

Huawei.H13-527_V5.0.v2026-08-10.q28

Exam Code:	H13-527_V5.0
Exam Name:	HCIP-Cloud Computing V5.0
Certification Provider:	Huawei
Free Question Number:	28
Version:	v2026-08-10
# of views:	103
# of Questions views:	281
https://www.freecram.net/torrent/Huawei.H13-527_V5.0.v2026-08-10.q28.html	

NEW QUESTION: 1

In HUAWEI CLOUD Stack, which of the following is not a method that an administrator can use to access OBS?

- A. Using the console to perform visual operations
- B. RESTful APIs
- C. Mounting an OBS bucket to a local ECS
- D. SDKs such as Java and Python

Answer: (SHOW ANSWER)

After correcting the wording, the best verified answer is C . Standard Huawei OBS access methods include the console , RESTful APIs , and SDKs . Huawei OBS documentation and related Huawei service guides consistently describe OBS as being accessed through API-based methods and software-development toolkits, and many Huawei cloud services explicitly reference RESTful and SDK interaction with OBS resources.

Option C is different. Although some Huawei cloud-native scenarios allow OBS buckets to be mounted to workloads in CCE through OBS volumes, that is a specialized container-storage integration scenario, not the standard administrator access method for OBS itself. Huawei's CCE documentation also warns about constraints and resource overhead when OBS buckets are mounted to workloads, which reinforces that this is not the normal general-purpose way an administrator "accesses OBS." So for exam-style classification, A , B , and D are the normal access methods, while C is not the standard answer. Therefore, the correct answer is C .

=====

NEW QUESTION: 2

In HUAWEI CLOUD Stack, the weak affinity policy of a VM group schedules VMs in the same group to one host when resources are sufficient. If resources are insufficient, the rule can be broken.

- A. TRUE
- B. FALSE

Answer: (SHOW ANSWER)

The correct answer is A (TRUE) . Huawei Cloud Stack compute resource management documentation explicitly describes the behavior of soft-affinity , which is the practical equivalent of weak affinity in scheduling policy terms. Huawei states that if host resources are sufficient, VMs in the same soft-affinity group will be scheduled to the same host . It further states that if resources are insufficient or the current host is faulty, the VMs can be scheduled to different hosts . This matches the wording of the question almost exactly.

Huawei ECS group documentation also explains that an ECS group can be configured with affinity, anti- affinity, weak affinity, or weak anti-affinity policies, which confirms that weak affinity is an officially recognized scheduling policy in Huawei Cloud Stack. The "weak" part means the scheduler prefers a placement rule but is allowed to violate it when necessary to ensure successful deployment or continuity.

Therefore, the statement accurately reflects Huawei's documented scheduling behavior: under normal conditions, VMs tend to be placed together, but if resources are inadequate, the rule may be broken. So the verified answer is TRUE .

NEW QUESTION: 3

Which of the following is a typical working scenario of a polling agent in OpenStack?

- A. Reporting alarms for high CPU usage
- B. Querying VM memory usage every 10 minutes
- C. Listening to VM deletion events
- D. Storing event logs

Answer: (SHOW ANSWER)

The correct answer is B. Querying VM memory usage every 10 minutes . OpenStack Telemetry documentation states that the Telemetry service polls service APIs every 10 minutes by default , and that this interval can be changed in polling.yaml. That is exactly the behavior of a polling agent: it periodically queries metrics rather than waiting for events.

Ceilometer measurement documentation also includes metrics such as memory.usage, which shows that VM memory usage is one of the kinds of resource data that polling-based telemetry can collect, depending on configuration and hypervisor support.

The other options describe different telemetry-related functions. Reporting alarms is associated more with alarming services such as Aodh, listening to VM deletion events corresponds to event-driven notification collection rather than polling, and storing event logs is a storage or event-backend function rather than the polling agent's role. Therefore, the typical example of a polling agent's work is periodic metric collection, making B the fully verified answer. This also aligns with Huawei Cloud Stack monitoring and O & M concepts built on OpenStack telemetry principles.

NEW QUESTION: 4

Which of the following statements are true about the network resource management function of Service OM?

- A. Administrators can check the status of ports, for example, VM NIC ports.
- B. Administrators can create virtual networks and private line access points.
- C. Administrators can directly view the status of physical and virtual networks.
- D. Administrators can create physical networks.

Answer: (SHOW ANSWER)

The best verified answer is A, B, and C. Huawei's Managing Network Resources documentation states that on Service OM, administrators can view the number of created physical networks and virtual networks, which supports option C. The same Huawei material also says administrators can collect information for handling VM network faults, and related VPC/API guidance shows Service OM visibility into ports, including VM NIC-related ports, which supports A.

Huawei documentation for Direct Connect further shows that administrators can create Direct Connect access points on Service OM before tenants can use Direct Connect. In training-style wording, this aligns with the idea of creating private line access points, so B is also correct.

Option D is the one that does not fit as a standard Service OM network resource management action in the same way. Huawei documentation emphasizes viewing physical network information and managing related network resources, but creation or low-level configuration of physical networks is generally handled through platform network configuration workflows rather than as a routine Service OM tenant/resource function.

Therefore, the correct set is A, B, C.

=====

NEW QUESTION: 5

Which of the following is not a core concept of OpenStack?

- A. Flexibility
- B. Openness
- C. Efficiency
- D. Scalability

Answer: ([SHOW ANSWER](#))

OpenStack is widely recognized for its openness, flexibility, and scalability . Official OpenStack documentation describes OpenStack as an open-source cloud platform that allows organizations to integrate different technologies and build clouds with a high degree of flexibility. It also explicitly highlights horizontal scalability , meaning compute, storage, and network resources can be expanded as demand grows.

In addition, OpenStack and related OpenInfra materials consistently emphasize openness as a defining characteristic of the platform.

By contrast, efficiency is certainly a benefit that cloud platforms may deliver, but it is not typically presented as one of the foundational core concepts of OpenStack in the same way as openness, flexibility, and scalability. Huawei-authored OpenStack content also discusses OpenStack in terms of openness and scalability, which aligns with the standard industry description rather than listing efficiency as a core concept.

Therefore, among the options given, Efficiency is the one that is not considered a core concept of OpenStack.

=====

NEW QUESTION: 6

In HUAWEI CLOUD Stack, FusionCare is deployed on a single node by default and runs on the PUB-SRV-01 management VM.

- A. TRUE
- B. FALSE

Answer: ([SHOW ANSWER](#))

The correct answer is B (FALSE) . Huawei documentation confirms that FusionCare is a deployed component used for inspection and troubleshooting , but the statement's deployment location is not reliably correct as written. In Huawei Cloud Stack maintenance and troubleshooting materials, Huawei explicitly references logging in to FusionCare separately and also notes cases where administrators must obtain the IP address of the node where FusionCare is deployed , which shows it is treated as a specific deployable node/service rather than a fixed default identity tied universally to PUB-SRV-01 .

More importantly, Huawei documentation for certificate replacement states that administrators may need the IP address of the PUB-SRV03 node where FusionCare is deployed , which directly contradicts the claim that FusionCare runs on PUB-SRV-01 by default. Huawei O & M rollback documents also mention checking whether FusionCare is deployed in each secondary region , indicating that its deployment is not always a single fixed-node default in the simplistic way the statement suggests.

So, although FusionCare is indeed part of Huawei Cloud Stack O & M, the statement is inaccurate because the node/location claim is not universally true.

Therefore, the verified answer is FALSE .

=====

NEW QUESTION: 7

Both heat_template_version and resources are mandatory if you want to write a HOT template.

- A. TRUE
- B. FALSE

Answer: ([SHOW ANSWER](#))

The statement is TRUE according to the official HOT specification. OpenStack Heat documentation shows the overall HOT structure and explains that heat_template_version identifies the YAML document as a HOT template of a specific version. This key is fundamental because it tells Heat which template format and feature set to validate.

The same HOT specification states that the resources section contains the declaration of the resources in the template and that this section with at least one resource should be defined in any HOT template , otherwise the template would not really do anything when instantiated. In other words, a valid HOT template intended to deploy a stack needs both the version declaration and a resource definition.

Other sections such as description, parameters, outputs, and conditions are documented as optional. That is why heat_template_version and resources stand out as the core required parts for normal HOT authoring. In Huawei Cloud Stack training contexts based on OpenStack Heat, these are treated as the mandatory structural elements of a HOT template. Therefore, the verified answer is A (TRUE) .

