

# Oracle

## Exam 1z0-460

### Oracle Linux 6 Implementation Essentials

Version: 6.0

[ Total Questions: 71 ]

**Question No : 1**

Finding a directive issued by your organization listing several important CVEs (Common Vulnerability and Exposures), you find one for Linux OpenSSH (CVE-2006-5764), which might apply to your oracle Linux systems. What command would help ensure that a patch has been applied to close this vulnerability on an Oracle Linux 6 system running OpenSSH server?

- A. yum listcves openssh
- B. rpm -qa | grep openssh | grep 5794
- C. rpm -q - - changelog openssh | grep 5794
- D. yum sec - list cves | grep 5794
- E. yum sec - list cvesApplied

**Answer: C**

**Explanation:** \*The command rpm -q --changelog rpm displays a detailed list of information (updates, configuration, modifications, etc.) about a specific package. This example shows information about the package rpm. However, only the last five change entries in the RPM database are listed. All entries (dating back the last two years) are included in the package itself. This query only works if CD 1 is mounted at /media/cdrom:

```
rpm -qp --changelog /media/cdrom/suse/i586/rpm-3*.rpm
```

\*Is the patch RPM suitable for my system?

To check this, first query the installed version of the package. For pine, this can be done with

```
rpm -q pine  
pine-4.44-188
```

**Question No : 2**

Which three parameters of a network interface can you modify by using the NetworkManager tool on your Oracle Linux 6 system?

- A. IPv4 settings
- B. IPv6 settings

- C. Netconsole settings
- D. MTU settings
- E. IP Proxy settings

**Answer: A,B,D**

**Explanation:**



Connection name: System eth0

Connect automatically

Wired 802.1x Security IPv4 Settings IPv6 Settings

MAC address: 52:54:00:26:9E:F1

MTU: automatic bytes

Available to all users

Cancel Apply...

[http://dsilas.fedorapeople.org/deployment-guide/html/images/Network\\_Configuration-NM-Edit\\_System\\_eth0-Wired.png](http://dsilas.fedorapeople.org/deployment-guide/html/images/Network_Configuration-NM-Edit_System_eth0-Wired.png)

Note:

1 Right-click the NetworkManager icon in the notification area at the top-right corner of the Red Hat desktop and click "Edit Connections."

2

Click the "System eth0" connection on the wired tab and click "Edit."

3

Click the "IPv4 Settings" tab.

**Question No : 3**

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 : 4**

View the exhibits.

```
[root@dbhost ~]# sar -u 2 5
Linux 2.6.39-100.5.1.el6uek.x86_64
x86_64 (1 CPU)
04:43:21 PM   CPU   %user   %nice   %system   %iowait   %steal   %idle
04:43:23 PM   all    36.82    0.00    50.25    0.00    0.00    12.94
04:43:25 PM   all    34.67    0.00    50.25    0.00    0.00    15.08
04:43:27 PM   all    33.50    0.00    52.00    0.00    0.00    14.50
04:43:29 PM   all    32.16    0.00    52.26    0.00    0.00    15.58
04:43:31 PM   all    31.16    0.00    50.25    0.00    0.00    18.59
Average:      all    33.67    0.00    51.00    0.00    0.00    15.33
[root@dbhost ~]#
```

```
top - 16:39:07 up 1 day, 2:34, 2 users, load average: 1.41, 0.64, 0.33
Tasks: 140 total, 2 running, 138 sleeping, 0 stopped, 0 zombie
Cpu(s): 36.8%us, 50.7%sy, 0.0%ni, 13.3%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 2056568k total, 1635148k used, 421420k free, 41452k buffers
Swap: 4128764k total, 0k used, 4128764k free, 567176k cached

  PID USER   PR   NI  VIRT  RES  SHR  S  %CPU  MEM%   TIME+  COMMAND
 5125 root    20    0 1886m 575m 22m  R  73.4  28.6   94:17.31 npviewer.bin
 1942 root    20    0  150m  45m 11m  S   6.0   2.2    2:48.36 Xorg
 2326 root    9   -11  495m 5044 3540  S   3.0   0.2   37:44.45 pulseaudio
 5059 root    20    0  879m 171m 29m  S   2.7   8.5   10:27.23 firefox
 5113 root    20    0  318m  14m 8368  S   2.0   0.7    7:06.51 plugin-containe
    1 root    20    0 19376 1508 1212  S   0.0   0.1    0:00.62 init
    2 root    20    0    0    0    0  S   0.0   0.0    0:00.00 kthreadd
    3 root    20    0    0    0    0  S   0.0   0.0    0:03.91 ksoftirqd/0
    5 root    20    0    0    0    0  S   0.0   0.0    0:00.00 kworker/u:0
    6 root    RT    0    0    0    0  S   0.0   0.0    0:00.00 migration/0
    7 root    RT    0    0    0    0  S   0.0   0.0    0:01.25 watchdog/0
    8 root    0  -20    0    0    0  S   0.0   0.0    0:00.00 cpuset
```

