PassQueen
a.k.a dumpsqueen.com

Huawei Certified Network Professional –Cloud DataCentre Operations

Huawei H31-522

Version Demo

Total Demo Questions: 28

Total Premium Questions: 288

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Data Center O&M Overview	93
Topic 2, Cloud Data Center O&M Service	35
Topic 3, Cloud Data Center O&M Process	18
Topic 4, Cloud Data Center O&M Management	28
Topic 5, Cloud Data Center O&M Tools	112
Topic 6, Mix Questions	2
Total	288

QUESTION NO: 1

Which of the following options are included on the Alarm Settings page of OperationCenter? (Multiple Choice)

A. Alarm Masking

B. Alarm Severity Redefinition

C. Alarm Notification

D. Alarm Transfer to Work Order

ANSWER: A B C

Explanation:

OperationCenter's Alarm Settings page provides configuration controls that determine how alarms are handled, interpreted, and communicated to operations personnel. **Alarm Masking** is included because administrators need to suppress selected alarms during planned maintenance, troubleshooting, or other periods in which known events should not generate operational noise. **Alarm Severity Redefinition** is also part of alarm settings, allowing the displayed severity of specified alarm types to be adjusted so that alarm prioritization reflects the organization's operational impact assessment. **Alarm Notification** belongs on this page because it defines how relevant alarm information is delivered to recipients, supporting timely awareness and response when alarm conditions occur. Together, these functions support the normal alarm-management lifecycle: reducing unnecessary alerts, ensuring the urgency classification is meaningful, and notifying the appropriate personnel. Huawei's enterprise support portal provides product documentation and operational guidance for ManageOne and OperationCenter features at Huawei ManageOne Support. General Huawei Cloud Stack documentation resources are also available through Huawei Cloud Stack Support.

QUESTION NO: 2

Which of the following are included in the management firmware upgrade process of the eSight server? (Multiple Choice)

A. Download the upgrade package

B. Automatically download the policy configuration

C. Upload the upgrade package

D. Perform the upgrade task

ANSWER: A B C D

Explanation:

The eSight firmware upgrade workflow includes obtaining the required upgrade software, importing it into the eSight environment, preparing the upgrade policy information, and running the upgrade operation. Therefore, **Download the upgrade package**, **Automatically download the policy configuration**, **Upload the upgrade package**, and **Perform the upgrade task** are all included in the overall management process. Downloading ensures that the appropriate, supported package is available from Huawei's software distribution channel. Uploading makes that package available to eSight for subsequent deployment and task execution. Policy configuration retrieval supplies the upgrade rules and associated configuration needed for controlled management, while the upgrade task is the execution stage in which eSight applies the selected firmware package to managed devices according to the configured policy. Together, these steps provide a managed and repeatable upgrade process rather than a manual device-by-device update. Huawei's eSight product documentation and software support resources provide the applicable release-specific procedures and upgrade packages: Huawei eSight support documentation and Huawei Enterprise software download center.

QUESTION NO: 3

Which of the following are the mainstream visualization platforms that can be managed by

ManageOne? (Multiple Choice)

A. FusionCompute

B. VMware vSphere

C. Hyper-V

D. FusionSphere OpenStack

ANSWER: A B D

Explanation:

ManageOne provides unified cloud-management capabilities across Huawei virtualization and cloud platforms as well as supported third-party virtualization environments. **FusionCompute** is Huawei's virtualization platform and is a core managed resource type in ManageOne deployments. ManageOne can connect to FusionCompute to orchestrate virtual-machine provisioning, manage compute and storage resources, monitor infrastructure, and expose services through a unified portal and service catalog.

VMware vSphere is also supported as a mainstream third-party virtualization environment. Through integration with vCenter Server, ManageOne can incorporate VMware-based resource pools into centralized service provisioning and operations management. This enables an organization to retain VMware infrastructure while managing cloud services alongside Huawei infrastructure.

FusionSphere OpenStack is Huawei's OpenStack-based cloud platform and is likewise managed by ManageOne. Its integration supports cloud-resource operations and tenant-facing service management across OpenStack domains. These supported platforms reflect ManageOne's role as a heterogeneous cloud-management layer, allowing consistent governance and lifecycle management across the relevant virtualization and cloud-resource domains. Huawei's ManageOne product information is available at Huawei ManageOne, and Huawei's FusionCompute documentation portal is available at Huawei FusionCompute Support.

QUESTION NO: 4

Which of the following are configuration items of eSight-OS? (Multiple Choice)

A. Disable the firewall in the operating system

B. Activate the Windows

C. Install the security patch of the operating system

D. Configure the network

ANSWER: A B C D

Explanation:

eSight-OS preparation includes making the underlying operating system ready for stable management-platform deployment and communication. "Activate the Windows" is a required operating-system readiness task where eSight is deployed on Windows, ensuring that the OS is properly licensed and can receive supported system services and updates. "Install the security patch of the operating system" is also an operating-system configuration item: current security patches reduce exposure to known vulnerabilities before the management system is placed into service. "Configure the network" is essential because eSight must have valid network parameters and reachability to managed devices, servers, and related services.

"Disable the firewall in the operating system" is included in the traditional eSight OS deployment checklist because an enabled host firewall can block the listening ports and service communications needed during installation and operation. In production, this should be implemented in accordance with the organization's security policy; where feasible, tightly scoped firewall rules for required eSight ports provide a safer equivalent to a broadly disabled firewall. These activities are all part of preparing the host OS and connectivity environment for eSight deployment. Huawei's eSight product and documentation resources are available at Huawei Enterprise Support: eSight and Huawei eSight product page.

QUESTION NO: 5

Which of the following is not involved in the service design management?

A. Service catalog management

B. Service level management

C. Supplier management

D. Financial management

ANSWER: D

Explanation:

Financial management is the correct answer because, in the ITIL v3 lifecycle model on which this terminology is based, Financial Management for IT Services belongs to the *Service Strategy* stage rather than Service Design. Its purpose is to support decisions about service funding, accounting, budgeting, charging, demand-related costs, and the financial value of IT services. These activities help an organization determine what services should be offered, how they should be funded, and whether they deliver appropriate business value before detailed service-design work begins.

Service Design uses the strategic direction and financial constraints established through service strategy to design services that can be transitioned and operated effectively. Financial Management therefore provides an important business input to the lifecycle, but it is not classified as a Service Design management process in the ITIL v3 process set. This distinction reflects ITIL's lifecycle separation between defining the business and financial strategy for services and designing the service solution, supporting processes, technology, suppliers, measurements, and governance arrangements. For ITIL lifecycle context, see AXELOS ITIL Service Management and the PeopleCert ITIL certification overview.

QUESTION NO: 6

To install FusionSphere OpenStack, the host CPU virtualization function must be enabled or else the VMs cannot be created.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. FusionSphere OpenStack relies on the host server's hardware-assisted virtualization capability—typically Intel VT-x or AMD-V—to run virtual machines through the KVM-based compute virtualization layer. This processor feature must be enabled in the server BIOS or UEFI before the compute host is deployed. When it is disabled, the host cannot provide the hardware virtualization support required for normal VM instantiation and execution, so virtual machines cannot be successfully created and run on that host. Enabling CPU virtualization is therefore a fundamental pre-installation and host-readiness requirement, alongside verifying compatible CPU architecture, memory, storage, and network resources. Huawei's FusionSphere OpenStack documentation identifies virtualization-capability checks as part of the environment and installation prerequisites for compute resources. For product documentation and deployment guidance, see the Huawei FusionSphere OpenStack product documentation and Huawei's FusionSphere OpenStack documentation portal.

QUESTION NO: 7

Both asynchronous and synchronous remote replication modes have no distance limitations.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct. Remote replication is constrained by the network characteristics between the primary and secondary storage systems, and synchronous replication is particularly sensitive to distance. In synchronous mode, a host write is not acknowledged as complete until the replicated data has also been committed at the remote site. The round-trip latency of the intersite link is therefore added to write I/O latency. As distance increases, propagation delay increases as well, which can materially affect application performance and limits the practical deployment distance for synchronous replication.

Asynchronous replication does not place the remote-write acknowledgement directly in the host write path. It can consequently support substantially longer distances and tolerate higher latency than synchronous replication. However, it still depends on available bandwidth, link stability, replication-cycle design, and the acceptable recovery point objective. It is not accurate to state that both modes have no distance limitations. Huawei storage replication guidance distinguishes synchronous and asynchronous replication according to latency, distance, and service requirements. See Huawei's OceanStor Remote Replication documentation and the Remote Replication overview.

QUESTION NO: 8

Which of the following statements is incorrect about service combination management?

A. In the service lifecycle, tracking investment works with other service management processes to ensure appropriate investment return.

B. It clarifies the service definition and ensures service achievements, thereby all design, transition, and operation activities are consistent with service values.

C. It builds an official complaint and upgrade procedure for customers.

D. Service combination is clear, and service requirements satisfied and service achievements supported by each service are clarified.

ANSWER: C

Explanation:

"It builds an official complaint and upgrade procedure for customers." is the incorrect statement. In this context, service combination management corresponds to service portfolio management: the discipline that governs an organization's services as investments and ensures that the set of services delivers intended business value. Its scope includes defining services, assessing their value and demand, balancing investment and risk, making lifecycle decisions, and maintaining a clear view of the services that are planned, available, or being retired. This governance enables service design, transition, and operation work to remain aligned with the value expected from each service.

A formal customer-complaint and escalation procedure is instead an operational customer-support and service-management control. It is normally associated with practices such as service desk, incident management, complaint handling, or customer relationship management. Although feedback and complaints can provide useful inputs when services are reviewed, establishing the procedure itself is not a core responsibility of service portfolio management. The distinction is important because portfolio decisions focus on selecting, prioritizing, funding, and governing services, whereas complaint escalation focuses on handling specific customer interactions and restoring satisfaction. See the IT Process Maps description of Service Portfolio Management and Atlassian's overview of service request management.

QUESTION NO: 9

ServiceCenter can be deployed on a VM provided by the FusionSphere OpenStack platform. ServiceCenter can be deployed in single-node mode or active/standby mode.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. Huawei ServiceCenter is a component used in FusionSphere/OpenStack-based cloud data-center environments to provide service management capabilities. Its deployment is supported on virtual machines supplied by the FusionSphere OpenStack resource pool, rather than requiring dedicated physical servers. This lets the component benefit from the cloud platform's standard VM lifecycle and resource-management capabilities.

The stated deployment models are also valid. A single-node deployment is suitable for smaller environments, lab scenarios, or cases in which high availability is not required. For production environments that need service continuity, ServiceCenter supports an active/standby deployment architecture. In this model, the standby instance is available to take over services when the active instance becomes unavailable, improving availability and reducing the impact of an individual VM or software failure. This aligns with Huawei's general cloud-platform design approach of providing both simplified single-instance deployments and redundant high-availability deployments according to the required reliability level. Huawei's FusionSphere documentation portal is available at Huawei FusionSphere OpenStack Support, and Huawei's cloud documentation resources are available through Huawei Enterprise Documentation.

QUESTION NO: 10

Which of the following are represented by AAA? (Multiple Choice)

A. Authentication

B. Authorization

C. Accounting

D. Audit

ANSWER: A B C

Explanation:

AAA is the standard access-control framework consisting of Authentication, Authorization, and Accounting. Authentication establishes the identity of a user or device attempting to access network resources, commonly by validating credentials such as a username and password, certificate, or other supported authentication method. Authorization is performed after identity verification and determines the services, privileges, command levels, VLANs, or other resource permissions that the authenticated entity may use. Accounting records session and usage information, such as login and logout times, session duration, commands, traffic consumption, or other auditable access records. Together, these functions enable centralized and policy-based control of administrative and user access, especially when implemented through protocols such as RADIUS or TACACS+. This three-part definition is the accepted AAA architecture used throughout enterprise networking, including Huawei device access-management configurations. The IETF AAA framework describes these functions and their role in network access control in RFC 3539. RADIUS, a commonly deployed protocol for carrying authentication, authorization, and accounting information, is specified in RFC 2865.

