Lpi.305-300.v2024-06-27.q42

Exam Code:	305-300
Exam Name:	LPIC-3 Exam 305: Virtualization and Containerization
Certification Provider:	Lpi
Free Question Number:	42
Version:	v2024-06-27
# of views:	1323
# of Questions views:	2553
https://www.dumpsfiles.com/files/Lpi/305-300/Lpi.305-300.v2024-06-27.q42	

NEW QUESTION: 1
Which of the following kinds of data cancloud-initprocess directly from user-data? (Choose three.)
A. Shell scripts to execute
B. Lists of URLs to import
C. ISO images to boot from
D. cloud-config declarations in YAML
E. Base64-encoded binary files to execute
Answer: A,B,D (LEAVE A REPLY)
Explanation
Cloud-init is a tool that allows users to customize the configuration and behavior of cloud instances during the boot process. Cloud-init can process different kinds of data that are passed to the instance via user-data, which is a mechanism provided by various cloud providers to inject data into the instance. Among the kinds of data that cloud-init can process directly from user-data are:
* Shell scripts to execute: Cloud-init can execute user-data that is formatted as a shell script, starting with the #!/bin/sh or #!/bin/bash shebang. The script can contain any commands that are valid in the shell environment of the instance. The script is executed as the root user during the boot process12.
* Lists of URLs to import: Cloud-init can import user-data that is formatted as a list of URLs, separated by newlines. The URLs can point to any valid data source that cloud-init supports, such as shell scripts, cloud-config files, or include files. The URLs are fetched and processed by cloud-init in the order they appear in the list13.
* cloud-config declarations in YAML: Cloud-init can process user-data that is formatted as a cloud-config file, which is a YAML document that contains declarations for various cloud-init modules. The cloud-config file can specify various aspects of the instance configuration, such as hostname, users, packages, commands, services, and more. The cloud-config file must start with the #cloud-config header14.
The other kinds of data listed in the question are not directly processed by cloud-init from user-data. They are either not supported, not recommended, or require additional steps to be processed. These kinds of data are:
* ISO images to boot from: Cloud-init does not support booting from ISO images that are passed as user-data. ISO images are typically used to install an operating system on a physical or virtual machine, not to customize an existing cloud instance. To boot from an ISO image, the user would need to attach it as a secondary disk to the instance and configure the boot order accordingly5.

* Base64-encoded binary files to execute: Cloud-init does not recommend passing binary files as user-data, as they may not be compatible with the instance's architecture or operating system.
Base64-encoding does not change this fact, as it only converts the binary data into ASCII characters. To execute a binary file, the user would need to decode it and make it executable on the instance6.
References:
* User-Data Formats - cloud-init 22.1 documentation
* User-Data Scripts
* Include File
* Cloud Config
* How to Boot From ISO Image File Directly in Windows
* How to run a binary file as a command in the terminal?.

NEW QUESTION: 2
Which of the following statements about the commandlxc-checkpointis correct?
A. It only works on stopped containers.
B. It creates a container image based on an existing container.
C. It creates a clone of a container.
D. It doubles the memory consumption of the container.
E. It writes the status of the container to a file.
Answer: E (LEAVE A REPLY)

NEW QUESTION: 3
FILL BLANK
What is the default path to the Docker daemon configuration file on Linux? (Specify the full name of the file,Including path.)
Answer:
/etc/docker/daemon.json

NEW QUESTION: 4
What is the purpose of the commandvagrant init?
A. It executes a provisioning tool in a running box.
B. It starts a Vagrant box.
C. It creates a Vagrant configuration file.
D. It installs Vagrant on a Linux host.
E. It downloads a Vagrant box.
Answer: C (LEAVE A REPLY)
Explanation
The command vagrant init is used to initialize the current directory to be a Vagrant environment by creating an initial Vagrantfile if one does not already exist1. The Vagrantfile contains the configuration settings for the Vagrant box, such as the box name, box URL, network settings, synced folders, provisioners, etc. The command vagrant init does not execute any provisioning tool, start any box, install Vagrant on a Linux host, or download any box. Those actions are performed by other Vagrant commands, such as vagrant provision, vagrant up, vagrant install, and vagrant box add, respectively. References:

