

# The Best 305-300 Dumps Discount & Leader in Certification Exams Materials & Fantastic 305-300 Formal Test



P.S. Free & New 305-300 dumps are available on Google Drive shared by PassReview: <https://drive.google.com/open?id=1bvOC3laXIU-BRh6akAw7os5Kdz3WOwaB>

In today's fast-paced world, having access to LPIC-3 Exam 305: Virtualization and Containerization (305-300) study material on the go is important. PassReview LPIC-3 Exam 305: Virtualization and Containerization (305-300) PDF questions are compatible with all smart devices, allowing you to study and prepare for the 305-300 Exam whenever and wherever you choose. Since you can access real Lpi 305-300 dumps in PDF from your smartphone or tablet, you can easily fit 305-300 exam preparation into your busy schedule.

The Lpi 305-300 Exam is intended for professionals who have a strong IT background, particularly in Linux and open-source technologies. It is ideal for system administrators, network administrators, DevOps engineers, and cloud architects who want to advance their careers in virtualization and containerization. LPIC-3 Exam 305: Virtualization and Containerization certification is recognized globally and is a valuable addition to an IT professional's resume. With the increasing demand for cloud computing and virtualization technologies, obtaining this certification can open up many opportunities for IT professionals in the industry.

>> 305-300 Dumps Discount <<

## Professional Lpi - 305-300 Dumps Discount

You don't know how to acquire a promotion quickly while you're trying to get a new job or already have one but need a promotion. The sole option is Lpi 305-300 certification, which makes it simple for you to advance in your career. Your skills will advance and your resume will be enhanced thanks to the Lpi 305-300 Certification.

## Lpi LPIC-3 Exam 305: Virtualization and Containerization Sample Questions (Q118-Q123):

### NEW QUESTION # 118

What is the name of the kernel module that is required to be loaded in order to use KVM on an Intel CPU architecture? (Specify ONLY the module name without any path information and with or without the module suffix.) Solution: `kvm-intel.ko` - or - `kvm-intel` - or - `kvm_intel.ko` - or - `kvm_intel`  
Determine whether the given solution is correct?

- A. Incorrect
- B. Correct

**Answer: B**

Explanation:

Kernel-based Virtual Machine (KVM) relies on hardware-assisted virtualization features provided by modern CPUs. On Intel CPU architectures, this support is enabled through the `kvm_intel` kernel module. Official KVM and Linux virtualization documentation clearly states that the required module for Intel processors is named `kvm_intel`, with the optional `.ko` suffix when referring to the kernel object file.

The provided solution lists multiple acceptable representations of the module name, including `kvm_intel` and `kvm_intel.ko`, both of

which are valid and correct. Linux kernel module naming conventions allow the module to be referenced with or without the .ko suffix when loading it using tools such as modprobe or lsmod.

Although variants using a hyphen (kvm-intel) are not the canonical kernel module name, the solution explicitly includes the correct and documented module name. Therefore, the solution correctly identifies the required kernel module for enabling KVM on Intel CPUs.

Virtualization documentation emphasizes that KVM functionality requires both the generic kvm module and the CPU-specific module (kvm\_intel for Intel or kvm\_amd for AMD). Hence, the determination that the solution is correct aligns with verified documentation.

### NEW QUESTION # 119

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.

### NEW QUESTION # 120

After setting up a data container using the following command:

```
docker create -v /data --name datastore debian /bin/true
```

how is an additional new container started which shares the /data volume with the datastore container?

- A. `docker run --volumes-from datastore --name service debian bash`
- B. `docker run --share-with datastore --name service debian bash`
- C. `docker run -v /data --name service debian bash`
- D. `docker run -v datastore:/data --name service debian bash`
- E. `docker run --volume-backend datastore -v /data --name service debian bash`

**Answer: A**

### NEW QUESTION # 121

Which of the following commands executes a command in a running LXC container?

- A. `lxc-batch`
- B. `lxc-attach`
- C. `lxc-run`
- D. `lxc-eval`
- E. `lxc-enter`

**Answer: B**

Explanation:

The command `lxc-attach` is used to execute a command in a running LXC container. It allows the user to start a process inside the container and attach to its standard input, output, and error streams. For example, the command `lxc-attach -n mycontainer -- ls -lh /home` will list all the files and directories in the `/home` directory of the container named `mycontainer`. The other options are not valid LXC commands. The command `lxc-batch` does not exist. The command `lxc-run` is an alias for `lxc-start`, which is used to start a container, not to execute a command in it. The command `lxc-enter` is also an alias for `lxc-attach`, but it is deprecated and should not be used. The command `lxc-eval` is also not a valid LXC command. References:



BRh6akAw7os5Kdz3WOwaB