QUESTION NO: 11

OperationCentercan monitor VM alarms.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In Huawei cloud data-center operations, OperationCenter provides centralized operations and maintenance monitoring for managed virtualized resources, including virtual machines. VM alarm monitoring is a core operational

capability: OperationCenter collects and displays alarm information so administrators can identify faults or abnormal conditions affecting VM services, assess their severity and status, and take appropriate corrective action. Centralizing these alarms avoids relying on each individual virtualization component or VM owner to discover issues independently, and supports faster event handling in a large cloud data center. In practice, this monitoring function is used alongside resource views, alarm filtering, and alarm-management workflows to maintain service availability and operational visibility. The capability aligns with Huawei's documented FusionSphere/OpenStack operations-management scope, where the platform integrates management and monitoring of cloud infrastructure resources and their associated alarms. Huawei's enterprise support portal provides product documentation and maintenance resources for FusionSphere OpenStack at Huawei FusionSphere OpenStack Support. Additional Huawei cloud-computing product documentation and operational resources are available through the Huawei Cloud Computing Support portal.

QUESTION NO: 12

Which of the following are included in the performance monitoring function of eSight? (Multiple Choice)

A. Real-time monitoring

B. Historical monitoring

C. In-band monitoring

D. Out-band monitoring

ANSWER: A B

Explanation:

eSight performance monitoring includes **Real-time monitoring** and **Historical monitoring**. Real-time monitoring provides current performance information for managed devices and services, such as resource utilization, interface traffic, response-related indicators, and other collected KPIs. This allows operations personnel to observe the present operating condition of the data center infrastructure and react promptly when performance thresholds or alarms indicate an emerging issue.

Historical monitoring retains and presents collected performance data over time. It supports trend analysis, comparison across time periods, capacity evaluation, and the generation of performance reports. By examining historical KPI changes, administrators can identify recurring peaks, forecast resource demand, and make evidence-based optimization decisions. Together, current visibility and long-term trend data are the core perspectives required for effective performance management in eSight. Huawei describes eSight as a unified operations and maintenance platform with monitoring and management capabilities for enterprise infrastructure; see the Huawei eSight product page and the Huawei Enterprise Support documentation portal.

QUESTION NO: 13

Which of the following are benefits of using a dedicated management network in a cloud data center? (Multiple Choice)

A. Separates management traffic from service traffic.

B. Reduces exposure of management interfaces.

C. Facilitates fault isolation for management communication.

D. Eliminates the need for access control.

ANSWER: A B C

Explanation:

A dedicated management network separates management traffic from tenant and service traffic. This separation improves security by limiting exposure of management interfaces, improves reliability by reducing contention with service traffic, and makes fault isolation easier because management-plane communication can be independently monitored and troubleshot.

A management network does not remove the need for access control. Administrative interfaces still require secure authentication, authorization, logging, and appropriate network security policies. The AAA framework is described in RFC 3539.

QUESTION NO: 14

Which of the following information should be included in an incident record? (Multiple Choice)

A. Affected service or resource

B. Incident symptoms and service impact

C. Handling actions and resolution result

D. Delete the incident record immediately after restoration.

ANSWER: A B C

Explanation:

An incident record should include the affected service or resource, the incident symptoms and impact, and the handling actions and resolution result. This information enables operations personnel to understand the scope of the fault, coordinate escalation, record restoration activities, and retain knowledge for future troubleshooting.

Recording the incident priority or severity is also necessary for managing response targets and escalation. Deleting the record immediately after service restoration is not appropriate because historical records are required for audit, trend analysis, problem management, and continual improvement. General incident-management practices are described by ITIL.