* 1: vagrant init - Command-Line Interface | Vagrant | HashiCorp Developer

NEW QUESTION: 5
After creating a new Docker network using the following command:
docker network create --driver bridge isolated_nw
which parameter must be added todocker createin order to attach a container to the network?
A. --eth0=isolated_nw
B. --alias=isolated_nw
C. --ethernet=isolated_nw
D. --network=isolated_nw
E. --attach=isolated_nw
Answer: D (LEAVE A REPLY)
Explanation
To attach a container to a network when creating it, the --network flag must be used with the name of the network as the argument. The --network flag specifies the network mode for the container. By default, the network mode is bridge, which means the container is connected to the default bridge network. However, if a custom network is created, such as isolated_nw in this case, the container must be explicitly attached to it using the --network flag. For example, to create a container named web1 and attach it to the isolated_nw network, the command would be:
docker create --name web1 --network isolated_nw nginx
The other options are not valid parameters for docker create. The --eth0, --ethernet, and --attach flags do not exist. The --alias flag is used to specify an additional network alias for the container on a user-defined network, but it does not attach the container to the network. References:
* docker network create | Docker Documentation1
* docker create | Docker Documentation
* Networking overview | Docker Docs2

NEW QUESTION: 6
FILL BLANK
What command is used to run a process in a new Linux namespace? (Specify ONLY the command without any path or parameters.)
Answer:
unshare

NEW QUESTION: 7
In order to use the optiondom0_memto limit the amount of memory assigned to the Xen Domain-0, where must this option be specified?
A. In the bootloader configuration, when Xen is booted.
B. In its Makefile, when Xen is built.
C. In its .config file, when the Domain-0 kernel is built.
D. In the configuration file /etc/xen/Domain-0.cfg, when Xen starts.
E. In any of Xen's global configuration files.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 8
Which of the following statements are true regarding a Pod in Kubernetes? (Choose two.)
A. All containers of a Pod run on the same node.
B. Pods are always created automatically and cannot be explicitly configured.
C. A Pod is the smallest unit of workload Kubernetes can run.
D. When a Pod fails, Kubernetes restarts the Pod on another node by default.
E. systemd is used to manage individual Pods on the Kubernetes nodes.
Answer: A,C (LEAVE A REPLY)
Explanation
A Pod in Kubernetes is a collection of one or more containers that share the same network and storage resources, and a specification for how to run the containers. A Pod is the smallest unit of workload Kubernetes can run, meaning that it cannot be divided into smaller units. Therefore, option C is correct. All containers of a Pod run on the same node, which is the smallest unit of computing hardware in Kubernetes. A node is a physical or virtual machine that hosts one or more Pods. Therefore, option A is also correct. Pods are not always created automatically and cannot be explicitly configured. Pods can be created manually using YAML or JSON files, or using commands like kubectl run or kubectl create. Pods can also be created automatically by higher-level controllers, such as Deployment, ReplicaSet, or StatefulSet. Therefore, option B is incorrect.
When a Pod fails, Kubernetes does not restart the Pod on another node by default. Pods are ephemeral by nature, meaning that they can be terminated or deleted at any time. If a Pod is managed by a controller, the controller will create a new Pod to replace the failed one, but it may not be on the same node. Therefore, option D is incorrect. systemd is not used to manage individual Pods on the Kubernetes nodes. systemd is a system and service manager for Linux operating systems that can start and stop services, such as docker or kubelet. However, systemd does not interact with Podsdirectly. Pods are managed by the kubelet service, which is an agent that runs on each node and communicates with the Kubernetes control plane. Therefore, option E is incorrect. References:
* Pods | Kubernetes
* What is a Kubernetes pod? - Red Hat
* What's the difference between a pod, a cluster, and a container?
* What are Kubernetes Pods? | VMware Glossary
* Kubernetes Node Vs. Pod Vs.Cluster: Key Differences - CloudZero

NEW QUESTION: 9
Which of the following commands executes a command in a running LXC container?
A. lxc-accach
B. lxc-batch
C. lxc-run
D. lxc-enter
E. lxc-eval
Answer: A (LEAVE A REPLY)
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 streams1. 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 mycontainer1. 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 it2. The command lxc-enter is also an alias for lxc-attach, but it is deprecated and should not be used3. The command lxc-eval is also not a valid LXC command. References:
* 1: Executing a command inside a running LXC - Unix & Linux Stack Exchange.
* 2: lxc-start: start a container. - SysTutorials.
* 3: lxc-attach: start a process inside a running container. - SysTutorials.