=====

NEW QUESTION: 8

Which of the following tools in HUAWEI CLOUD Stack is mainly used to automate standardized O & M tasks and execute orchestration workflows?

A. CMDB

B. AutoOps

C. FusionCare

D. CloudNetDebug

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed 150 to 200 words of Explanation From HUAWEI CLOUD Computing Topics:

The correct answer is B. AutoOps . In HUAWEI CLOUD Stack, AutoOps is the automation platform used for standardized O & M task execution. Its core value is that it can take individual atomic operations and organize them into automated workflows, allowing repeated maintenance and troubleshooting procedures to be executed in a consistent, efficient, and less error-prone manner. This makes it highly suitable for routine operational tasks, fault handling, and platform maintenance scenarios where manual execution would be slower and more risky.

The other options have different roles. CMDB is mainly used for configuration and dependency information management. FusionCare is more focused on inspection, health check, and troubleshooting support.

CloudNetDebug is designed for network fault diagnosis, especially for virtual network path troubleshooting.

None of these primarily serve as the orchestration and automation engine for O & M workflows.

Therefore, when the requirement is to automate standardized operations and execute orchestration-based maintenance tasks in HUAWEI CLOUD Stack, the best answer is AutoOps , so the correct option is B .

NEW QUESTION: 9

In HUAWEI CLOUD Stack, the management of BMSs depends on the component Ironic , which manages physical servers throughout their lifecycle. (Fill in the blank.)

Answer:

Ironic

Explanation:

The correct fill-in answer is Ironic . Huawei's official HUAWEI CLOUD Stack documentation clearly states that Ironic provides BMS management services by working with components such as Nova and Neutron

. Huawei's BMS usage guide also explicitly refers to BMS lifecycle management operations , including actions such as starting, stopping, forcibly stopping, and other lifecycle tasks. This confirms that Ironic is the component responsible for managing physical servers through their lifecycle in a Huawei Cloud Stack environment.

This is consistent with standard OpenStack architecture as well, where Ironic is the OpenStack Bare Metal service used to provision and manage physical machines similarly to how Nova manages virtual machines. In Huawei Cloud Stack, BMS is not handled as a simple ECS variant; instead, it relies on Ironic for provisioning, state transitions, and lifecycle orchestration of dedicated hardware resources.

Therefore, after correcting the sentence, the blank should be filled with Ironic , and the fully verified answer is Ironic .

NEW QUESTION: 10

Which of the following is not a core value proposition of Horizon in OpenStack?

- A. API eventual consistency
- B. View consistency
- C. API backward compatibility
- D. Scalability

Answer: ([SHOW ANSWER](#))

Horizon is the web-based dashboard for OpenStack, designed to provide a consistent graphical interface for managing cloud resources. In practice, Horizon focuses on presenting a consistent view of services to users and administrators, and OpenStack components commonly maintain API backward compatibility to support upgrades and continued interoperability. OpenStack documentation also shows that Horizon can be used in larger and multi-region environments, which supports the idea that scalability is relevant to its operational value.

However, API eventual consistency is not a Horizon value proposition. "Eventual consistency" is a concept found in distributed storage systems such as OpenStack Swift/Object Storage, where updates may take time to become visible everywhere. OpenStack documentation explicitly uses this term in the context of Object Storage behavior, not Horizon dashboard design. Because Horizon is a management dashboard rather than a distributed storage consistency model, API eventual consistency does not belong as one of its core value propositions. Therefore, the correct answer is A . This is the best-supported conclusion from official OpenStack documentation and aligns with Huawei Cloud training-style OpenStack topic coverage.

NEW QUESTION: 11

Which of the following is not an advantage of BMS over ECS?

- A. Physical resource isolation for high security
- B. No virtualization layer, low performance loss
- C. Quick and elastic adjustment of CPUs and memory
- D. Direct access to physical hardware such as GPUs and FPGAs

Answer: ([SHOW ANSWER](#))

The correct answer is C. Quick and elastic adjustment of CPUs and memory . Huawei's official HUAWEI CLOUD Stack materials describe Bare Metal Server (BMS) as dedicated physical servers that provide excellent or remarkable computing performance and direct access to physical resources. Huawei also explains that Ironic works with Nova and Neutron to manage BMS resources, highlighting the physical- server nature of the service. These characteristics support advantages such as physical isolation , no hypervisor performance overhead , and direct access to specialized hardware .

