.sh be executed?
When will the script /usr/local/bin/example.sh be executed?

A. At 14:15 local time. January till May.
B. At 15:14 local time, 1st to 5th day of month.
C. At 14:15 local time, February till June.
D. At 14:15 local time, 1st to 5th day of month.
E. At 14:15 local time, Monday to Friday.

Answer: E

Explanation:

Question: 234

Which of the following environment variables can be defined in locale. conf ? (Choose TWO correct answers.)

LC_ALL

A. LC_USERNAME
B. LC_U?F8
C. LC_GEOGRAPHY
D. LC_TIME

Answer: A, E

Explanation: