

# GETCERTKEY



**GETCERTKEY**

100% guarantee you pass IT cert exam!

## Instant Update

We are checking our exam questions all the time.



Security & Privacy



24/7 customer support

## Free Demo Download

Try before you buy, Download a free sample of any of our exam questions and answers.



## One Year Free Update

Free update is available within One Year after your purchase.



<http://www.getcertkey.com>

No help, Full refund!

**Exam** : **305-300**

**Title** : LPIC-3 Exam 305:  
Virtualization and  
Containerization

**Vendor** : Lpi

**Version** : DEMO

**NO.1** Which of the following types of guest systems does Xen support? (Choose two.)

- A. Foreign architecture guests (FA)
- B. Paravirtualized guests (PVI)
- C. Emulated guests
- D. Container virtualized guests
- E. Fully virtualized guests

**Answer:** B E

Explanation:

Xen supports two types of guest systems: paravirtualized guests (PV) and fully virtualized guests (HVM).

\* Paravirtualized guests (PV) are guests that have been modified to run on the Xen hypervisor. They use a special kernel that communicates with the hypervisor through hypercalls, and use paravirtualized drivers for I/O devices. PV guests can run faster and more efficiently than HVM guests, but they require the guest operating system to be ported to Xen and to support the Xen ABI12.

\* Fully virtualized guests (HVM) are guests that run unmodified operating systems on the Xen hypervisor. They use hardware virtualization extensions, such as Intel VT-x or AMD-V, to create a virtual platform for the guest. HVM guests can run any operating system that supports the hardware architecture, but they incur more overhead and performance penalties than PV guests. HVM guests can also use paravirtualized drivers for I/O devices to improve their performance12.

The other options are not correct. Xen does not support foreign architecture guests (FA), emulated guests, or container virtualized guests.

\* Foreign architecture guests (FA) are guests that run on a different hardware architecture than the host.

For example, running an ARM guest on an x86 host. Xen does not support this type of virtualization, as it would require emulation or binary translation, which are very complex and slow techniques3.

\* Emulated guests are guests that run on a software emulator that mimics the hardware of the host or another platform. For example, running a Windows guest on a QEMU emulator. Xen does not support this type of virtualization, as it relies on the emulator to provide the virtual platform, not the hypervisor. Xen can use QEMU to emulate some devices for HVM guests, but not the entire platform14.

\* Container virtualized guests are guests that run on a shared kernel with the host and other guests, using namespaces and cgroups to isolate them. For example, running a Linux guest on a Docker container. Xen does not support this type of virtualization, as it requires the guest operating system to be compatible with the host kernel, and does not provide the same level of isolation and security as hypervisor-based virtualization56.

:

Xen Project Software Overview - Xen

Xen ARM with Virtualization Extensions - Xen

Xen Project Beginners Guide - Xen

QEMU - Xen

Docker overview | Docker Documentation

What is a Container? | App Containerization | VMware

**NO.2** What are some benefits of using Packer for image building? (Select all that apply)

- A. Built-in monitoring and logging
- B. Improved security through automated patching
- C. Faster provisioning of virtual machines
- D. Consistency in machine images across different environments

**Answer:** C D

Explanation:

Packer is designed to automate machine image creation, which leads to faster provisioning of virtual machines (C) because prebuilt images eliminate repetitive configuration steps at deployment time. Another primary benefit is consistency across environments (D), ensuring that development, testing, and production systems are built from the same image.

Packer does not provide built-in monitoring, and while it can apply patches during image creation, automated patching is not its primary or guaranteed function.

Thus, the correct answers are C and D.

**NO.3** Which of the following statements are true regarding VirtualBox?

- A. It is a hypervisor designed as a special kernel that is booted before the first regular operating system starts.
- B. It only supports Linux as a guest operating system and cannot run Windows inside a virtual machine.
- C. It requires dedicated shared storage, as it cannot store virtual machine disk images locally on block devices of the virtualization host.
- D. It provides both a graphical user interface and command line tools to administer virtual machines.
- E. It is available for Linux only and requires the source code of the currently running Linux kernel to be available.

**Answer:** D

Explanation:

VirtualBox is a hosted hypervisor, which means it runs as an application on top of an existing operating system, not as a special kernel that is booted before the first regular operating system starts<sup>1</sup>. VirtualBox supports a large number of guest operating systems, including Windows, Linux, Solaris, OS/2, and OpenBSD<sup>1</sup>. VirtualBox does not require dedicated shared storage, as it can store virtual machine disk images locally on block devices of the virtualization host, or on network shares, or on iSCSI targets<sup>1</sup>. VirtualBox provides both a graphical user interface (GUI) and command line tools (VBoxManage) to administer virtual machines<sup>1</sup>. VirtualBox is available for Windows, Linux, macOS, and Solaris hosts<sup>1</sup>, and does not require the source code of the currently running Linux kernel to be available. References:

\* Oracle VM VirtualBox: Features Overview

**NO.4** What happens when the following command is executed twice in succession?

```
docker run -tid -v data:/data debian bash
```

- A. The container resulting from the second invocation can only read the content of /data/ and cannot change it.
- B. Each container is equipped with its own independent data volume, available at /data/ in the respective container.
- C. Both containers share the contents of the data volume, have full permissions to alter its content

and mutually see their respective changes.

