DUMPSARENA

Dell ECS Deploy 2023 Exam

EMC D-ECS-DY-23

Version Demo

Total Demo Questions: 8

Total Premium Questions: 86

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, ECS Concepts	29
Topic 2, ECS Installation	26
Topic 3, ECS Administration	24
Topic 4, ECS Troubleshooting	7
Total	86

QUESTION NO: 1

Which command is used to update the hare switch firmware?

A. updateArista

B. flashboot

C. firmwaredownload

D. update firmware

ANSWER: A

Explanation:

`updateArista` is the ECS service command used to update firmware on the Arista switches that provide the Hare network. In an ECS deployment, Hare is the internal, high-speed node-to-node network, and its switching infrastructure is managed through ECS-specific operational procedures rather than through a generic switch CLI command. The command coordinates the firmware-update workflow expected for the supported Arista hardware, helping ensure that the appropriate firmware image and deployment process are used consistently across the relevant switches. Firmware maintenance should be performed only as part of the documented ECS upgrade or service procedure, with the required prerequisites, maintenance planning, and post-update verification completed. This avoids introducing an unsupported switch state or disrupting internal ECS connectivity. Dell publishes ECS product documentation and support materials, including installation, maintenance, and upgrade guidance, through the Dell ECS documentation portal. Additional current product advisories, software, and support resources are available from the Dell ECS support page.

QUESTION NO: 2

An ECS storage administrator creates a bucket to be used by the NFS clients. The development team tries to write some objects using the S3 browser on the bucket. They call the storage administrator to inform them that they are experiencing write operation errors.

What is a possible reason for the errors?

A. Bucket is Read only

B. Secret key used is not correct

C. Object user needs write credentials

D. Base URL was not configured

ANSWER: A

Explanation:

Bucket is Read only is a possible reason for the failed S3 write operations. ECS applies access controls to buckets and objects, including permissions that govern whether clients can create, replace, or delete object data. If the bucket has been configured as read-only, an S3 client such as an S3 browser can authenticate and may be able to list or retrieve permitted content, but its PUT or other modifying requests are rejected because the bucket does not allow write activity. The error is therefore consistent with an access-policy condition on the target bucket rather than with the fact that NFS clients use it. Before enabling application writes, the administrator should verify the bucket's configured access mode and ensure that its intended data-access design permits write requests through S3. This is especially important for buckets used in a multiprotocol environment, where the required level of access must be deliberately assigned for each supported client workflow. Dell ECS documentation describes ECS as an object-storage platform supporting S3-compatible access and file protocols, with access administration forming part of the data-access configuration. See the Dell ECS product page and the Dell ECS documentation portal.

QUESTION NO: 3

Compliance monitoring is enabled on an ECS system. An operations team receives a BAD event even though no confirmed retention-policy violation has been reported.

Which compliance states can cause the BAD event?

A. READY and DEGRADED

B. SUSPECT and UNKNOWN

C. BAD and NOT READY

D. DEGRADED and UNTRACKED

ANSWER: B

Explanation:

A BAD event is generated when the monitored compliance state is SUSPECT or UNKNOWN. In either state, ECS cannot affirm a known healthy compliance condition: SUSPECT requires investigation, while UNKNOWN means the required state cannot be determined from the available information.

A BAD event is therefore not limited to a confirmed violation. The event is intended to make uncertain or potentially noncompliant conditions visible for investigation.

QUESTION NO: 4

How do ECS object users authenticate to the ECS OpenStack Swift Service when running the Version 1 protocol?

A. .PEA file

B. Login and password

C. Scoped token

D. Simple token

ANSWER: B

Explanation:

Login and password is correct because Swift Version 1 authentication is a credential-based flow. A client submits the object user's login name and password (commonly represented in Swift requests by the authentication user and key headers) to the authentication endpoint. After the credentials are successfully validated, the service returns an authentication token and the storage URL that the client uses for subsequent object operations. In ECS, an object user is created under an ECS namespace and has credentials that applications use to access ECS-supported object protocols, including Swift. The token is therefore the result of successful Version 1 credential authentication and is presented on later requests; it is not the initial authentication mechanism. This follows the standard Swift authentication sequence, in which the client obtains a token by supplying user credentials and then uses that token to access the account's storage endpoint. See the OpenStack Swift authentication API documentation and Dell's ECS documentation and support resources for ECS object-access documentation.

QUESTION NO: 5

A company is integrating Dell EMC Isilon CloudPools with the ECS and would like to enable CloudPools for disaster recovery failover. Which deep copy option should be used?

A. Allow

B. Force

C. Stub sync

D. Deny

ANSWER: B

Explanation:

Force is the appropriate deep-copy setting when CloudPools is being used to support disaster-recovery failover. This setting ensures that CloudPools performs a deep copy of the file's cloud-resident content rather than relying only on a stub representation or leaving the operation dependent on the file's existing local state. As a result, the required data is fully materialized for the copy operation, providing a usable and complete data set at the recovery location.

For a failover design, the recovery environment must be able to access the actual file contents after the protected data has been transferred. Forcing deep copy provides that assurance for data managed through the CloudPools tiering relationship with ECS. This behavior is particularly important for files that have already been archived, where the local namespace entry may exist but the primary file data resides in the object-store tier. Dell's CloudPools documentation describes deep-copy behavior and the management of archived file data, while the ECS documentation provides the object-storage platform context for the integration.