NEW QUESTION: 10
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 (LEAVE A REPLY)
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 starts1. VirtualBox supports a large number of guest operating systems, including Windows, Linux, Solaris, OS/2, and OpenBSD1. 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 targets1. VirtualBox provides both a graphical user interface (GUI) and command line tools (VBoxManage) to administer virtual machines1. VirtualBox is available for Windows, Linux, macOS, and Solaris hosts1, and does not require the source code of the currently running Linux kernel to be available. References:
* Oracle VM VirtualBox: Features Overview

NEW QUESTION: 11
What is the purpose of the packer inspect subcommand?
A. Execute commands within a running instance of a Packer image.
B. Retrieve files from an existing Packer image.
C. List the artifacts created during the build process of a Packer image.
D. Show usage statistics of a Packer image.
E. Display an overview of the configuration contained in a Packer template.
Answer: E (LEAVE A REPLY)

NEW QUESTION: 12
Which CPU flag indicates the hardware virtualization capability on an AMD CPU?
A. PVM
B. SVM
C. HVM
D. VIRT

E. VMX
Answer: B (LEAVE A REPLY)

NEW QUESTION: 13
Which of the following values would be valid in the FROM statement in aDockerfile?
A. ubuntu:focal
B. docker://ubuntu: focal
C. registry:ubuntu:focal
D. file:/tmp/ubuntu/Dockerfile
E. http://docker.example.com/images/ubuntu-focal.iso
Answer: A (LEAVE A REPLY)
Explanation
The FROM statement in a Dockerfile specifies the base image from which the subsequent instructions are executed1. The value of the FROM statement can be either an image name, an image name with a tag, or an image ID1. The image name can be either a repository name or a repository name with a registry prefix2. For example, ubuntu is a repository name, and docker.io/ubuntu is a repository name with a registry prefix2. The tag is an optional identifier that can be used to specify a particular version or variant of an image1. For example, ubuntu:focal refers to the image with the focal tag in the ubuntu repository2. The image ID is a unique identifier that is automatically generated when an image is built or pulled1. For example, sha256:9b0dafaadb1cd1d14e4db51bd0f4c0d56b6b551b2982b2b7c637ca143ad605d2 is an image ID3.
Therefore, the only valid value in the FROM statement among the given options is ubuntu:focal, which is an image name with a tag.
The other options are invalid because:
* docker://ubuntu:focal is not a valid image name format. The docker:// prefix is used to specify a transport protocol, not a registry prefix4.
* registry:ubuntu:focal is not a valid image name format. The registry prefix should be a valid hostname or IP address, not a generic term2.
* file:/tmp/ubuntu/Dockerfile is not a valid image name format. The file: prefix is used to specify a local file path, not an image name5.
* http://docker.example.com/images/ubuntu-focal.iso
is not a valid image name format.
The
http:// prefix is used to specify a web URL, not an image name
5.
References:
* 1: Dockerfile reference | Docker Docs
* 2: docker - Using FROM statement in dockerfile - Stack Overflow
* 3: How to get the image id from a docker image - Stack Overflow
* 4: skopeo - Docker Registry v2 API tool - Linux Man Pages (1)
* 5: How to build a Docker image from a local Dockerfile? - Stack Overflow

NEW QUESTION: 14
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 (LEAVE A REPLY)
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 container12. If the volume data does not exist, it is created automatically3. 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.

NEW QUESTION: 15
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 (LEAVE A REPLY)
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 accessed1. 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 network2. 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