By contrast, Huawei's ECS Advantages documentation explicitly says that with ECS, administrators can modify an ECS flavor, including the number of CPUs and memory size . That kind of quick and elastic CPU/memory adjustment is therefore an ECS advantage , not a BMS-over-ECS advantage.

So after correcting the option formatting, A , B , and D are all real BMS strengths, while C belongs more to virtualized ECS elasticity. Therefore, the verified answer is C .

=====

NEW QUESTION: 12

HUAWEI CLOUD Stack operation administrators and VDC administrators both can create custom approval processes, but only the approval processes created by operation administrators are globally visible.

A. TRUE

B. FALSE

Answer: ([SHOW ANSWER](#))

The statement is TRUE. Huawei Cloud Stack separates management responsibilities between operation administrators and VDC administrators. Huawei's service-management documentation shows that operation administrators work at the platform-wide level and can publish cloud services and associate them with approval processes, which reflects their broader administrative scope. By contrast, VDC administrators work within the boundary of their own VDC and manage service usage and permissions inside that tenant-like scope. This architecture strongly supports the rule that both roles can participate in approval-process customization, but the visibility scope is different.

In practice, this means an approval process created by an operation administrator can be used as a global process across the platform, while a process created by a VDC administrator is limited to the VDC context rather than being exposed to all tenants or all service domains. This aligns with Huawei Cloud Stack's layered governance model, where global service governance belongs to operation administrators and tenant-level governance belongs to VDC administrators. Therefore, the correct answer is A (TRUE).

=====

NEW QUESTION: 13

Which of the following OpenStack authentication methods uses a pair of private and public keys for signing and verification?

A. UUID

B. Local authentication

C. PKI

D. Fernet

Answer: ([SHOW ANSWER](#))

The correct answer is PKI. In OpenStack Keystone, PKI (Public Key Infrastructure) tokens are cryptographically signed and specifically rely on a public/private key pair. Official OpenStack Keystone documentation states that PKI tokens are signed documents and that the Identity service uses public/private key pairs and certificates to create and validate them. It also explains that PKI and PKIZ tokens can be verified offline using Keystone's public signing key, which directly matches the wording in the question about signing and verification.

By contrast, UUID tokens are simple persisted identifiers and do not inherently depend on public/private key signing. Fernet tokens are also a Keystone token format, but they are based on symmetric cryptography and shared key repositories rather than a public/private key pair. OpenStack documentation explicitly distinguishes Fernet and UUID from the deprecated PKI token mechanism. Therefore, among the listed options, the authentication method that uses a pair of private and public keys is clearly PKI. This aligns with standard OpenStack identity-service architecture and is consistent with Huawei Cloud Stack training topics that build on Keystone authentication concepts.

NEW QUESTION: 14

Which of the following data is best suited for storage in Gnocchi on OpenStack?

A. VM creation event logs

B. Network port configuration changes

C. User operation audit records

D. Changes of the VM CPU usage over time

Answer: ([SHOW ANSWER](#))

The correct answer is D. Changes of the VM CPU usage over time . OpenStack documentation describes Gnocchi as a time-series storage and resource indexing service . In the Telemetry pipeline documentation, OpenStack explicitly states that when the Gnocchi publisher is enabled, measurement and resource information is pushed to Gnocchi for time-series optimized storage . That makes Gnocchi the right component for storing metrics that change continuously over time, such as CPU utilization, memory usage, disk I/O, and similar performance measurements.

By contrast, the other options describe events, configuration changes, or audit-style records , which are not the primary use case for Gnocchi. OpenStack's telemetry retrieval guidance recommends storing metric data in Gnocchi, while event data should be stored in an event-oriented service such as Panko. This clearly separates time-series metrics from event or log records.

Therefore, in Huawei Cloud Stack and OpenStack telemetry scenarios, Gnocchi is best suited for historical metric trends , making VM CPU usage over time the most appropriate choice.

=====

NEW QUESTION: 15

Which of the following IT service models is most suitable for daily office work of small- and medium-sized enterprises, where enterprises do not need to pay attention to software installation and maintenance?