QUESTION NO: 15

Which of the following steps are included in the Ope ration Center software package installation process? (Multiple Choice)

A. Obtain the software package.

B. Install OperationCenter.

C. Perform OS security hardening.

D. Verify the installation.

ANSWER: A B C D

Explanation:

The OperationCenter deployment workflow includes obtaining the required software package, installing OperationCenter, performing operating-system security hardening, and verifying the completed installation. Installation begins with acquiring the approved package and associated materials from Huawei's software-delivery channel so that the deployed version and its dependencies match the planned environment. The installation step then deploys the OperationCenter components and applies the required configuration for the management platform.

Operating-system security hardening is also an installation-process activity because OperationCenter is deployed on supporting hosts whose baseline security configuration must meet the platform's operational and security requirements before production use. This includes applying the prescribed OS-level controls and maintaining the secure operating environment on which the service runs. Finally, installation verification confirms that services are started correctly, required components are reachable, and the platform can be accessed and operated as expected. These steps form an end-to-end process: obtain trusted installation media, deploy the software, establish the secure host baseline, and validate service availability. Huawei's cloud-computing product documentation and enterprise support portal provide the applicable installation guides, release packages, and deployment prerequisites: Huawei FusionSphere OpenStack Support and Huawei Enterprise Support.

QUESTION NO: 16

Which of the following statements about the INTEL hardware virtualization is incorrect?

A. VT-x refers to the hardware-assisted virtualization technology of INTEL.

B. VT-d refers to the I/O assisted virtualization technology.

C. VT-c refers to the network-assisted virtualization technology.

D. SR-IOV refers to the VMDq technology.

ANSWER: D

Explanation:

“SR-IOV refers to the VMDq technology.” is the incorrect statement. Single Root I/O Virtualization (SR-IOV) and Virtual Machine Device Queues (VMDq) are separate Ethernet virtualization capabilities, even though both can be used to improve network I/O for virtual machines. SR-IOV is a PCI Express specification that enables one physical PCIe device, such as a network adapter, to expose a physical function and multiple lightweight virtual functions. Those virtual functions can be assigned directly to VMs, reducing hypervisor involvement in the I/O data path. VMDq, by contrast, is a NIC receive-side classification and queuing capability that directs packets into separate queues associated with virtual machines or virtual interfaces. Therefore, SR-IOV does not mean, define, or “refer to” VMDq; the technologies have different architectures and purposes. Intel’s overview of virtualization technologies describes hardware-assisted virtualization capabilities, while the PCI-SIG specification overview identifies SR-IOV as an I/O virtualization mechanism: Intel Virtualization Technology and PCI-SIG Single Root I/O Virtualization.

QUESTION NO: 17

Which of the following functions is used to automatically restart a VM on another host when the host running the VM fails?

A. HA

B. Live migration

C. DRS

D. Snapshot

ANSWER: A

Explanation:

HA is correct. High Availability (HA) monitors host status in a virtualization cluster. When a host fault is detected, HA attempts to restart the VMs that were running on the failed host on other available hosts in the cluster. This reduces service interruption caused by a physical host failure.

Live migration is used to move a running VM during planned maintenance or load balancing and normally requires the source host to remain available. DRS is used for resource scheduling, and snapshot is a point-in-time data protection mechanism. Huawei FusionCompute documentation and product resources are available from Huawei FusionCompute Support.

QUESTION NO: 18

On which of the following pages, client logout upon timeout can be disabled?

A. User management

B. Account policy

C. Idle timeout settings

D. User session inspection

ANSWER: C

Explanation:

Idle timeout settings is correct because client logout caused by inactivity is controlled through the management system's session or idle-timeout configuration. This page governs how long a logged-in client may remain inactive before the system automatically terminates the session. Disabling timeout-based logout therefore requires changing the idle-timeout behavior, typically by disabling the timeout feature or configuring the applicable idle-timeout value according to the platform's permitted settings.