NEW QUESTION: 16
What is the purpose ofcloud-init?
A. Replace common Linux inic systems, such as systemd or SysV init.
B. Assign an laaS instance to a specific computing node within a cloud.
C. Standardize the configuration of infrastructure services, such as load balancers or virtual firewalls in a cloud.
D. Orchestrate the creation and start of multiple related laaS instances.
E. Prepare the generic image of an laaS instance to fit a specific instance's configuration.
Answer: E (LEAVE A REPLY)
Explanation
Cloud-init is a tool that processes configurations and runs through five stages during the initial boot of Linux VMs in a cloud. It allows users to customize a Linux VM as it boots for the first time, by applying user data to the instance. User data can include scripts, commands, packages, files, users, groups, SSH keys, and more.
Cloud-init can also interact with various cloud platforms and services, such as Azure, AWS, OpenStack, and others. The purpose of cloud-init is to prepare the generic image of an laaS instance to fit a specific instance's configuration, such as hostname, network, security, and application settings. References:
* Cloud-init - The standard for customising cloud instances
* Understanding cloud-init - Azure Virtual Machines
* Tutorial - Customize a Linux VM with cloud-init in Azure - Azure Virtual Machines

Valid 305-300 Dumps shared by TrainingDump.com for Helping Passing 305-300 Exam! TrainingDump.com now offer the **newest 305-300 exam dumps**, the TrainingDump.com 305-300 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingDump.com 305-300 dumps with Test Engine here: https://www.trainingdump.com/Lpi/305-300-practice-exam-dumps.html (**125** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 17
Ifdocker stackis to be used to run a Docker Compose file on a Docker Swarm, how are the images referenced in the Docker Compose configuration made available on the Swarm nodes?
A. docker stack transfers the image from its local Docker cache to each Swarm node.
B. docker stack instructs the Swarm nodes to pull the images from a registry, although it does not upload the images to the registry.
C. docker stack triggers the build process for the images on all nodes of the Swarm.
D. docker stack passes the images to the Swarm master which distributes the images to all other Swarm nodes.
E. docker stack builds the images locally and copies them to only those Swarm nodes which run the service.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 18
How does Packer interact with system images?
A. Packer has to be installed within the target image and is executed during the image's first boot in order to execute preparation tasks.

B. Packer installs a client within the image which has to be run periodically via cron in order to retrieve the latest template from the Packer server and apply it locally.
C. Packer periodically connects through the network to the Packer daemons of all running Packer images in order to re-apply the whole template to the running instance.
D. Packer downloads and extracts an image in order to make changes to the image's file system, repack the modified image and upload it again.
E. Packer creates an instance based on a source image, prepares the instance through a network connection and bundles the resulting instance as a new system image.
Answer: E (LEAVE A REPLY)
Explanation
Packer is a tool that automates the creation of identical machine images for multiple platforms from a single source configuration. Packer works by creating an instance based on a source image, which is a pre-existing image that serves as a starting point. Packer then connects to the instance through a network connection, such as SSH or WinRM, and runs various commands and scripts to install and configure software within the instance. Packer then shuts down the instance and creates a new system image from it, which can be used to launch new instances. Packer supports many platforms, such as AWS, Azure, VMware, Docker, and others. Packer does not install any software or run any daemon within the target image, nor does it periodically connect to the running instances to re-apply the template. Packer also does not modify the source image directly, but creates a new image from the modified instance. References:
* Packer by HashiCorp
* HashiCorp Packer - Build Automated Machine Images
* Introduction | Packer | HashiCorp Developer

NEW QUESTION: 19
In an IaaS cloud, what is a common method for provisioning new computing instances with an operating system and software?
A. Each new instance is connected to the installation media of a Linux distribution and provides access to the installer by logging in via SSH.
B. Each new instance is created based on an image file that contains the operating system as well as software and default configuration for a given purpose.
C. Each new instance is a clone of another currently running instance that includes all the software, data and state of the original instance.
D. Each new instance is connected via a VPN with the computer that started the provisioning and tries to PXE boot from that machine.
E. Each new instance contains a minimal live system running from a virtual CD as the basis from which the administrator deploys the target operating system.
Answer: B (LEAVE A REPLY)
Explanation
In an IaaS cloud, the most common method for provisioning new computing instances is to use an image file that contains a pre-installed operating system and software. This image file is also known as a machine image, a virtual appliance, or a template. The image file can be customized for a specific purpose, such as a web server, a database server, or a development environment. The image file can be stored in a repository or a library that is accessible by the cloud provider or the user. When a new instance is requested, the cloud provider copies the image file to a virtual disk and attaches it to the instance. The instance then boots from the

virtual disk and runs the operating system and software from the image file. This method is faster and more efficient than installing the operating system and software from scratch for each new instance. It also ensures consistency and reliability across multiple instances that use the same image file. References:
* LPI Virtualization and Containerization Exam Objectives, Topic 305.1: Virtualization Concepts and Theory, Objective: Describe the concept of machine images and templates
* LPI Virtualization and Containerization Study Guide, Chapter 1: Virtualization Concepts and Theory, Section: Machine Images and Templates
* LPI LPIC-3 305 Certification Sample Questions and Practice Exam, Question 10: In an IaaS cloud, what is a common method for provisioning new computing instances with an operating system and software?