- A.** IaaS
- B.** MaaS
- C.** SaaS
- D.** PaaS

Answer: ([SHOW ANSWER](#))

The correct answer is C. SaaS. In the Software as a Service model, the cloud provider delivers a complete application that users access over the network, while the provider is responsible for installation, maintenance, upgrades, and much of the operational management. This makes SaaS especially suitable for daily office work, such as email, document collaboration, office suites, and business productivity applications.

Organizations, especially small- and medium-sized enterprises, can use the software directly without managing the underlying infrastructure or platform.

By comparison, IaaS mainly provides infrastructure resources such as compute, storage, and networking, so the enterprise still needs to manage operating systems and applications. PaaS provides an application development and runtime platform, which is useful for developers but not the best fit for general office users.

MaaS is not the standard answer in this classic cloud-service-model classification.

In Huawei Cloud computing training and in mainstream cloud architecture, SaaS is the service model that best matches a scenario where customers want to use software immediately and avoid the burden of software deployment and maintenance. Therefore, the verified answer is C. SaaS.

NEW QUESTION: 16

Which of the following is not a service plugin in Neutron?

- A.** L3 Service Plugin
- B.** ML2 Plugin
- C.** VPN Service Plugin
- D.** LB Service Plugin

Answer: ([SHOW ANSWER](#))

The correct answer is B. ML2 Plugin . In OpenStack Neutron, ML2 stands for Modular Layer 2 and is the main core networking plugin framework used to support different network types and mechanism drivers.

Official OpenStack documentation describes ML2 as a plug-in with components for network types and mechanisms, not as a service plugin .

Service plugins in Neutron are different. They add higher-level networking services beyond basic Layer 2 connectivity. OpenStack deployment documentation shows optional service areas such as load balancing and virtual private network services, and Neutron's router-related functionality is commonly associated with the L3 service plugin . These are examples of service plugins layered on top of the core networking plugin model.

So the key distinction is this: ML2 is the core plugin architecture , while L3, VPN, and load balancing belong to the service-plugin category. In Huawei Cloud Stack networking topic coverage, this separation between the core Neutron plugin and optional service plugins is also fundamental. Therefore, the option that is not a service plugin is B. ML2 Plugin .

=====

Valid H13-527_V5.0 Dumps shared by EduDump.com for Helping Passing H13-527_V5.0 Exam! EduDump.com now offer the **newest H13-527_V5.0 exam dumps**, the EduDump.com H13-527_V5.0 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com H13-527_V5.0 dumps with Test Engine here: https://www.edudump.com/exams/Huawei/H13-527_V5.0/premium/ (62 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 17

In HUAWEI CLOUD Stack, which of the following is the main function of Service OM?

- A. Infrastructure hardware monitoring
- B. Centralized cross-product monitoring
- C. Cloud platform deployment and upgrade
- D. Resource pool and basic cloud service management

Answer: (SHOW ANSWER)

The correct answer is D. Resource pool and basic cloud service management . Huawei's official License Guide includes a dedicated section called "Resource Pool & Basic Cloud Services" and explicitly instructs administrators to enter Service OM to perform management actions there. This directly ties Service OM to the management plane for resource pools and basic cloud services , rather than positioning it primarily as a hardware-monitoring tool or a deployment platform. Huawei Cloud Stack O & M materials also show Service OM being used for operational management tasks such as viewing VMs, host groups, cloud services, and system resource status. In contrast, deployment and upgrade are associated with dedicated installation and upgrade tools, not Service OM as the main function.

Similarly, centralized monitoring across products is broader than Service OM's core purpose and is covered by other monitoring and O & M capabilities in the Huawei Cloud Stack architecture.

Therefore, among the available choices, the description that best matches Huawei's official positioning of Service OM is resource pool and basic cloud service management . That makes D the most accurate and fully verified answer.

=====

NEW QUESTION: 18

Which of the following cannot be checked using HUAWEI CLOUD Stack FusionCare?

- A. Service OM
- B. eSight
- C. CCE
- D. Workspace

Answer: (SHOW ANSWER)

The best verified answer is D. Workspace . Huawei’s official health-check documentation for FusionCare shows that FusionCare supports unified health checks for multiple Huawei Cloud Stack components and services. The published health-check item lists explicitly include Service OM , eSight , and CCE and SWR , which confirms that those services are within FusionCare's supported inspection scope.

