COH-150 MCQs
COH-150 Exam Questions
COH-150 Practice Test
COH-150 TestPrep
COH-150 Study Guide

PASS

killexams.com

COHESITY

COH-150

Cohesity Certified Protection Associate - MultiCloud (COH150)

ORDER FULL VERSION

Question: 1118

Which command would you use to check the status of a CloudArchive job in Cohesity?

A. list-cloudarchive-jobs
B. get-cloudarchive-status
C. show-cloudarchive-job-status
D. check-cloudarchive-job

Answer: A

Explanation: The command `list-cloudarchive-jobs` allows users to view the status of CloudArchive jobs and their details in Cohesity.

Question: 1119

You are reviewing a CloudSpin operation that failed due to insufficient permissions. What action should be taken to resolve this issue?

A. Increase the storage capacity in the cloud
B. Restart the CloudSpin service
C. Review and modify user access permissions
D. Change the destination cloud provider

Answer: C

Explanation: To resolve permission issues that caused the CloudSpin operation to fail, it is important to review and modify user access permissions to ensure that the necessary rights are granted for the operation.

Question: 1120

An e-commerce company operating a hybrid multicloud setup with Cohesity clusters on-premises and in Azure faces a regional outage in their primary data center, requiring invocation of DRaaS to spin up EC2 instances from replicated snapshots while adhering to PCI DSS requirements for immutable backups. Which SiteContinuity parameters and settings must be applied in the DR plan to ensure pay-as-you-go AWS infrastructure activation and rapid recovery of 200 Kubernetes pods?

A. Set up DRaaS failover orchestration with AWS Auto Scaling Groups parameterized for pod deployment, enabling immutable WORM snapshots with 90-day retention via FortKnox integration

B. Configure near-sync replication policies in SiteContinuity with RTO sliders set to 10 minutes, using CloudReplicate to propagate changes to Azure as a tertiary site for PCI compliance auditing

C. Activate SaaS-based automated failback with ML-driven recommendations for clean recovery points, specifying EC2 instance types (e.g., m5.4xlarge) and integrating with Azure Site Recovery for hybrid validation

D. Enable CDP with log consolidation every 30 seconds and post-failover Bash scripts to orchestrate Kubernetes pod restarts using kubectl apply commands on EKS clusters

Answer: A, D

Explanation: DRaaS in SiteContinuity leverages pay-as-you-go AWS EC2 via Auto Scaling Groups for dynamic pod scaling, with FortKnox ensuring PCI DSS immutable WORM snapshots for 90-day retention during failover. CDP with 30-second log consolidation provides near-zero RPO for Kubernetes data, while Bash scripts automate post-failover kubectl commands to restart 200 pods efficiently, minimizing downtime in hybrid environments.

Question: 1121

In a scenario where data is tiered to an external cloud storage, which of the following statements is true regarding data retrieval?

A. Retrieving data from the cloud incurs additional costs and may take time.
B. Data retrieval from the cloud is instant and does not incur any delays.
C. Data cannot be retrieved once it is tiered to the cloud.
D. All data in the cloud is automatically replicated back to the on-premises system.

Answer: A

Explanation: When data is tiered to an external cloud storage, retrieving that data can incur additional costs and may take time, as it involves moving data back from the cloud to the on-premises environment.

Question: 1122

Which feature of CloudSpin allows users to create multiple copies of a VM for testing and development purposes?

A. Snapshot management
B. Backup scheduling
C. Dev/test clone
D. Policy enforcement

Answer: C

Explanation: The dev/test clone feature of CloudSpin allows users to create multiple copies of a VM for testing and development purposes, enabling safe experimentation without affecting the original VM.

Question: 1123

A Cohesity setup for climate modeling tiers sim outputs to Azure Geo Zone Redundant Cool at 79% utilization with 110-day coldness, but model iterations fail from 7-hour rehydrations. Which zonal configs would minimize, supporting climate-scale sims and GZRS resilience?

A. Enable zone-aware prefetch with --zone-prefetch enabled --zones 1,2,3
B. Drop coldness to 95 days and use LRS for intra-zone speed
C. Configure GZRS with standard rehydration only
D. Pre-tier models to Hot during iteration phases

Answer: A, B

Explanation: Zone-aware --zone-prefetch enabled --zones 1,2,3 anticipates rehydrations for faster iterations, scaling with GZRS. 95-day coldness with LRS boosts speed for modeling. GZRS standard is default, and pre-tiering to Hot increases costs unnecessarily.

Question: 1124

DR process for CloudRetrieve from Azure Archive to AWS cluster, 320TB load. Steps for zero-downtime failover simulation?