This is a session-management control rather than a user-identity or authorization setting. It applies to the client's active web session and is intended to centrally manage the duration of inactivity that is accepted before logout. Such controls are part of the operational security configuration used by cloud data-center management platforms to balance administrator convenience with protection against unattended authenticated sessions. Huawei's enterprise support portal provides product documentation and configuration guidance for these management functions: Huawei Enterprise Support. Huawei also describes its cloud-computing and data-center management portfolio at Huawei Cloud Computing & Data Center.

QUESTION NO: 19

Which of the following actions should be performed before a planned maintenance window for a cloud service? (Multiple Choice)

A. Assess the service impact and risks.

B. Prepare and verify a rollback plan.

C. Notify affected users and service owners.

D. Perform the maintenance without creating a change record.

ANSWER: A B C

Explanation:

Before planned maintenance, operations personnel should assess service impact, prepare a rollback plan, and notify affected users or service owners. An impact assessment identifies dependent services, affected resource pools, risks, and the expected duration of the maintenance activity. A rollback plan provides a controlled method to restore the previous service state if the implementation fails or produces an unacceptable impact.

Notification enables users and stakeholders to prepare for the service interruption or operational restriction. Performing maintenance without recording the activity is not appropriate because the change record provides traceability, approval evidence, implementation details, and closure information. These practices align with controlled change enablement and service operations. See ITIL best-practice resources and Huawei Enterprise Support.

QUESTION NO: 20

FusionCompute consists of CNA and VRM. VRM uses VNA deployed on CNA nodes to manage al virtualization resources.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In Huawei FusionCompute, the Compute Node Agent (CNA) is installed on each physical compute server and supplies the CPU, memory, storage, and network capabilities used by virtual machines. The Virtual Resource Manager (VRM) is the centralized management component: it performs functions such as resource scheduling, lifecycle management, monitoring, and configuration management across the virtualization environment.

VRM does not directly manipulate the local virtualization capabilities of every compute server. Instead, it communicates with the Virtualization Node Agent (VNA) deployed on CNA hosts. The VNA acts as the local management agent and provides the interface through which VRM obtains host status and carries out management operations on CNA-side virtualization resources. This controller-and-agent design lets a VRM centrally manage a pool of CNAs while the host-level agent executes or reports operations locally. Therefore, the statement correctly describes the relationship among VRM, CNA, and VNA in the FusionCompute architecture. Huawei identifies FusionCompute as its virtualization platform for centrally managing compute resources and virtual machines; see the Huawei FusionCompute product page and the Huawei FusionCompute support portal.

QUESTION NO: 21

Which of the following methods can be used to create VMs in FusionCompute? (Multiple Choice)

A. Create a bare VM.

B. Create a VM using an image file.

C. Create a VM using a template.

D. Use an existing VM to clone a VM.

ANSWER: A B C D

Explanation:

FusionCompute supports several VM-provisioning paths to suit different operational needs. **Create a bare VM.** is the standard manual deployment method: an administrator defines the VM's compute, storage, and network parameters and can subsequently install an operating system. **Create a VM using an image file.** is also supported after an appropriate image has been imported into the image library; this enables operating-system or application images to be used as the basis for a new VM. **Create a VM using a template.** provides repeatable, standardized deployment from a prepared VM configuration and disk image. **Use an existing VM to clone a VM.** creates another VM based on an existing source VM, which is useful when a deployed workload must be replicated quickly. These methods collectively cover manual provisioning, image-based deployment, template-based deployment, and replication of an existing VM. Huawei FusionCompute documentation describes VM lifecycle and provisioning functions within its virtualization-management capabilities; see Huawei FusionCompute product documentation and Huawei FusionCompute overview.