By contrast, the Huawei monitoring materials mention Workspace Overview under broader monitoring capabilities, not under the FusionCare-supported health-check list shown in the cited documents. In other words, Workspace appears in Huawei’s overall monitoring ecosystem, but it is not evidenced in the cited FusionCare support list the way Service OM, eSight, and CCE are. That makes Workspace the option that cannot be checked using FusionCare in this question set.

So based on Huawei’s documented FusionCare-supported inspection items, A , B , and C are covered, while D is not shown as a FusionCare check target. Therefore, the verified answer is D .

=====

NEW QUESTION: 19

Match the following Huawei Cloud Stack Cloud Configuration Service (CCS) components with their functions.

CCS API		Provide administrators with visualized CCS tagging.
CCS Console		Provide the service logic of CCS tagging.
CCS Client		Provide administrators with cloud service tagging using CCS command lines.

Answer:

CCS API	CCS Console	Provide administrators with visualized CCS tagging.
CCS Console	CCS API	Provide the service logic of CCS tagging.
CCS Client	CCS Client	Provide administrators with cloud service tagging using CCS command lines.

Explanation:

- * CCS API # Provide the service logic of CCS tagging
- * CCS Console # Provide administrators with visualized CCS tagging

* CCS Client # Provide administrators with cloud service tagging using CCS command lines Comprehensive and Detailed 150 to 200 words of Explanation

From HUAWEI CLOUD Computing Topics:

This drag-and-drop match follows the normal architectural roles of Huawei Cloud Stack CCS components.

The CCS API is the backend service interface, so it is responsible for the service logic of CCS tagging . In cloud platforms, APIs handle request processing, tag creation, tag query, and tag-management actions exposed to other components or clients. The CCS Console is the graphical management interface, so its role is to provide administrators with visualized CCS tagging . This means administrators can perform tag- related operations through a web-based UI rather than command lines. The CCS Client is the command-line access component, so it is used to provide administrators with cloud service tagging using CCS command lines .

This mapping is also consistent with standard Huawei Cloud Stack component design principles: API for backend logic , Console for visual management , and Client for CLI-based operations . In practical cloud O & M and resource governance, these three components complement each other by supporting programmatic integration, web-based administration, and command-line automation. Therefore, the verified matching is:


CCS API # service logic , CCS Console # visualized tagging , and CCS Client # command-line tagging .

NEW QUESTION: 20

Match the following Auto Scaling (AS) concepts with their descriptions in Huawei Cloud Stack.

AS group		Specify the conditions for triggering an AS action.
AS policy		A set of instances with the same application scenario
AS configuration		Define the specifications of the instance to be added.
AS action		Its purpose is to keep the number of instances in the application system the same as expected.

Answer:



Explanation:

* AS group # A set of instances with the same application scenario

* AS policy # Specify the conditions for triggering an AS action

* AS configuration # Define the specifications of the instance to be added

* AS action # Its purpose is to keep the number of instances in the application system the same as expected

Explanation From HUAWEI CLOUD Computing Topics:

This matching is consistent with Huawei's official Auto Scaling concept definitions. Huawei states that an AS group consists of a collection of instances that apply to the same application scenario , so that description belongs to AS group . Huawei also explains that an AS policy specifies the conditions for triggering a scaling action , which makes the second match straightforward.

Huawei's Auto Scaling documentation further distinguishes AS configuration from policy and action. The configuration defines the template or specifications for the instances that will be added during scaling, so define the specifications of the instance to be added correctly matches AS configuration . Finally, Huawei defines a scaling action as adding or removing ECS instances from an AS group, and states that its purpose is to keep the number of instances in the application system at the expected level. That directly matches the final description.

So the verified drag-and-drop mapping is: group = instance set , policy = trigger condition , configuration

= instance specification , and action = maintain expected instance count .

NEW QUESTION: 21

In HUAWEI CLOUD Stack, the maximum size of a Scalable File System (SFS) file is 240 TB. The capacity of a file system can be increased to _____. (Fill in the blank.)

Answer:

PB level

Explanation:

The correct fill-in answer is PB level . Huawei's official HUAWEI CLOUD Stack SFS documentation states that the maximum capacity of a single file is 240 TB , and it further explains that the file system capacity can be scaled to the PB level . This is the exact Huawei wording used to describe the upper expansion capability of SFS.