A. Simulate: --dry-run-retrieve=true with load balancer attach
B. Orchestrate: Runbook --failover-zero-dt=true --test-mode
C. Sync: --delta-sync=true pre-full for ongoing changes

D. Validate: --ha-check=true post-sim for cluster health

Answer: B, C, D

Explanation: Runbook --failover-zero-dt=true simulates in test-mode. --delta-sync=true minimizes data gap. --ha-check validates health post. Dry-run tests retrieve but not full failover.

Question: 1125

What is a critical aspect to consider when using CloudSpin for data recovery in a multicloud environment?

A. Data must be restored to the original location
B. It only supports AWS cloud services
C. Data encryption is not supported
D. Recovery time objectives (RTO) can be significantly reduced

Answer: D

Explanation: CloudSpin can significantly reduce recovery time objectives (RTO) by enabling quick restoration of virtual machines in a multicloud environment.

Question: 1126

In a scenario where a VM is cloned for development purposes using CloudSpin, which of the following should be considered regarding the original VM's data?

A. The cloned VM will have access to the original VM's data.
B. The cloned VM will not affect the original VM's performance.
C. The data in the cloned VM will be automatically encrypted.
D. The original VM's data will be deleted after cloning.

Answer: B

Explanation: The cloned VM will not affect the original VM's performance, allowing development work to proceed without impacting production workloads.

Question: 1127

In a scenario where a financial services firm experiences a ransomware attack on their primary VMware vSphere cluster, they initiate a failover to AWS using Cohesity SiteContinuity for DR. The cluster has replicated snapshots with near-sync replication enabled, and the DR plan includes automated orchestration for 50 critical VMs. Which of the following steps must be executed in sequence during the failover process to ensure minimal data loss and seamless transition to Amazon EC2 instances?

A. Activate the DR plan in the SiteContinuity GUI, triggering format conversion from VMDK to AMI
B. Manually provision EC2 instances in the AWS account prior to failover initiation
C. Validate the recoverability index using ML-based recommendations on the replicated snapshots
D. Enable continuous data protection (CDP) post-failover to capture any delta changes during the outage

Answer: A,C,D

Explanation: Activating the DR plan in the SiteContinuity GUI automatically triggers the failover orchestration, including seamless format conversion from VMDK to AMI for compatibility with Amazon EC2, ensuring the VMs spin up ready-to-run in the cloud. Validating the recoverability index with machine learning-based recommendations identifies the cleanest point-in-time snapshot for failover, minimizing data loss to near-zero RPO. Enabling continuous data protection (CDP) post-failover ensures ongoing capture of any incremental changes during the outage, facilitating efficient failback without data gaps. Manual provisioning of EC2 instances is unnecessary as SiteContinuity automates the spin-up using pay-as-you-go AWS infrastructure in the customer's account.

Question: 1128

In R&D, a pharma company sets up dev/test clones of 110 lab VMs in Azure from GCP for drug simulation. Which steps integrate with Azure Batch for HPC jobs?

A. Clone with --batchPoolId sim-pool --vmSize Standard_D4s_v3 --nodeCount 20 --taskJson /path/to/sim-task.json
B. Register Azure Batch with --accountName batch-pharma --key primary --endpoint https://batch-pharma.eastus.batch.azure.com
C. Configure policy with --hpcIntegration true --jobQueue high_priority --dataTransfer azcopy --masking proprietary_compounds
D. Runbook with --submitBatch az batch job create --poolId sim-pool

Answer: A, B

Explanation: Cloning to --batchPoolId with --vmSize and --nodeCount schedules HPC simulations. Registering Batch account with --accountName and --endpoint enables API-driven job submission.

Question: 1129

During an unplanned outage in a retail chain's data center, the Cohesity team executes a CloudSpin operation to recover 40 POS VMs to Azure Stack HCI. The VMs have custom drivers and require guest customization. What sequential steps ensure a successful operation with RTO under 20 minutes?

A. Trigger failover from Cohesity UI selecting Azure Stack target, auto-converting VMDK to VHDX with --customize-guest=true for driver injection
B. Pre-run 'cohesity-cli agent install --target=azurestack --physical' on recovery VMs, then monitor via dashboard with threshold alerts for 95% completion
C. Validate network mappings in the wizard to bridge POS VLANs, executing post-spin PowerShell scripts for license reactivation via Azure Automation
D. Initiate manual replication sync with 'cohesity-cli replicate force --job=', importing VHDX directly to HCI storage pools before launch

Answer: A, C

Explanation: UI-triggered failover with guest customization handles driver issues for quick VMDK-to-VHDX conversion. Network mapping and post-spin PowerShell ensure connectivity and licensing, achieving sub-20-minute RTO for POS recovery.

Question: 1130

A user is tasked with ensuring that all data sent to the cloud is both compressed and encrypted. What should they do first?

A. Enable data encryption
B. Enable data compression
C. Set a data retention policy
D. Configure a protection policy

Answer: D

Explanation: Configuring a protection policy is the first step, which allows the user to enable both data compression and encryption settings for the data being sent to the cloud.

Question: 1131

What is a key consideration when choosing between CloudRetrieve and traditional recovery methods?

A. Cost of storage
B. Data type compatibility
C. Speed of recovery
D. User training requirements

Answer: C

Explanation: The speed of recovery is a key consideration when choosing between CloudRetrieve and traditional recovery methods, as CloudRetrieve typically offers faster access to archived data.

Question: 1132

An organization is planning to use Cohesity to manage their backup and archive strategies in a multi-cloud environment. What is a significant advantage of using Cloud Edition in this scenario?

A. Increased hardware maintenance
B. Seamless integration with public clouds
C. Limited scalability
D. Higher operational costs

Answer: B

Explanation: The Cloud Edition offers seamless integration with public cloud services, making it advantageous for organizations managing backup and archive strategies in a multi-cloud environment.

Question: 1133

Which of the following best describes the role of age-based policies in managing data tiers?

A. They restrict access to data
B. They dictate when data should be moved to lower-cost storage
C. They determine how often data is backed up
D. They enhance data encryption

Answer: B

Explanation: Age-based policies dictate when data should be moved to lower-cost storage, allowing organizations to manage their storage resources effectively based on data age.

Question: 1134

When configuring an external target for tiered data, which of the following cloud storage types is NOT typically supported by Cohesity?

A. Amazon S3
B. Microsoft Azure Blob Storage
C. Google Drive
D. IBM Cloud Object Storage

Answer: C

Explanation: Cohesity typically supports cloud storage types such as Amazon S3, Microsoft Azure Blob Storage, and IBM Cloud Object Storage, but Google Drive is not typically supported for tiered data.

Question: 1135

In a DevOps pipeline, CloudTier offloads container images to MinIO S3 after builds, but CI/CD pulls fail due to recall queues. Which options resolve?

A. Prioritize recall queues with "priority: HIGH" in policy for build views.
B. Use multi-threaded recall with "threads: 16" for image layers.

C. Disable encryption for internal MinIO to speed transit.
D. Enable prefetch caching for frequent image tags.

Answer: A, B

Explanation: High-priority queues in policies fast-track CI/CD recalls. 16-thread multi-recall handles layered images concurrently, reducing pipeline delays.

Question: 1136

For a retail chain's CloudArchive Direct setup archiving POS data directly to GCP Coldline, the config includes no local backups and metadata indexing. What External Target options support hybrid full/granular restores to either AWS or on-premises clusters?

A. Target Protocol: S3-Compatible with multi-region replication enabled
B. Metadata Retention: Infinite with index format JSON for cross-cloud search
C. Recovery Endpoint: Configure IAM for AWS STS assume-role in retrieve policy
D. Tiering Policy: Auto-transition to Archive after 30 days with recall fee optimization

Answer: A, B, C

Explanation: S3-Compatible protocol for GCP Coldline allows multi-region replication, supporting restores to AWS or on-premises. Infinite metadata retention with JSON indexing enables consistent search across environments. IAM configuration for AWS STS assume-role in the retrieve policy secures cross-cloud access for full or granular operations. Tiering optimizes costs but does not impact restore endpoints.

Question: 1137

What is the primary benefit of using External Targets for archiving in Cohesity?

A. Enhanced security features
B. Cost-effective storage
C. Improved data retrieval speed
D. Automatic backup scheduling

Answer: B

Explanation: External Targets allow organizations to utilize cost-effective storage solutions, such as public clouds, for archiving data, which can significantly reduce

overall storage costs.

Question: 1138

In a telecom provider's setup, Cohesity archives call detail records (CDRs) from Cassandra clusters to external Wasabi targets with v2 incremental forever, incorporating AI-driven anomaly detection for fraud. The workflow runs bi-hourly with 99.999% durability SLA. Which features optimize for high-volume CDRs and detection?

A. Target config: --wasabi true --v2 forever --anomaly-ai enabled, durability 99.999
B. Policy: bihourly-run, retain 2y, ai-scan post-archive with fraud-pattern thresholds at 0.1%
C. CLI: `policy_ai --source cassandra --target wasabi://cdr-bucket --incremental v2 --ai-fraud true --threshold 0.1 --durability 5nines`
D. Set metadata enrichment for CDR timestamps, integrating with external SIEM via webhook alerts on anomalies

Answer: A, C

Explanation: The target config --wasabi true --v2 forever --anomaly-ai enabled and 99.999 durability supports high-volume CDR archival with built-in AI for fraud patterns in incremental workflows. The CLI `policy_ai --source cassandra --target wasabi://cdr-bucket --incremental v2 --ai-fraud true --threshold 0.1 --durability 5nines` tunes bi-hourly runs and 0.1% thresholds, ensuring SLA compliance and proactive detection.

KILLEXAMS.COM

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.

Find Exam
Search your required exam

Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.