

Oracle

Exam 1z0-460

Oracle Linux 6 Implementation Essentials

Version: Demo

[Total Questions: 10]

Question No : 1

The DBA tells you that the system is not overloaded but you can tell that the system is actively swapping. What command would you run to show this information to the DBA?

- A. # iotop
- B. # iostat 5 10
- C. # cat /proc/meminfo
- D. # vmstat 5 10

Answer: B

Explanation: *iostat - Report Central Processing Unit (CPU) statistics and input/output statistics for devices, partitions and network filesystems (NFS).

*The iostat command is used for monitoring system input/output device loading by observing the time the devices are active in relation to their average transfer rates. The iostat command generates reports that can be used to change system configuration to better balance the input/output load between physical disks.

Incorrect:

Not A: Related to kernel and processes.

*iotop - simple top-like I/O monitor

*iotop watches I/O usage information output by the Linux kernel (requires 2.6.20 or later) and displays a table of current I/O usage by processes or threads on the system.

*iotop displays columns for the I/O bandwidth read and written by each process/thread during the sampling period. It also displays the percentage of time the thread/process spent while swapping in and while waiting on I/O. For each process, its I/O priority (class/level) is shown. In addition, the total I/O bandwidth read and written during the sampling period is displayed at the top of the interface.

Not C: related to RAM usage.

*The entries in the /proc/meminfo can help explain what's going on with your memory usage, if you know how to read it.

*High-Level Statistics

MemTotal: Total usable ram (i.e. physical ram minus a few reserved bits and the kernel binary code)

MemFree: Is sum of LowFree+HighFree (overall stat)

MemShared: 0; is here for compat reasons but always zero.

Buffers: Memory in buffer cache. mostly useless as metric nowadays

Cached: Memory in the pagecache (diskcache) minus SwapCache

SwapCache: Memory that once was swapped out, is swapped back in but still also is in the swapfile (if memory is needed it doesn't need to be swapped out AGAIN because it is already in the swapfile. This saves I/O)

Not D:vmstat - Report virtual memory statistics

Question No : 2

A system administrator wants to view all running processes on the system in real time, to find out what RAM has been allocating to each process. What system command should be used?

- A. ps -ef
- B. ps ax
- C. top
- D. meminfo

Answer: C

Explanation: top - display Linux tasks

The top program provides a dynamic real-time view of a running system. It displays system summary information as well as a list of tasks currently being managed by the Linux kernel. The types of system summary information shown and the types, order and size of information displayed for tasks are all user configurable and that configuration can be made persistent across restarts.

Incorrect:

Not A, Not B:

ps displays information about a selection of the active processes.

To see every process on the system using standard syntax:

ps -e

ps -ef

```
ps -eF  
ps -ely
```

To see every process on the system using BSD syntax:

```
ps ax  
ps axu
```

Not D:meminfo - provide information about memory

The meminfo() function provides information about virtual and physical memory particular to the calling process. The user or developer of performance utilities can use this information to analyze system memory allocations and develop a better understanding of the factors affecting application performance.

Question No : 3

Your Oracle Linux system has two network interfaces – eth0 and eth1. You have to change the netmask and the IP address of the eth1 network interface. Which configuration file would you edit to make these changes?

- A. /proc/net/ifcg-eth1
- B. /etc/sysconfig/network-scripts/eth1
- C. /etc/sysconfig/network/ifcg-eth1
- D. /etc/sysconfig/network-scripts/ifcg-eth1

Answer: D

Explanation: The "/etc/sysconfig/network-scripts/ifcfg-eth0" file holds the network configuration for the "eth0" adapter. If you have multiple network adapters, you would expect additional configuration files (eth1, eth2 etc.).

Note:

*The "/etc/sysconfig/network" file holds top-level networking configuration, including the hostname and gateway settings.

Question No : 4

View the output below.

As a root user, you run the two ulimit commands as shown in the output below. Why does the second ulimit command fail as shown in the output?

```
[root@dbhost ~]# ulimit -H -n
4096
[root@dbhost ~]#
[root@dbhost ~]# ulimit -S -n 4099
bash: ulimit: open files: cannot modify limit: Invalid argument
[root@dbhost ~]#
```

- A. The ulimit command cannot be run from the bash shell.
- B. The ulimit command syntax is not correct.
- C. The soft limit value of file description cannot be set greater than the hard limit value.
- D. The soft limit of file descriptions' value should always be less than 1024.

Answer: C

Explanation: Any user can set a soft limit to any value less than or equal to the hard limit. Any user can lower a hard limit. Only a user with appropriate privileges can raise or remove a hard limit.

Note:

*limit, ulimit, unlimit– set or get limitations on the system resources available to the current shell and its descendents

* Syntax:

ulimit [- [HS] [c | d | f | n | s | t | v]] limit

-H

Displays or sets a hard limit.

-S

Displays or sets a soft limit.

Reference: man ulimit

Question No : 5

You have to aggregate two network interfaces, eth0 and eth1, into a single logical interface such as bond0. Which option shows the four configuration files that need to be configured to set up this bonding?

- A. /etc/sysconfig/network-scripts/ifcfg-bond0
/etc/sysconfig/network-scripts/ifcfg-eth0
/etc/sysconfig/network-scripts/idfg-eth1
/proc/bonding.conf
- B. /etc/sysconfig/network-scripts/ifcfg-bond0
/etc/sysconfig/network-scripts/ifcfg-eth0
/etc/sysconfig/network-scripts/idfg-eth1
/etc/modprobe.d/bonding.cfg
- C. /etc/sysconfig/network/ifcfg-bond0
/etc/sysconfig/network-scripts/ifcfg-eth0
/etc/sysconfig/network-scripts/ifcfg-eth1
/etc/modprobe.d/bonding.conf
- D. /etc/sysconfig/network-scripts/ifcfg-bond0
/etc/sysconfig/network-scripts/eth0
/etc/sysconfig/network-scripts/eth1
/etc/bonding.conf

Answer: C

Explanation: *Step #1: Create a Bond0 Configuration File

Red Hat Enterprise Linux (and its clone such as CentOS) stores network configuration in /etc/sysconfig/network-scripts/ directory. First, you need to create a bond0 config file as follows:

```
# vi /etc/sysconfig/network-scripts/ifcfg-bond0
```

*Step #2: Modify eth0 and eth1 config files

Open both configuration using a text editor such as vi/vim, and make sure file read as follows for eth0 interface# vi /etc/sysconfig/network-scripts/ifcfg-eth0

*Step # 3: Load bond driver/module

Make sure bonding module is loaded when the channel-bonding interface (bond0) is brought up. You need to modify kernel modules configuration file:

For each configured channel bonding interface, there must be a corresponding entry in your new **/etc/modprobe.d/bonding.conf** file.

Question No : 6

The sshd service running and you execute the following command:

```
# chkconfig sshd off
```

What happens when you run this chkconfig command?

- A. The sshd service disabled only for runlevel 5.
- B. The sshd service id disabled for runlevels 2, 3, 4, and 5, but the ssh service is still available until the next reboot.
- C. The sshd service is disabled for runlevels 2, 3, 4, and 5 and ssh service is stopped.
- D. The sshd service is disabled only for current runlevel.

Answer: B

Question No : 7

DTrace is being ported from Solaris to Oracle Linux. Which three statements are true for the DTrace tool?

- A. DTrace allows static and dynamic tracing of your applications and your kernel.
- B. DTrace tool is used to compile debug kernel modules and device drivers
- C. DTrace allows you to dynamically define probe points on the fly.
- D. DTrace probes and probe points are usually defined by the user using scripts written in a language called D.
- E. DTrace tool is based on the strace Linux tool and includes both user and kernel strace features.