Reference documentation: Dell PowerScale OneFS documentation and Dell ECS documentation.

QUESTION NO: 6 - (SIMULATION)

What is the correct order of steps used to create a VDC federation?

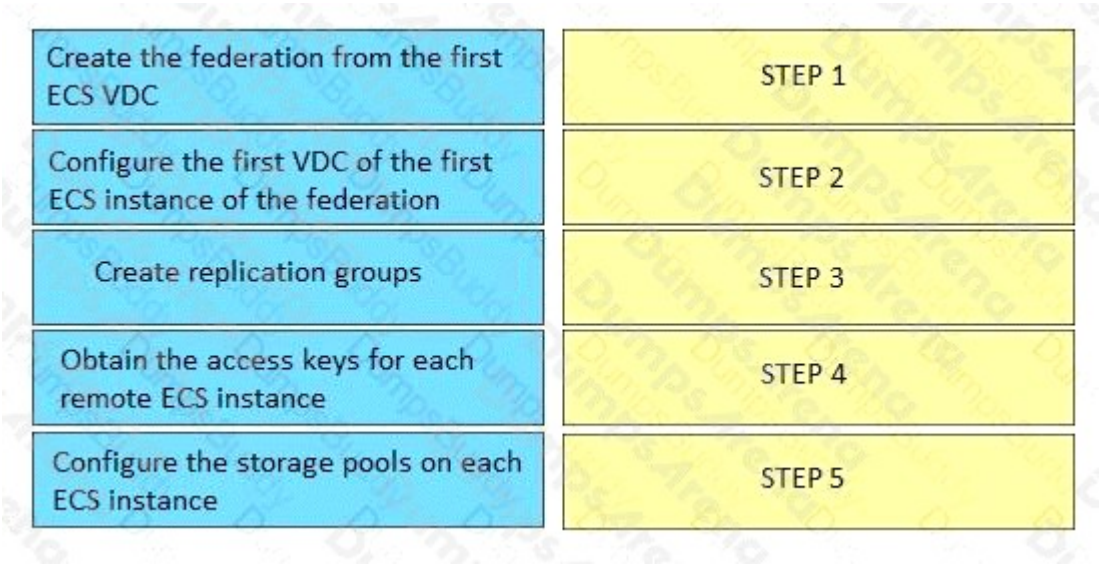

ANSWER: See the explanation for the answer

Explanation:

The correct sequence for creating a VDC federation is:

1. **Step 1:** Configure the storage pools on each ECS instance.
2. **Step 2:** Create replication groups.
3. **Step 3:** Obtain the access keys for each remote ECS instance.
4. **Step 4:** Create the federation from the first ECS VDC.
5. **Step 5:** Configure the first VDC of the first ECS instance of the federation.

Storage pools and replication groups must exist before federation configuration. The remote ECS access keys are then required to establish the federation and configure its initial VDC.

QUESTION NO: 7

During pre-ECS software installation, which system check should be manually performed?

A. Platform data

B. Rack networking

C. Network separation

D. Disk configuration

ANSWER: B

Explanation:

Rack networking should be manually checked before ECS software installation because ECS depends on a correctly cabled and configured network foundation from the first stage of deployment. The installer uses the supplied network information to configure node interfaces and establish the communication required among ECS nodes, the management network, and the customer-accessible data network. Before beginning installation, personnel should physically verify the rack's network cabling, switch connections, link status, and the availability of the required network paths and addressing. This validation is especially important because a cabling or switch-port issue can prevent node discovery, initial cluster communication, or access to ECS management services after deployment. A manual rack-networking check confirms that the physical environment matches the deployment plan before software configuration starts, reducing the likelihood of an installation interruption caused by infrastructure conditions that software cannot reliably validate in advance. Dell's ECS documentation and product-support resources provide the installation prerequisites and network-planning guidance used for this predeployment verification: Dell ECS documentation and Dell ECS product documentation.

QUESTION NO: 8

What services are reachable from all nodes and must be included in every ECS data center?

A. NTP and SNMP

B. DHCP and PXE

C. DNS and NTP

D. SMTP and FTP

ANSWER: C

Explanation:

DNS and NTP are required shared infrastructure services for an ECS data center because every ECS node needs reliable name resolution and synchronized time. DNS enables nodes and ECS services to resolve the fully qualified domain names used for internal communication, management access, and configured external endpoints. This dependency must be reachable consistently from all nodes so that a node does not rely on locally maintained or inconsistent hostname mappings.

NTP provides a common, accurate clock across the ECS deployment. Consistent time is essential to distributed-system operation: it allows event records, audit information, logs, certificates, authentication flows, and operational troubleshooting data to be interpreted in the correct sequence across all nodes. ECS deployment planning therefore includes identifying DNS and NTP servers that are accessible on the customer network from each node in the data center. These services are external network prerequisites rather than services that are expected to be supplied by individual ECS nodes. Dell's ECS product documentation is available through the Dell ECS documentation portal, and Dell's ECS support resources are available at Dell ECS Support Matrix.