NEW QUESTION: 20
Which directory is used bycloud-initto store status information and configuration information retrieved from external sources?
A. /var/lib/cloud/
B. /etc/cloud-init/cache/
C. /proc/sys/cloud/
D. /tmp/.cloud/
E. /opt/cloud/var/
Answer: (SHOW ANSWER)
Explanation
cloud-init uses the /var/lib/cloud/ directory to store status information and configuration information retrieved from external sources, such as the cloud platform'smetadata service or user data files. The directory contains subdirectories for different types of data, such as instance, data, handlers, scripts, and sem. The instance subdirectory contains information specific to the current instance, such as the instance ID, the user data, and the cloud-init configuration. The data subdirectory contains information about the data sources that cloud-init detected and used. The handlers subdirectory contains information about the handlers that cloud-init executed.
The scripts subdirectory contains scripts that cloud-init runs at different stages of the boot process, such as per-instance, per-boot, per-once, and vendor. The sem subdirectory contains semaphore files that cloud-init uses to track the execution status of different modules and stages. References:
* Configuring and managing cloud-init for RHEL 8 - Red Hat Customer Portal
* vsphere - what is the linux file location where the cloud-init user ...

NEW QUESTION: 21
Which of the following values are valid in thefirmwareattribute of a<os>element in a libvirt domain definition? (Choose two.)
A. pcie
B. virtio
C. bios
D. scsi
E. efi
Answer: (SHOW ANSWER)

NEW QUESTION: 22
What does IaaS stand for?

A. Intelligence as a Service
B. Information as a Service
C. Instances as a Service
D. Infrastructure as a Service
E. Integration as a Service
Answer: D (LEAVE A REPLY)

NEW QUESTION: 23
If aDockerfilecontains the following lines:
WORKDIR /
RUN cd /tmp
RUN echo test > test
where is the filetestlocated?
A. /ting/test within the container image.
B. /root/tesc within the container image.
C. /test within the container image.
D. /tmp/test on the system running docker build.
E. test in the directory holding the Dockerf ile.
Answer: (SHOW ANSWER)
Explanation
The WORKDIR instruction sets the working directory for any subsequent RUN, CMD, ENTRYPOINT, COPY and ADD instructions that follow it in the Dockerfile1. The RUN instruction executes commands in a new layer on top of the current image and commits the results2. The RUN cd command does not change the working directory for the next RUN instruction, because each RUN command runs in a new shell and a new environment3. Therefore, the file test is created in the root directory (/) of the container image, not in the /tmp directory. References:
* Dockerfile reference: WORKDIR
* Dockerfile reference: RUN
* difference between RUN cd and WORKDIR in Dockerfile

NEW QUESTION: 24
If aDockerfilecontains the following lines:
WORKDIR /
RUN cd /tmp
RUN echo test > test
where is the filetestlocated?
A. /test within the container image.
B. /root/tesc within the container image.
C. /ting/test within the container image.
D. test in the directory holding the Dockerf ile.
E. /tmp/test on the system running docker build.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 25
Which of the following services can QEMU provide in a user network? (Choose three.)
A. CIFS
B. AppleTalk
C. BGP
D. DHCP
E. TFTP
Answer: (SHOW ANSWER)