**D.** The original content of the container image data is available in both containers, although changes stay local within each container.

**E.** The second command invocation fails with an error stating that the volume data is already associated with a running container.

**Answer:** C

Explanation:

The command `docker run -tid -v data:/data debian bash` creates and runs a new container from the debian image, with an interactive terminal and a detached mode, and mounts a named volume data at /data in the container<sup>12</sup>. If the volume data does not exist, it is created automatically<sup>3</sup>. If the command is executed twice in succession, two containers are created and run, each with its own terminal and process ID, but they share the same volume data. This means that both containers can access, modify, and see the contents of the data volume, and any changes made by one container are reflected in the other container. Therefore, the statement C is true and the correct answer. The statements A, B, D, and E are false and incorrect, as they do not describe the behavior of the command or the volume correctly. References:

\* 1: [docker run | Docker Docs](#).

\* 2: [Docker run reference | Docker Docs - Docker Documentation](#).

\* 3: [Use volumes | Docker Documentation](#).

\* [4]: [How to Use Docker Run Command with Examples - phoenixNAP](#).

**NO.5** Which of the following is true about LXC?

**A.** It is only compatible with Windows

**B.** It does not use the Linux kernel

**C.** It is a container runtime

**D.** It is primarily used for virtual machine management

**Answer:** C

Explanation:

LXC (Linux Containers) is a container runtime that uses Linux kernel features such as namespaces and cgroups to provide lightweight, isolated environments. According to containerization documentation, LXC allows multiple Linux systems (containers) to run on a single host while sharing the same kernel. LXC is Linux-specific, uses the Linux kernel directly, and is not a virtual machine manager. Therefore, the correct answer is C.

**NO.6** What is Packer?

**A.** A container orchestration tool

**B.** An image building tool

**C.** A configuration management tool

**D.** A virtualization platform

**Answer:** B

Explanation:

Packer is an image automation and building tool developed by HashiCorp. According to virtualization and DevOps documentation, Packer is used to create machine images for multiple platforms from a single configuration file.

Packer is not an orchestration platform, configuration management tool, or hypervisor. Its primary

role is to ensure consistent, repeatable, and automated image creation, making B the correct answer.

**NO.7** When setting up a KVM virtualization host, which one of the following components is NOT required?

- A. bridgeutils
- B. kvm kernel modules
- C. qemu
- D. libvirt
- E. virsh

**Answer:** E

Explanation:

When configuring a KVM-based virtualization host, several core components are mandatory to enable and manage virtual machines. According to KVM and virtualization documentation, KVM kernel modules are essential because they provide hardware-assisted virtualization support within the Linux kernel. QEMU is required to perform hardware emulation and manage virtual machine execution. Libvirt acts as the virtualization management API, enabling centralized and secure control of virtual machines. Additionally, bridgeutils is commonly required to configure network bridges, allowing virtual machines to communicate with external networks.

However, virsh is not strictly required. Virsh is a command-line utility that interacts with libvirt to manage virtual machines, but it is only a management interface, not a core dependency. Virtual machines can still be created and managed using alternative tools such as virt-manager, Ansible, OpenStack, or custom API-based solutions without virsh being installed.

Virtualization documentation clearly distinguishes between essential backend components (KVM, QEMU, libvirt) and optional management tools (virsh). Therefore, while virsh is widely used and highly recommended for administrative convenience, it is not a mandatory component for a functional KVM virtualization host.

**NO.8** Which of the following tools is not typically used for VM provisioning and image creation?

- A. Kubernetes
- B. Vagrant
- C. Packer
- D. cloud-init

**Answer:** A

Explanation:

VM provisioning and image creation tools are used to build, configure, and initialize virtual machine images. According to virtualization documentation, Packer is specifically designed for automated image creation, Vagrant is used for provisioning reproducible development environments, and cloud-init is used to configure instances at first boot.

Kubernetes, however, is a container orchestration platform, not a VM provisioning or image creation tool.

It manages containerized workloads after infrastructure has already been provisioned.

Documentation clearly separates infrastructure provisioning tools from application orchestration platforms

. As such, Kubernetes is not typically used for VM provisioning or image creation.

Therefore, the correct answer is A.

**NO.9** The command `virsh vol-list vms` returns the following error:

error: failed to get pool 'vms'

error: Storage pool not found: no storage pool with matching name 'vms ' Given that the directory `/vms` exists, which of the following commands resolves this issue?