QUESTION NO: 22

Which parameters define a backup policy? (Multiple Choice)

A. Backup type

B. Backup frequency

C. Backup window

D. Retention period

ANSWER: A B C D

Explanation:

A backup policy defines both what form of backup is created and the schedule and lifecycle governing that backup. **Backup type** specifies the backup method or category to apply, which determines how protected copies are created. **Backup frequency** establishes the recurrence interval for automatic backup jobs, ensuring that protection occurs regularly rather

than only through manual intervention. **Backup window** defines the permitted time period in which scheduled backup processing can start or run, allowing backup activity to be aligned with lower-demand operational periods. **Retention period** specifies how long generated backup copies are retained before they become eligible for deletion, directly supporting recovery requirements and storage-capacity management.

These settings work together as the core elements of a policy: backup behavior, execution schedule, execution timing, and backup lifecycle. Huawei Cloud documentation describes backup policies as configurable schedules with retention settings, and Huawei storage documentation similarly identifies scheduling and retention as policy-controlled protection attributes. See Huawei Cloud CBR documentation and Huawei Cloud EVS backup policy documentation.

QUESTION NO: 23

When a user creates a VM on the FusionSphere OpenStack web client, if the disk space is insufficient, which of the following items should be expanded?

A. Ceilometer-data

B. Image

C. Image-cache

D. Swift

ANSWER: B

Explanation:

Image should be expanded. In FusionSphere OpenStack, the Image storage area supports the image service repository used during virtual-machine provisioning. When a user creates a VM through the web client, the selected base image must be accessed and prepared so that the compute service can create the VM's system disk from that image. If this image-related storage area has insufficient capacity, image handling and VM creation can fail because the required image data cannot be stored or processed successfully.

Expanding **Image** provides additional capacity for the image repository and allows the platform to continue supporting image-based VM deployment. This is particularly important when administrators upload larger operating-system images, retain multiple image versions, or create VMs from images that require significant temporary or prepared image space during provisioning. OpenStack's Image service (Glance) is responsible for discovering, registering, and retrieving VM images, while the Compute service uses an image as the source for an instance's root disk. Huawei's FusionSphere implementation follows this OpenStack image-driven provisioning model. See the OpenStack Image service documentation and the OpenStack instance-launch documentation for the relevant image and instance-creation workflow.

QUESTION NO: 24

Which of the following statements about DR is incorrect?

A. DR cannot retrieve lost data.

B. DR cannot recover damaged data.

C. DR cannot restore services.

D. DR ensures business continuity.

ANSWER: C

Explanation:

"DR cannot restore services." is the incorrect statement. Disaster recovery (DR) is specifically intended to restore the availability of applications, systems, and supporting infrastructure after a disaster or major service interruption. A DR solution maintains or rebuilds the recovery environment and uses replicated system data to bring protected workloads back online at a recovery site, enabling the organization to resume critical business operations within its recovery objectives. The exact

recovery process can include failover to a standby environment, restoration of virtual machines and application dependencies, network redirection, and later failback to the production environment. Thus, restoring services is a core outcome of DR planning and implementation, rather than something DR is unable to do. Huawei Cloud's Cloud Disaster Recovery Service documentation describes DR as a service for protecting cloud servers and supporting recovery after faults or disasters, including failover and failback capabilities. Huawei's broader business-continuity guidance also treats disaster recovery as a key capability for maintaining or resuming essential services following disruptive events. See Huawei Cloud Disaster Recovery Service documentation and Huawei Cloud Cloud Disaster Recovery Service overview.

QUESTION NO: 25

The FusionCompute distributed virtual switch (DVS) is similar to a two-layer physical switch. Which of the following is used for DVS to connect VMs and physical networks? (Multiple Choice)

A. Port group

B. VLAN pool and subnet

C. Uplink group

D. Downlink group

ANSWER: A C

Explanation:

In FusionCompute, a distributed virtual switch provides the virtual Layer 2 forwarding plane that joins VM virtual NICs to the external physical network through host physical NICs. **Port group** is the VM-facing logical networking construct on a DVS. A VM virtual NIC is connected to a port group, which applies consistent network attributes and policies, such as VLAN-related configuration, to the virtual ports used by attached VMs. This lets multiple VMs access the same logical network while being managed centrally through the distributed switch.