NEW QUESTION: 26
Which of the following statements are true about container-based virtualization? (Choose two.)
A. Each container runs its own operating system kernel.
B. Different containers may use different distributions of the same operating system.
C. Container-based virtualization relies on hardware support from the host system's CPU.
D. All containers run within the operating system kernel of the host system.
E. Linux does not support container-based virtualization because of missing kernel APIs.
Answer: (SHOW ANSWER)
Explanation
Container-based virtualization is a method of operating system-level virtualization that allows multiple isolated user spaces (containers) to run on the same host system1. Each container shares the same operating system kernel as the host, but has its own file system, libraries, and processes2. Therefore, the statements A and C are false, as containers do not run their own kernels or rely on hardware support from the CPU. The statement E is also false, as Linux does support container-based virtualization through various technologies, such as cgroups, namespaces, LXC, Docker, etc12. The statement B is true, as different containers may use different distributions of the same operating system, such as Debian, Ubuntu, Fedora, etc., as long as they are compatible with the host kernel3. The statement D is also true, as all containers run within the operating system kernel of the host system, which provides isolation and resource management for them12. References:
* 1: Containerization (computing) - Wikipedia.
* 2: What are containers? | Google Cloud.
* 3: What is Container-Based Virtualization? - StackHowTo.

NEW QUESTION: 27
What is true aboutcontainerd?
A. It is the initial process run at the start of any Docker container.
B. It uses rune to start containers on a container host.
C. It requires the Docker engine and Docker CLI to be installed.
D. It runs in each Docker container and provides DHCP client functionality
E. It is a text file format defining the build process of containers.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 28

In order to use the optiondom0_memto limit the amount of memory assigned to the Xen Domain-0, where must this option be specified?

A. In the bootloader configuration, when Xen is booted.

B. In any of Xen's global configuration files.

C. In its .config file, when the Domain-0 kernel is built.

D. In the configuration file /etc/xen/Domain-0.cfg, when Xen starts.

E. In its Makefile, when Xen is built.

Answer: A (LEAVE A REPLY)

Explanation

The option dom0_mem is used to set the initial and maximum memory size of the Domain-0, which is the privileged domain that starts first and manages the unprivileged domains (DomU) in Xen. The option dom0_mem must be specified in the bootloader configuration, such as GRUB or GRUB2, when Xen is booted.

This ensures that the Domain-0 kernel can allocate memory for storing memory metadata and network related parameters based on the boot time amount of memory. If the option dom0_mem is not specified in the bootloader configuration, the Domain-0 will use all the available memory on the host system by default, which may cause performance and security issues. References:

* Managing Xen Dom0s CPU and Memory
* Xen Project Best Practices
* Dom0 Memory - Where It Has Not Gone

NEW QUESTION: 29

Which command within virsh lists the virtual machines that are running on the current host?

A. l view

B. list-vm

C. list

D. show

E. list-all

Answer: (SHOW ANSWER)

Explanation

The command virsh list is used to list all running domains (VMs) on the current host. The command virsh list --all can be used to list both active and inactive domains. The other options are not valid virsh commands. The command virsh list is a basic command that lists all running domains (VMs). You can also list all configured VMs by adding the --all option. This is useful if you want to see all VMs configured in the target hypervisor that you can use on subsequent commands1. References:

* 1: 8 Linux virsh subcommands for managing VMs on the command line | Enable Sysadmin.

NEW QUESTION: 30

Which of the following statements is true regarding the following output ofxl list:

Name	ID	Mem	VCPUs	State	Time(s)
Domain-0	0	1024	1	r-----	498.7
Debian	2	1024	1	--p---	783.5
Ubuntu	6	1024	1	-b----	313.6
CentOS	7	2048	2	r-----	455.1

dumpsfiles.com
Linux Professional Institute

A. CentOS is the domain which has consumed the most CPU time.
B. Both Debian and Ubuntu require xl commands to start running.
C. The domain with ID 2 usesPara virtualization.
D. Ubuntu is idle or waiting for I/O.
E. It is necessary to use the xl command to change Ubuntu's state to running.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 31
Which of the following services can QEMU provide in a user network? (Choose three.)
A. DHCP
B. BGP
C. CIFS
D. AppleTalk
E. TFTP
Answer: (SHOW ANSWER)
Explanation
QEMU can provide some network services in a user network, which is a mode of networking that does not require any administrator privilege to run. The user network uses the SLIRP TCP/IP emulator to create a virtual NAT'ted subnet, with a DHCP server started by QEMU that gives out IP addresses to the guest machines and puts the host on 10.0.2.21. QEMU can also provide a TFTP server in the user network, which can be used to boot the guest machines from a network image. The TFTP server can be configured with the -tftp option2. QEMU does not provide BGP, CIFS, or AppleTalk services in the user network. BGP is a routing protocol that is used to exchange routing information between autonomous systems on the Internet3. CIFS is a file-sharing protocol that is used to access files and printers on a network4. AppleTalk is a deprecated network protocol suite that was used by Apple devices5. These services require more advanced networking features than the user network can offer, such as bridging, routing, or tunneling.
References:
* Documentation/Networking - QEMU
* QEMU/Networking - Wikibooks, open books for an open world
* Border Gateway Protocol - Wikipedia
* Common Internet File System - Wikipedia
* AppleTalk - Wikipedia

