



1Z0-460^{Q&As}

Oracle Linux 6 Implementation Essentials

Pass Oracle 1Z0-460 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/1z0-460.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

Which two statements describe Oracle ksplice?

- A. Ksplice actively applies kernel errata updates to the on-disk image and after a subsequent reboot changes take effect.
- B. Ksplice can be used to update Oracle Database software installed on a running Oracle Linux system.
- C. Ksplice actively applies kernel errata updates to the running Oracle Linux kernel image.
- D. Ksplice only works with the Red Hat Compatible Kernel.
- E. Ksplice allows customers to remain current with their OS vulnerability patches while at the same time minimizing downtime.

Correct Answer: CE

Ksplice is a free software extension of the Linux kernel which allows system administrators to apply security patches to a running kernel without having to reboot the operating system (C, not A) (a technique broadly referred to as Dynamic Software Updating). It only supports patches that don't make significant semantic changes to kernel data structures. Ksplice has been implemented for Linux on the x86-32 and x86-64 architectures. It was developed by Ksplice, Inc. until 21 July 2011, when Oracle acquired Ksplice and started offering support for Oracle Linux (not D). Support for Red Hat Enterprise Linux was dropped and turned into a free 30-day trial for RHEL customers as an incentive to migrate to Oracle Linux Premier Support.

QUESTION 2

Identify the two kernels that are shipped with Oracle Linux 6.

- A. Unbreakable Enterprise Kernel
- B. Unbeatable Enterprise Kernel 11g
- C. Red Hat Compatible Kernel
- D. Linux Compatible Kernel
- E. Solaris 11 Container Kernel

Correct Answer: AC

Oracle Linux 6 ships with two sets of kernel packages:

*

Unbreakable Enterprise Kernel [kernel-uek-2.6.32-100.28.5.el6] Only available on the x86_64 (64 bit) platform Installed and booted by default

*

Red Hat compatible Kernel [kernel-2.6.32-71.el6] Installed by default



Reference: Oracle Linux 6 Release Notes

QUESTION 3

Which two features are available with the Unbreakable kernel R2, but not with the Red Hat Compatible Kernel?

- A. Oracle Clusterware for Linux
- B. Up to 4-petabyte cluster volumes with OCFS2
- C. Ksplice zero downtime patching
- D. Transparent Huge Pages support (that is, 2 MB instead of 4 KB)

Correct Answer: AB

Oracle's Unbreakable Enterprise Kernel KEY FEATURES:

*

Modern kernel based on 2.6.32, optimized by Oracle for server deployments

*

Includes OCFS (Oracle Cluster File System) 2.1.6 for clustered volumes

*

Includes OFED 1.5.1

*

Advanced NUMA support

*

New diagnostic and tracing tools, including performance counters

*

Complete data integrity checking from application to disk

*

Hardware fault management

QUESTION 4

Examine the sar command below. Your Oracle Linux system has one CPU. What does the runq- sz column of this output convey about your system?



```
[root@dbhost ~]# sar -q 2 3
Linux 2.6.39-100.5.1.el6uek.x86_64 (dbhost.example.com)      _x86_64_
(1 CPU)

11:49:13 AM      runq-sz      plist-sz      ldavg-1      ldavg-5      ldavg-15
11:49:15 AM              3          324          3.03          1.45          0.61
11:49:17 AM              3          324          3.03          1.45          0.61
11:49:19 AM              3          324          3.03          1.45          0.61
Average:              3          324          3.03          1.45          0.61
[root@dbhost ~]#
```

- A. CPU is bottleneck because the run queue size is greater than the number of CPUs on your system.
- B. The average of three processes are only using the CPU on your system and hence the CPU is not a bottleneck.
- C. CPU is not a bottleneck because the run queue size indicates the number of CPU bound processes on your system.
- D. CPU is a bottleneck because the run queue size indicates that adequate memory is not allocated.

Correct Answer: A

*

Use the sar -q command to report the following information:

The Average queue length while the queue is occupied.

The percentage of time that the queue is occupied.

*

The following list describes the output from the -q option.

runq-sz The number of kernel threads in memory that are waiting for a CPU to run. Typically, this value should be less than 2. Consistently higher values mean that the system might be CPU-bound.