Answer: A,C,D

Explanation: A:DTrace is a comprehensive dynamic tracing framework created by Sun Microsystems for troubleshooting kernel and application problems on production systems in real time.

C:Key benefits and features of DTrace on Oracle Linux include:

/Designed to work on finding performance bottlenecks

/(C)Dynamically enables the kernel with a number of probe points, improving ability to service software

/Enables maximum resource utilization and application performance

/Fast and easy to use, even on complex systems with multiple layers of software

D:Testers write tracing programs (also referred to as scripts) using the D programming language (not to be confused with other programming languages named "D"). The language, a subset of C, includes added functions and variables specific to tracing. D programs resemble awk programs in structure; they consist of a list of one or more probes (instrumentation points), and each probe is associated with an action. These probes are comparable to a pointcut in aspect-oriented programming.

Question No : 8

Which two statements describe the capabilities used with the Unbreakable Enterprise Kernel?

- A. Existing Red Hat Enterprise Linux 5 and 6 customers need to reinstall Oracle Linux to use the Unbreakable Enterprise Kernel.
- B. The Unbreakable Enterprise kernel is the default kernel starting with Oracle Linux 5.6.
- C. The Unbreakable Enterprise kernel is required when using multithreaded CPUs.
- D. Oracle Clusterware, OCFS2, and the Enterprise Manager pack for Linux support are included with Oracle Linux Basic and Premier support.
- E. Switching between the Red Hat Compatible kernel and the Unbreakable Enterprise kernel is simple process of changing kernels and glibc.

Answer: D,E

Explanation: *Commercial technical support is available through Oracle's Oracle Linux Support program, which supports Oracle Linux, and existing RHEL or CentOS installations(i.e. without reinstallation).

Note:

*The Unbreakable Enterprise Kernel Release 2 is Oracle's second major release of its heavily tested and optimized operating system kernel for Oracle Linux 5 and Oracle Linux 6.

Unbreakable Enterprise Kernel Release 2 is based on the mainline Linux kernel version 3.0.16 and boasts a wide range of new features and improvements relevant for enterprise workloads.

Incorrect:

Not A, not B:Unbreakable Enterprise Kernel Release 2 can be installed on Oracle Linux 5 Update 8 or newer, as well as on Oracle Linux 6 Update 2 or newer.

Question No : 9

You want to allow multiple users the write access to files within the same directory, in addition, you want all the new files created in this directory to be of the required group instead of the primary ID of the user who creates the file. How do you accomplish this?

- A. Set the setgid bit on the directory.
- B. Change the group owner of the new files manually.
- C. Run a cron job to change the group owner.
- D. Change the primary group ID of every user to the required group.

Answer: A

Explanation: Linux: SETGID on directory

SETGID stands for SET Group ID. We can use the command `chmod` to set the group ID bit for a directory.

```
chmod g+s mydir
```

or with numeric mode:

```
chmod 2775 mydir
```

After the change, the permission of the directory "mydir" becomes "drwxrwsr-x".

```
drwxrwsr-x 3 zen zen 4096 2010-03-18 19:57 mydir
```

But what is so special about setting the group ID for a directory? The trick is that when another user creates a file or directory under such a directory "mydir", the new file or directory will have its group set as the group of the owner of "mydir", instead of the group of the user who creates it.

For example, if user2 belongs to the groups "user2" (main group) and "zen", and he creates a file "newfile" under the directory "mydir", "newfile" will be owned by the group of "zen" instead of user2's main group ID "user2".

Question No : 10

The ASMLib on Linux us designed to _____.

- A. Make managing ASM volume easier.
- B. Make managing ASM volumes easier and provide a performance improvement over ASM volumes alone.
- C. Provide load balancing across multiple volumes.
- D. Allow the user to list contents of the volume from the OS command line.

Answer: B

Explanation: ASMLib is an optional support library for the Automatic Storage Management feature of the Oracle Database.

ASMLib allows an Oracle Database using ASM more efficient and capable access to the disk groups it is using.