Uplink group is the physical-network-facing construct. It associates the DVS with the hosts' physical NIC resources and defines the uplink path over which traffic leaves or enters the virtual environment. Traffic from a VM therefore enters through its port group, is switched by the DVS, and reaches the physical network through the uplink group. Together, these two constructs provide the required VM-side access and physical-side connectivity, mirroring the access-port and uplink roles commonly found in a Layer 2 physical switching design. Huawei's FusionCompute product and support resources describe the platform's virtualized network architecture and configuration model: Huawei FusionCompute and FusionCompute Technical Support.

QUESTION NO: 26

Which of the following actions should be performed before implementing a major change in a cloud data center? (Multiple Choice)

A. Perform an impact and risk assessment.

B. Obtain required change approval.

C. Prepare and verify a rollback plan.

D. Notify stakeholders only after the change fails.

ANSWER: A B C

Explanation:

A major change can affect service availability, security, performance, and recoverability. Therefore, the change must be assessed for technical and business risk, approved through the established change-management process, and supported by a tested rollback plan. These controls ensure that the organization understands the possible impact, has authorized the implementation, and can restore service if the intended result is not achieved.

Notification only after an implementation fails is not an acceptable change-management practice. Relevant stakeholders should be informed before a change is implemented, especially where a service window or possible customer impact exists. This approach aligns with IT service-management practices for controlled change enablement. See AXELOS ITIL resources.

QUESTION NO: 27

Which of the following statements about RPO is incorrect?

A. RPO is short for recovery point objective.

B. RPO is measured based on data.

C. RPO is the time taken to recover data.

D. In theory, the shorter the RPO is, the better.

ANSWER: C

Explanation:

“RPO is the time taken to recover data.” is the incorrect statement. Recovery point objective (RPO) defines the maximum acceptable amount of data loss measured in time: specifically, it is the point in time to which data must be restored after an interruption. For example, an RPO of one hour means that, in a worst-case event, the organization accepts losing up to one hour of changes made since the last recoverable copy or replication point. RPO therefore drives backup frequency, snapshot scheduling, and replication design. A shorter RPO generally requires more frequent protection or near-continuous replication, reducing the possible data-loss window. The duration required to resume services or restore data is instead described by the recovery time objective (RTO). RTO measures the maximum acceptable downtime and guides the selection of recovery procedures and infrastructure capabilities. Huawei disaster-recovery planning similarly distinguishes data-loss objectives from service-recovery-time objectives. See Huawei’s disaster recovery overview and the NIST contingency-planning guidance for the definitions and use of recovery objectives.

QUESTION NO: 28

In the multi-tenant, resource sharing scenario, which of the following issues causes customer information and processes to be attacked by another customer?

A. Isolation failure

B. Information disclosure

C. Data loss

D. Relationship invalidity

ANSWER: A

Explanation:

Isolation failure is correct because a multi-tenant cloud environment relies on strict logical and, where applicable, physical separation between each tenant’s workloads, data, identities, networks, and management operations. If that separation fails, one customer may gain an unintended path to another customer’s virtual machines, containers, storage, network traffic, or processes. This is the core cross-tenant security risk described in the question: shared infrastructure is acceptable only when the platform’s isolation controls reliably prevent one tenant from affecting or accessing another.

Effective tenant isolation is enforced through mechanisms such as hypervisor and container isolation, virtual network segmentation, access-control policies, encryption, resource quotas, and secure management-plane authorization. Cloud operators must continuously validate those controls, promptly remediate platform vulnerabilities, and monitor for anomalous cross-tenant access attempts. Huawei Cloud describes virtualization and tenant isolation as foundational security capabilities in its cloud security architecture; the Cloud Security White Paper provides relevant details at Huawei Cloud Security White

Paper. Further general guidance on secure multi-tenancy and isolation controls is available from CISA's Cloud Security Technical Reference Architecture.