This reflects the design goal of SFS in Huawei Cloud Stack: to provide managed, shared file storage for applications that require large-scale file-system growth over time. In other words, the size of one individual file and the total expandable capacity of the file system are two different things. A single file has a stated upper limit, but the whole file system can continue to grow far beyond that limit, reaching petabyte-scale capacity. Therefore, after correcting the sentence structure, the blank should be filled with PB level, which is the official Huawei description of the scalable upper capacity for the file system.

=====

NEW QUESTION: 22

When registering an image on Service OM, if the image is bigger than 6 GB, select FTP or NFS for uploading the image. (Fill in the blank.)

Answer:

FTP

Explanation:

The correct fill-in answer is FTP. Huawei documentation for image registration on Service OM shows that images can be uploaded using different modes, including local upload, FTP, and NFS. The guidance indicates that for larger image files, administrators should use a network-based upload mode rather than simple local browser upload. In the Huawei document snippet, the upload options explicitly include NFS, and the paired option used in this context is FTP, which matches the exam sentence structure: "select ____ or NFS." This makes sense operationally in HUAWEI CLOUD Stack because large image files are more reliably transferred through shared file or server-based methods, especially in enterprise O & M scenarios. Local browser-based uploads are less suitable for very large images due to session stability, browser limits, and transfer reliability. Using FTP or NFS improves image registration efficiency and reduces the chance of interrupted uploads during image onboarding. Therefore, after correcting the sentence, the blank should be filled with FTP.

=====

NEW QUESTION: 23

Keystone users are assigned roles, and each role has different permissions.

A. TRUE

B. FALSE

Answer: (SHOW ANSWER)

In OpenStack, Keystone is the Identity service responsible for authentication and authorization. A core Keystone concept is that permissions are not usually granted directly to a user; instead, users are assigned one or more roles, and those roles are associated with a project, domain, or system scope. OpenStack's official Keystone documentation states that the Identity service assigns roles to users, and a user can have different roles in different projects. It also explains that default roles such as admin, member, and reader are used to control what actions a user can perform. This means the statement is correct: Keystone users are assigned roles, and those roles determine their permissions.

This is also consistent with Huawei Cloud's identity and permission model, where permissions are controlled through assigned authorization policies and identities use tokens containing identity and permission information when accessing services. In practical Huawei Cloud Stack and OpenStack-based environments, role-based access control is the standard method for defining what a user is allowed to do. Therefore, the correct answer is TRUE .

NEW QUESTION: 24

In HUAWEI CLOUD Stack, which of the following is a function of Destination Network Address Translation (DNAT)?

A. Converting the private IP address that an Elastic Cloud Server (ECS) uses to access the Internet

B. Sharing a private IP address among multiple servers

C. Filtering secure public network traffic to a VPC

D. Forwarding requests from public IP addresses to private IP addresses of servers in the cloud

Answer: D (LEAVE A REPLY)

The correct answer is D . Huawei's official HUAWEI CLOUD Stack NAT Gateway documentation explains that DNAT is used to map the private IP address and port of a cloud server in a VPC to a public IP address and port by binding an EIP. In practical terms, DNAT allows requests that arrive at a public IP address to be forwarded to the private IP address of a server inside the cloud network. That matches option D exactly.

Option A describes SNAT , not DNAT. Huawei explicitly states that SNAT translates a private IP address into an extranet/public IP address so cloud servers can access the Internet. Option B is also related to SNAT behavior, because Huawei notes that multiple cloud servers in a VPC can share an EIP for outbound Internet access through SNAT. Option C describes a filtering or firewall-style function rather than NAT translation.

Therefore, in Huawei Cloud Stack networking, DNAT is specifically used for inbound mapping from public addresses to private cloud-server addresses , so the verified answer is D .

=====

NEW QUESTION: 25

Which of the following components can be used to process Network Address Translation (NAT) and route forwarding in Neutron?

- A. L2 Agent
- B. L3 Agent
- C. NAT Agent
- D. DHCP Agent

Answer: (SHOW ANSWER)

The correct answer is B. L3 Agent . Official OpenStack Neutron documentation states that routers provide virtual Layer 3 services such as routing and NAT , and that the Networking service uses a layer-3 agent to manage routers. OpenStack also explicitly says that the neutron-l3-agent uses the Linux IP stack and iptables to perform L3 forwarding and NAT . This directly matches the wording of the question.