Examine the output of sar command and the top command in the Exhibits. Which statement is the correct interpretation of this data?

- A. The system is running low on swap space and memory.
- B. CPU is Idle and the system has plenty of free memory available.
- C. The CPU utilization is high and one process is using most of the CPU.
- D. The system is idle with very little memory, CPU, and I/O utilization.

**Answer: C**

**Explanation:** From the top exhibit we see that npviewer.gin uses 73.4% of the available CPU.

Note:

\*sar - Collect, report, or save system activity information.

\*The sar command writes to standard output the contents of selected cumulative activity counters in the operating system. The accounting system, based on the values in the count and interval parameters, writes information the specified number of times spaced at the specified intervals in seconds.

\*sar -u 2 5

Report CPU utilization for each 2 seconds. 5 lines are displayed.

### Question No : 5

Which three steps are involved in the installing Ksplice on servers that will be updated?

- A. You have to be logged as user “ksplice” on the server you want to prepare and install Oracle Ksplice on.
- B. The Uptrack package needs access directly or through a proxy to the Oracle public yum repository (<http://public-yum.oracle.com/>) to download the required packages for the uptrack-\* utilities to be able to work correctly.
- C. Download the install-uptrack script using the “wget -N <http://www.ksplice.com/uptrack/install-uptrack>” command.
- D. You have to be logged in as user “root” on the server you want to prepare and install Oracle ksplice on.
- E. The uptrack package will set up a yum repository (/etc/yum.repos.d/ksplice-uptrack.repo) and download the required package for the uptrack-\* utilities to be able to work correctly.
- F. Download the ksplice ISO image from <https://edelivery.oracle.com/linux> and then run the “sh install-uptrack” script from the ISO image.

**Answer: B,C,D**

**Explanation:** B:Your system must have access to the internet to install Ksplice. If you are using a proxy, set the proxy in your shell:

```
export http_proxy=http://proxy.company.com:port
export https_proxy=http://proxy.company.com:port
```

CD:Once you have an access key run the following commands as root, replacing YOUR\_ACCESS\_KEY with the access key you received upon sign-up:

```
wget -N https://www.ksplice.com/uptrack/install-uptrack
sh install-uptrack YOUR_ACCESS_KEY
```

uptrack-upgrade-y

E:If you'd like Ksplice Utrack to automatically install updates as they become available, run:

```
sh install-uptrack YOUR_ACCESS_KEY --autoinstall
```

in place of the above install-uptrack command, or set "autoinstall = yes" in your /etc/uptrack/uptrack.conf after installation.

Reference: Ksplice,Installation instructions

### Question No : 6

What happens when the following command is run?

```
# authconfig - - passalgo = md5 - - update
```

- A. It produces the MD5 checksum of the input data.
- B. It configures the MD5 checksum for newly authored documents
- C. It converts the stdio input to MD5 algorithm.
- D. It changes the user password hashing algorithm to MD5.