Valid 305-300 Dumps shared by TrainingDump.com for Helping Passing 305-300 Exam! TrainingDump.com now offer the **newest 305-300 exam dumps**, the TrainingDump.com 305-300 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingDump.com 305-300 dumps with Test Engine here: https://www.trainingdump.com/Lpi/305-300-practice-exam-dumps.html (**125** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 32
Which file in acgroupdirectory contains the list of processes belonging to thiscgroup?
A. pids
B. members
C. procs
D. casks
E. subjects
Answer: C (LEAVE A REPLY)
Explanation
The file procs in a cgroup directory contains the list of processes belonging to this cgroup. Each line in the file shows the PID of a process that is a member of the cgroup. A process can be moved to a cgroup by writing its PID into the cgroup's procs file. For example, to move the process with PID 24982 to the cgroup cg1, the following command can be used: echo 24982 > /sys/fs/cgroup/cg1/procs1. The file procs is different from the file tasks, which lists the threads belonging to the cgroup. The file procs can be used to move all threads in a thread group at once, while the file tasks can be used to move individual threads2.
References:
* Creating and organizing cgroups cgroup2 - GitHub Pages
* Control Groups - The Linux Kernel documentation

NEW QUESTION: 33
Which of the following values are valid in the type attribute of a<domain>element in a libvirt domain definition? (Choose two.)
A. proc
B. namespace
C. kvm
D. cgroup
E. lxc
Answer: C,E (LEAVE A REPLY)
Explanation
The type attribute of a <domain> element in a libvirt domain definition specifies the hypervisor used for running the domain. The allowed values are driver specific, but include "xen", "kvm", "hvf" (since 8.1.0 and QEMU 2.12), "qemu" and "lxc"1. Therefore, the valid values among the options are C. kvm and E. lxc. KVM stands for Kernel-based Virtual Machine, which is a full virtualization solution for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V)2. LXC stands for Linux Containers, which is an operating system-level virtualization method for running multiple isolated Linux systems (containers) on a single control host3. The other options are not valid values for the type attribute, asthey are either not hypervisors or not supported by libvirt.
References:http://libvirt.org/formatdomain.html

https://libvirt.org/formatcaps.html

NEW QUESTION: 34
Which of the following devices exist by default in an LXC container? (Choose three.)
A. /dev/log
B. /dev/kmem
C. /dev/urandom
D. /dev/console
E. /dev/root
Answer: A,C,D (LEAVE A REPLY)

NEW QUESTION: 35
Which of the following values would be valid in the FROM statement in aDockerfile?
A. registry:ubuntu:focal
B. ubuntu:focal
C. http://docker.example.com/images/ubuntu-focal.iso
D. file:/tmp/ubuntu/Dockerfile
E. docker://ubuntu: focal
Answer: B (LEAVE A REPLY)

NEW QUESTION: 36
FILL BLANK
What LXC command lists containers sorted by their CPU, block I/O or memory consumption? (Specify ONLY the command without any path or parameters.)
Answer:
lxc-top
Explanation
LXD supports the following network interface types for containers: macvlan, bridged, physical, sriov, and ovn1. Macvlan creates a virtual interface on the host that is connected to the same network as the parent interface2. Bridged connects the container to a network bridge that acts as a virtual switch3. Physical attaches the container to a physical network interface on the host2. Ipsec and wifi are not valid network interface types for LXD containers. References:
* 1: Bridge network - Canonical LXD documentation
* 2: How to create a network - Canonical LXD documentation
* 4: LXD containers and networking with static IP - Super User

NEW QUESTION: 37
In an IaaS cloud, what is a common method for provisioning new computing instances with an operating system and software?
A. Each new instance is created based on an image file that contains the operating system as well as software and default configuration for a given purpose.
B. Each new instance is connected to the installation media of a Linux distribution and provides access to the installer by logging in via SSH.