The other options do not fit. The L2 agent handles local virtual switch configuration and Layer 2 networking behavior, not routing and NAT. The DHCP agent provides DHCP services to tenant networks. " NAT Agent

" is not the standard Neutron component name used in official OpenStack architecture. OpenStack documentation consistently identifies the L3 agent as the component responsible for router namespace management, forwarding, floating IP behavior, and source NAT functions.

This is also consistent with Huawei Cloud Stack networking architecture topics, where routing and NAT functions are tied to Neutron Layer 3 services.

Therefore, the fully verified answer is B. L3 Agent .

NEW QUESTION: 26

In HUAWEI CLOUD Stack, which of the following features reflect the elasticity of ECSs?

- A. Quick creation of instances based on IMS images
- B. Automatic scaling of instances in combination with AS
- C. Specification downgrade (for example, by reducing the number of CPU cores) to reduce costs
- D. Online upgrade of CPUs and memory
- E. Resource isolation at the physical server level

Answer: (SHOW ANSWER)

The correct answers are A, B, and D. Huawei's official ECS documentation highlights several elasticity- related capabilities. It states that Auto Scaling (AS) can automatically add or remove ECSs based on service demand, which directly supports option B. Huawei also states that ECS flavors can be modified, including changes to CPU and memory, which supports D as an elastic scaling feature. In addition, Huawei's IMS documentation explains that images can be used to create ECSs in batches and, together with AS, support dynamic scale-out, which makes A a valid reflection of ECS elasticity as well.

Option C is not well supported as stated, because Huawei documentation emphasizes flavor modification and online flavor change mainly in the context of scaling resources, especially increasing CPUs and memory, not

"specification downgrade" as a core elasticity feature. Option E refers to physical isolation, which is a BMS characteristic rather than an ECS elasticity feature. Therefore, the best verified answer is A, B, D.

NEW QUESTION: 27

In OpenStack, HOT and CFN templates both support YAML. CFN templates are more suitable for OpenStack deployments.

A. TRUE

B. FALSE

Answer: ([SHOW ANSWER](#))

The statement is FALSE . Official OpenStack Heat documentation says that when the Orchestration service was originally introduced, it worked with AWS CloudFormation templates , which are in JSON format. It then states that Heat also runs HOT templates written in YAML . This directly disproves the claim that both HOT and CFN templates support YAML in the same native sense.

OpenStack Heat documentation also explains that HOT was introduced to replace the Heat CloudFormation-compatible format (CFN) as the native format supported by Heat over time. That means HOT, not CFN, is the format more suitable for OpenStack-native deployments. The documentation further notes that YAML is easier to write, parse, and maintain, which is one reason HOT became the preferred OpenStack format.

Although internal parsing code shows the API can accept a JSON- or YAML-formatted document in some contexts, the OpenStack-native orchestration model clearly favors HOT and identifies CFN as the older CloudFormation-compatible format. So the combined statement is incorrect, and the correct answer is B (FALSE) .

=====

NEW QUESTION: 28

Which of the following services is most suitable for encrypting connections between cloud resources and local data centers of small- and medium-sized enterprises?

A. EIP

B. VPCEP

C. Direct Connect

D. Virtual Private Network (VPN)

Answer: D ([LEAVE A REPLY](#))

The correct answer is D. Virtual Private Network (VPN) . Huawei's official HUAWEI CLOUD Stack VPN documentation states that a VPN connection is an encrypted communications channel established between the VPN gateway in a VPC and the gateway in an external or local data center. Huawei also describes VPN as a way to connect a local data center to cloud resources through a secure and reliable encrypted tunnel over the Internet . This directly matches the requirement in the question for encrypting connections between cloud resources and local data centers.

By contrast, Direct Connect provides a dedicated, high-speed, low-latency connection, but the cited Huawei description emphasizes dedicated connectivity and stability rather than the VPN-style encrypted tunnel itself.

EIP is only a public IP addressing service, and VPCEP is used for private access to services, not site-to-cloud encrypted interconnection. For small- and medium-sized enterprises that want a practical, secure, and relatively simple hybrid-cloud connection, Huawei documentation supports VPN as the best fit. Therefore, the verified answer is D .

H13-527_V5.0 dumps with Test Engine here: https://www.edudump.com/exams/Huawei/H13-527_V5.0/premium/ (62 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)