- A. `dd if=/dev/zero of=/vms bs=1 count=0 flags=name:vms`
- B. `libvirt-poolctl new --name=/vms --type=dir --path=/vms`
- C. `qemu-img pool vms:/vms`
- D. `virsh pool-create-as vms dir --target /vms`
- E. `touch /vms/.libvirtpool`

**Answer:** D

Explanation:

The command `virsh pool-create-as vms dir --target /vms` creates and starts a transient storage pool named `vms` of type `dir` with the target directory `/vms`. This command resolves the issue of the storage pool not found error, as it makes the existing directory `/vms` visible to libvirt as a storage pool. The other commands are invalid because:

\* `dd if=/dev/zero of=/vms bs=1 count=0 flags=name:vms` is not a valid command syntax. The `dd` command does not take a `flags` argument, and the output file `/vms` should be a regular file, not a directory<sup>3</sup>.

\* `libvirt-poolctl new --name=/vms --type=dir --path=/vms` is not a valid command name. There is no such command as `libvirt-poolctl` in the libvirt package<sup>4</sup>.

\* `qemu-img pool vms:/vms` is not a valid command syntax. The `qemu-img` command does not have a `pool` subcommand, and the `vms:/vms` argument is not a valid image specification<sup>5</sup>.

\* `touch /vms/.libvirtpool` is not a valid command to create a storage pool. The `touch` command only creates an empty file, and the `.libvirtpool` file is not recognized by libvirt as a storage pool configuration file<sup>6</sup>.

References:

1: `virsh` - difference between `pool-define-as` and `pool-create-as` - Stack Overflow

2: `dd(1)` - Linux manual page - man7.org

3: 12.3.3. Creating a Directory-based Storage Pool with `virsh` - Red Hat Customer Portal

4: libvirt - Linux Man Pages (3)

5: `qemu-img(1)` - Linux manual page - man7.org

6: `touch(1)` - Linux manual page - man7.org

**NO.10** How can data be shared between several virtual machines running on the same Linux-based host system?

- A. By writing data to the file system since all virtual machines on the same host system use the same file system.
- B. By mounting other virtual machines' file systems from `/dev/virt-disks/remote/`.
- C. By setting up a ramdisk in one virtual machine and mounting it using its UUID in the other VMs.
- D. By using a network file system or file transfer protocol.
- E. By attaching the same virtual hard disk to all virtual machines and activating EXT4 sharing extensions on it.

**Answer:** D

Explanation:

The correct way to share data between several virtual machines running on the same Linux-based host system is by using a network file system or file transfer protocol. A network file system (NFS) is a distributed file system protocol that allows a user on a client computer to access files over a network in a manner similar to how local storage is accessed<sup>1</sup>. A file transfer protocol (FTP) is a standard network protocol used for the transfer of computer files between a client and server on a computer network<sup>2</sup>. Both methods allow data to be shared between virtual machines regardless of their underlying file systems or virtualization technologies. The other options are incorrect because they either do not work or are not feasible. Option A is wrong because each virtual machine has its own file system that is not directly accessible by other virtual machines. Option B is wrong because there is no such device as `/dev/virt-disks/remote/` that can be used to mount other virtual machines' file systems. Option C is wrong because a ramdisk is a volatile storage device that is not suitable for sharing data between virtual machines. Option E is wrong because attaching the same virtual hard disk to multiple virtual machines can cause data corruption and conflicts, and EXT4 does not have any sharing extensions that can prevent this. References:<https://kb.vmware.com/s/article/1012706>  
<https://bing.com/search?q=data+sharing+between+virtual+machines>

**NO.11** What is the purpose of the kubelet service in Kubernetes?

- A. Provide a command line interface to manage Kubernetes.
- B. Build a container image as specified in a Dockerfile.
- C. Manage permissions of users when interacting with the Kubernetes API.
- D. Run containers on the worker nodes according to the Kubernetes configuration.
- E. Store and replicate Kubernetes configuration data.

**Answer:** D

Explanation:

The purpose of the kubelet service in Kubernetes is to run containers on the worker nodes according to the Kubernetes configuration. The kubelet is an agent or program that runs on each node and communicates with the Kubernetes control plane. It receives a set of PodSpecs that describe the desired state of the pods that should be running on the node, and ensures that the containers described in those PodSpecs are running and healthy. The kubelet also reports the status of the node and the pods back to the control plane. The kubelet does not manage containers that were not created by Kubernetes. References:

\* Kubernetes Docs - kubelet

\* Learn Steps - What is kubelet and what it does: Basics on Kubernetes

**NO.12** Which of the following tools is used to interact with XenStore?

- A. xendo
- B. oxs
- C. xstore
- D. xentore-ls
- E. xl store

**Answer:** D

Explanation:

XenStore is a shared configuration and state database used by Xen domains. According to Xen documentation, tools such as `asxenstore-ls`, `xenstore-read`, and `xenstore-write` are used to interact directly with XenStore.

Although option D contains a minor typographical error ("xentore-ls" instead of xenstore-ls), it clearly refers to the correct and documented utility. None of the other options represent standard XenStore interaction tools.

Therefore, despite the typo, D is the correct answer.