%runocc

The percentage of time that the dispatch queues are occupied.

swpq-sz

Swap queue of processes for the sar command.

%swpocc

Swap queue of processes for the sar command.

QUESTION 5

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

Correct Answer: B

*

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 6



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.

Correct Answer: ACD

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 7

View the exhibit.



```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE: You have a /boot partition. This means that
#           all kernel and initrd paths are relative to /boot/, eg.
#           root (hd0,0)
#           kernel /vmlinuz-version ro root=/dev/mapper/vg_dbhost-lv_root
#           initrd /initrd-[generic-]version.img
#boot=/dev/sda
default=0
timeout=8
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Oracle Linux Server (2.6.39-100.5.1.el6uek.x86_64)
    root (hd0,0)
    kernel /vmlinuz-2.6.39-100.5.1.el6uek.x86_64 root=/dev/mapper/vg_dbhost-
lv_root ro rd_NO_LUKS LANG=en_US.UTF-8 rd_NO_MD quiet SYSFONT=latarcyrheb-sun16
rd_LVM_LV=vg_dbhost/lv_swap rhgb KEYBOARDTYPE=pc KEYTABLE=us rd_LVM_LV=vg_dbhost
/lv_root rd_NO_DM numa=off
    initrd /initramfs-2.6.39-100.5.1.el6uek.x86_64.img
title Oracle Linux Server (2.6.32-300.11.1.el6uek.x86_64)
    root (hd0,0)
    kernel /vmlinuz-2.6.32-300.11.1.el6uek.x86_64 ro root=/dev/mapper/vg_dbh
ost-lv_root rd_NO_LUKS LANG=en_US.UTF-8 rd_NO_MD quiet SYSFONT=latarcyrheb-sun16
rd_LVM_LV=vg_dbhost/lv_swap rhgb KEYBOARDTYPE=pc KEYTABLE=us rd_LVM_LV=vg_dbh
ost-lv_root rd_NO_DM numa=off
```

Examine the grub.conf file snippet in the Exhibit. Which statement is true if your Linux system boots by using this grub.conf file?

- A. GRUB will boot, by default, the first kernel entry of this grub.conf file.
- B. GRUB will prompt you to select the kernel to be booted because the default parameter is set to 0.
- C. GRUB will boot the kernel specified in the inittab file of the system.
- D. GRUB will boot, by default, the second kernel entry of this grub.conf file.

Correct Answer: A

According the grub .conf file you have got 8 seconds to choose whether to boot the first entry.

Now if you want to change, and let say you want the second grub entry as booting system by default, just change the line: default=0 by default=1

QUESTION 8

Examine the output of the yum.conf configuration file from o Oracle Linux 6 server.



```
# cat /etc/yum.conf
[main]
cachedir=/var/cache/yum/$basearch/$releasever
keepcache=0
debuglevel=2
logfile=/var/log/yum.log
exactarch=1
obsoletes=1
gpgcheck=1
plugins=1
installonly_limit=3
```

What does the "installonly_limit=3" option mean?

- A. Only three users can use the yum command simultaneously to install package on Oracle Linux 6 server.
- B. Only three Linux packages can be installed in one invocation of the yum command.
- C. Only three repository channels can be used for installing Linux packages.
- D. Only three versions of packages listed in installonlypkgs can be installed simultaneously.

Correct Answer: D

installonly_limit

Number of packages listed in installonlypkgs to keep installed at the same time. Setting to 0 disables this feature. Default is '3'. Note that this functionality used to be in the "installonlyn" plugin, where this option was altered via. tokeep. Note that as of version 3.2.24, yum will now look in the yumdb for a installonly attribute on installed packages. If that attribute is "keep", then they will never be removed.

Reference: yum.conf

QUESTION 9

The crond daemon checks each command to see whether it should be run in the current _____.

- A. Second
- B. Minute
- C. Hour
- D. Day

Correct Answer: B

The cron utility then wakes up every minute, examining all stored crontabs, checking each command to see if it should be run in the current minute.



QUESTION 10

Which option determines whether a system uses NIS, local files, DNS, or a combination as the source of information, and also order of the source?

- A. /etc/resolv.conf
- B. /etc/idap.conf
- C. /etc/nsswitch.conf
- D. /etc/yp.conf

Correct Answer: C

The Name Service Switch (NSS) configuration file, /etc/nsswitch.conf, is used by the GNU C Library to determine the sources from which to obtain name-service information in a range of categories, and in what order. Each category of information is identified by a database name. The file is plain ASCII text, with columns separated by spaces or tab characters. The first column specifies the database name. The remaining columns describe the order of sources to query and a limited set of actions that can be performed by lookup result.

Incorrect: Not A: In most Unix-like operating systems and others that implement the BIND Domain Name System (DNS) resolver library, the resolv.conf configuration file contains information that

Not D: /etc/yp.conf - NIS binding configuration file

Reference: nsswitch.conf

QUESTION 11

As a system administrator, you run the system-config-network tool and make changes to the configuration. You change the hostname and the DNS search path settings. Which two files will these changes be written into?

- A. "/etc/sysconfig/network" and "/etc/resolv.conf" files
- B. "/etc/sysconfig/network" and "etc/nsswitch.conf/" files
- C. "/etc/sysconfig/netconfig" and "/etc/resolv.conf" files
- D. "etc/sysconfig/network-scripts/network" and "/etc/resolv.conf" files

Correct Answer: C

The system-config-network-tui and system-config-network commands start a text-based network configuration tool.

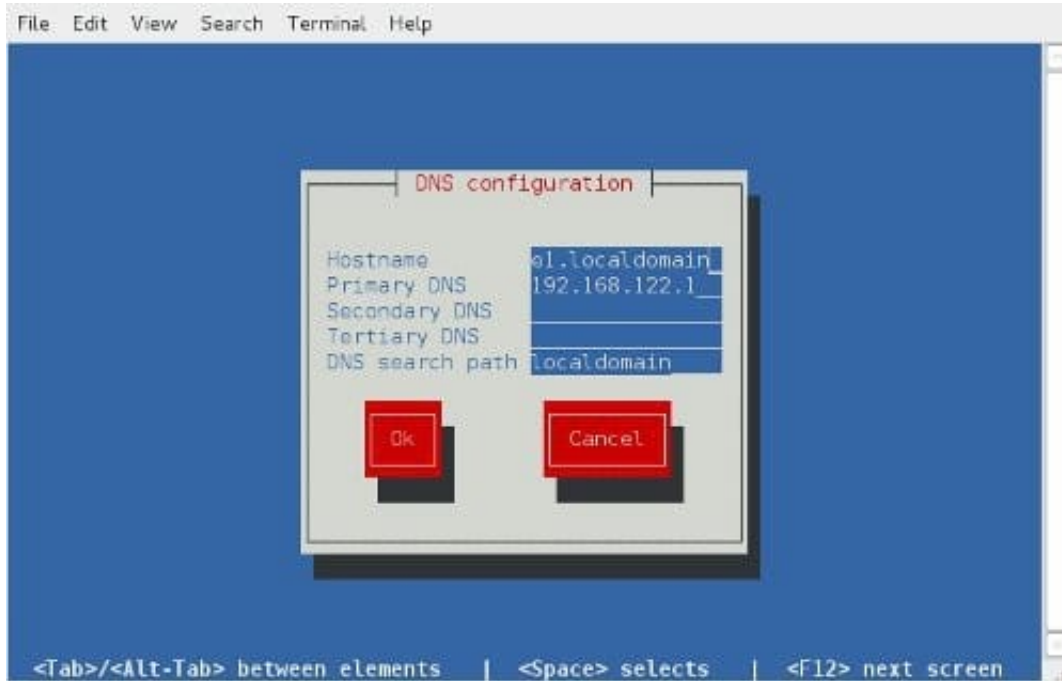
Navigate using the "tab", "arrow" and "return" keys. The "Device configuration" option gives a list of network devices.

Selecting the device allows you to edit the adapter's network configuration, which is saved to the "/etc/



sysconfig/network-scripts/ifcfg-eth0" file.

The "DNS configuration" option on the first screen allows you to modify the configuration in the "/etc/sysconfig/network" and "/etc/resolv.conf" files.



QUESTION 12

You have to find out the version of bash shell package installed on your Oracle Linux system. Which yum command can help you find the version information?

- A. # yum version bash
- B. # yum showversion bash
- C. # yum info bash
- D. # yum listversion bash

Correct Answer: C

Example:

```
[oracle@oraclelinux6 ~]$ yum info bash
```

```
Loaded plugins: refresh-packagekit, security
Installed Packages
Name : bash Arch : x86_64 Version : 4.1.2 Release : 9.el6_2
Size : 3.0 M Repo : installed
From repo : ol6_u3_base
Summary : The GNU Bourne Again shell
URL : http://www.gnu.org/software/bash
License : GPLv3+
Description : The GNU Bourne Again shell (Bash) is a shell or
command language : interpreter that is compatible with the Bourne
shell (sh). Bash : incorporates useful features from the Korn shell
(ksh) and the
```

C : shell (csh). Most sh scripts can be run by bash without : modification.



QUESTION 13

As a root user, you executed the following command on your Oracle Linux 6 server:

```
[root@host] # strace - o /tmp/diag.out sh diag.sh
```

Which statement describes the purpose of this command?

- A. It collects the memory and swap space metrics when the diag.sh scripts runs.
- B. It collects the operating system metrics when the diag.sh script runs.
- C. It records the memory usage and CPU usage information of the processes when the diah.sh script runs.
- D. It records the system calls, which are called by the processes when the diag.sh script runs.

Correct Answer: D

*

strace - trace system calls and signals

In the simplest case strace runs the specified command until it exits. It intercepts and records the system calls which are called by a process and the signals which are received by a process. The name of each system call, its arguments and its return value are printed on standard error or to the file specified with the -o option.

*

strace is a useful diagnostic, instructional, and debugging tool. System administrators, diagnosticians and trouble-shooters will find it invaluable for solving problems with programs for which the source is not readily available since they do not need to be recompiled in order to trace them. Students, hackers and the overly-curious will find that a great deal can be learned about a system and its system calls by tracing even ordinary programs. And programmers will find that since system calls and signals are events that happen

at the user/kernel interface, a close examination of this boundary is very useful for bug isolation, sanity checking and attempting to capture race conditions.

Each line in the trace contains the system call name, followed by its arguments in parentheses and its return value. An example from stracing the command `cat /dev/null` is:

```
open("/dev/null", O_RDONLY) = 3
```

Errors (typically a return value of -1) have the errno symbol and error string appended. `open("/foo/bar",`

`O_RDONLY) = -1 ENOENT (No such file or directory)` Signals are printed as a signal symbol and a signal string. An excerpt from stracing and interrupting the command `sleep 666` is:

```
sigsuspend([]
```

```
--- SIGINT (Interrupt) ---+++ killed by SIGINT +++
```

Reference: `man strace`



QUESTION 14

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/ifcfg-eth1 /proc/bonding.conf
- B. /etc/sysconfig/network-scripts/ifcfg-bond0 /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-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

Correct Answer: C

*

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 15

Examine the following commands: # groupadd project # mkdir /usr/share/project # chown R root.project /usr/share/project # qpasswd a scott project # qpasswd a foo project # chmod 2775 /usr/share/project

Based on the commands, which statement is correct?

- A. Any new file created in the /usr/share/project folder by user scott will not be possible for foo to modify it.



B. All members of the project group need the administrator's help to change the file permission every time users write new files in the /usr/share/project folder.

C. Files created by all members of the project group in the /usr/share/project folder will get the same group permission as the folder itself.

D. The chmod command can only take a tree-digit argument.

Correct Answer: A

* A permission of "2755" for a directory means that everyone has read and execute permission, while the file owner and members of the file's group additionally have write permission. And any files or subdirectories created in that directory will inherit the parent directory's group id.

Incorrect:

Not D: chmod 2775 is a valid command.

[1Z0-460 Study Guide](#)

[1Z0-460 Exam Questions](#)

[1Z0-460 Braindumps](#)