C. Each new instance is a clone of another currently running instance that includes all the software, data and state of the original instance.
D. Each new instance contains a minimal live system running from a virtual CD as the basis from which the administrator deploys the target operating system.
E. Each new instance is connected via a VPN with the computer that started the provisioning and tries to PXE boot from that machine.
Answer: (SHOW ANSWER)

NEW QUESTION: 38
Which file format is used by libvirt to store configuration data?
A. INI-style text files
B. SQLite databases
C. XML files
D. Java-like properties files
E. Text files containing key/value pairs
Answer: (SHOW ANSWER)
Explanation
Libvirt uses XML files to store configuration data for objects in the libvirt API, such as domains, networks, storage, etc. This allows for ease of extension in future releases and validation of documents prior to usage.
Libvirt does not use any of the other file formats listed in the question. References:
* libvirt: XML Format
* LPIC-3 Virtualization and Containerization: Topic 305.1: Virtualization Concepts and Theory

NEW QUESTION: 39
Which of the following statements in aDockerfileleads to a container which outputs hello world? (Choose two.)
A. ENTRYPOINT "echo Hello World"
B. ENTRYPOINT ["echo hello world"]
C. ENTRYPOINT ["echo", "hello", "world"]
D. ENTRYPOINT "echo", "Hello", "World*
E. ENTRYPOINT echo Hello World
Answer: (SHOW ANSWER)

NEW QUESTION: 40
Which of the following commands boots a QEMU virtual machine using hardware virtualization extensions?
A. qvirt -create -drive file=debian.img -cdrom debian.iso -m 1024 -boot d -driver hvm
B. vm -kvm -drive file=debian.img -cdrom debian.iso -m 1024 -boot d
C. qemu-hw -create -drive file=debian.img -cdrom debian.iso -m 1024 -boot d
D. qemu -accel kvm -drive file-debian.img -cdrom debian.iso -m 1024 -boot d
E. qvm start -vmx -drive file=debian.img -cdrom debian.iso -m 1024 -boot d
Answer: D (LEAVE A REPLY)
Explanation

The correct command to boot a QEMU virtual machine using hardware virtualization extensions is qemu -accel kvm -drive file-debian.img -cdrom debian.iso -m 1024 -boot d. This command uses the -accel option to specify the hardware accelerator to use, which in this case is kvm. KVM is a full virtualization solution for Linux on x86 hardware containing virtualization extensions (Intel VT or AMD-V)1. The -drive option specifies the disk image file to use, which in this case is debian.img. The -cdrom option specifies the ISO image file to use as a CD-ROM, which in this case is debian.iso. The -m option specifies the amount of memory to allocate to the virtualmachine, which in this case is 1024 MB. The -boot option specifies the boot order, which in this case is d, meaning to boot from the CD-ROM first. References: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/virtualization_deployment_and_ https://fedoraproject.org/wiki/Virtualization

NEW QUESTION: 41
Which of the following tasks are part of a hypervisor's responsibility? (Choose two.)
A. Isolate the virtual machines and prevent unauthorized access to resources of other virtual machines.
B. Create filesystems during the installation of new virtual machine quest operating systems.
C. Provide host-wide unique PIDs to the processes running inside the virtual machines in order to ease inter-process communication between virtual machines.
D. Map the resources of virtual machines to the resources of the host system.
E. Manage authentication to network services running inside a virtual machine.
Answer: (SHOW ANSWER)

NEW QUESTION: 42
What is the purpose of the commandvagrantinit?
A. It starts a Vagrant box.
B. It executes a provisioning tool in a running box.
C. It installs Vagrant on a Linux host.
D. It downloads a Vagrant box.
E. It creates a Vagrant configuration file.
Answer: E (LEAVE A REPLY)

Valid 305-300 Dumps shared by TrainingDump.com for Helping Passing 305-300 Exam! TrainingDump.com now offer the **newest 305-300 exam dumps**, the TrainingDump.com 305-300 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingDump.com 305-300 dumps with Test Engine here: https://www.trainingdump.com/Lpi/305-300-practice-exam-dumps.html (**125** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)