**Answer: D**

**Explanation:** To configure the Linux system to use theMD5algorithm, enter:

```
# authconfig --passalgo=MD5--update
```

Note: The default algorithm for storing password hashes in /etc/shadow is MD5. I was told to use SHA-512 hashing algorithm. How do I set password hashing using the SHA-256 and SHA-512 under CentOS or Redhat Enterprise Linux 5.4?

You need to use authconfig command to setup SHA-256/512 hashing. This command provides a simple method of configuring /etc/sysconfig/network to handle NIS, as well as /etc/passwd and /etc/shadow, the files used for shadow password support. Basic LDAP, Kerberos 5, and SMB (authentication) client configuration is also provided.

## Display Current Hashing Algorithm

Type the following command:

```
# authconfig --test | grep hashing
```

Sample outputs:

```
password hashing algorithm is md5
```

## Configure Linux Server To Use The SHA-512

To configure the Linux system to use the SHA-512 algorithm, enter:

```
# authconfig --passalgo=sha512 --update
```

**Question No : 7**

As user bob, you have logged in to the system on a terminal and issued the following command to make the top command run in the background.

```
[bob@host - ] top&
```

You exit from the terminal and log back as bob into the system at the same terminal. How is the background job affected?

- A. The background job gets the foreground as soon as bob logs into the system.
- B. The background job is suspended temporarily from the job and resumes when user bob logs back in to the system.
- C. The background job starts running in the background again as soon as bob logs into the system.
- D. The background job will not be affected.
- E. The background job is deleted from the job pool and does not get listed using jobs command.

**Answer: E**

**Question No : 8**



Which rpm command can be used to find the package that owns the /etc/rsyslog.conf file?

- A. rpm -query /etc/rsyslog.conf
- B. rpm -gf /etc/rsyslog.conf
- C. rpm -q1 /etc/rsyslog.conf
- D. rpm -q /etc/rsyslog.conf

**Answer: B**

**Explanation:** Package Selection Options include:

-f <file>

Query package owning <file>

Note:

rpm -q — What does it do?

One of the nice things about using RPM is that the packages you manage don't end up going into some kind of black hole. Nothing would be worse than to install, upgrade, and erase several different packages and not have a clue as to what's on your system. In fact, RPM's query function can help you get out of sticky situations like:

\*You're poking around your system, and you come across a file that you just can't identify. Where did it come from?

\*Your friend sends you a package file, and you have no idea what the package does, what it installs, or where it originally came from.

\*You know that you installed XFree86 a couple months ago, but you don't know what version, and you can't find any documentation on it.

The list could go on, but you get the idea. The rpm -q command is what you need. If you're the kind of person that doesn't like to have more options than you know what to do with, rpm -q might look imposing. But fear not. Once you have a handle on the basic structure of an RPM query, it'll be a piece of cake.

On your Oracle Linux 6 system, you have to configure the eth0 network interface to 100 MB/sec, half duplex without trying to autonegotiate. Which command will help you configure this requirement?

- A. # ifconfig eth0 speed 100 autoneg off duplex half
- B. # ethtool interface eth0 speed 100 autoneg off duplex half
- C. # ifconfig interface eth0 speed 100 autoneg off duplex half
- D. # ethtool -seth0speed 100 autoneg off duplex half

**Answer: D**

**Explanation:** When I have a device that is acting up, I tend to run: ethtool ethX, check the Supported link modes, the Link partner advertised link modes and the actual speed and Duplex. If my Supported link mode is set low (say 10/Half for some reason) but my switch supports 1000baseT/Full then I'll use ethtool -s ethX to change my ethernet settings to 1000baseT/Full. Just about anything you see from: ethtool ethX, can be changed with ethtool -S ethX. In this case you would use the following:

```
ethtool -S eth0 speed 1000 duplex full autoneg on
```

### Question No : 10

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: