DUMPSBOSS.

Hitachi Ops Center Protection

Hitachi HQT-6711

Version Demo

Total Demo Questions: 19

Total Premium Questions: 196

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

How can Hitachi Ops Center Protector's backup functionality be enhanced?(ChooseTwo)

A. Integrating with cloud storage solutions

B. Enabling real-time synchronization

C. Regularly testing backup integrity

D. Increasing the compression of backup files

ANSWER: A C

Explanation:

Integrating with cloud storage solutions enhances Hitachi Ops Center Protector by extending backup-copy retention beyond the local environment. Protector supports a range of copy targets and policies, enabling organizations to use cloud-based object storage as part of a resilient, off-site data-protection strategy. This improves recovery preparedness by maintaining an additional copy that is geographically and operationally separated from the protected primary systems. Hitachi describes Ops Center Protector as providing policy-driven data protection, copy management, and recovery capabilities across physical, virtual, and cloud environments.

Regularly testing backup integrity is equally important because a backup has operational value only when it can be reliably read and restored. Scheduled validation and recovery testing help identify corruption, incomplete protection policies, credential or connectivity changes, and recovery-process issues before an actual outage occurs. Testing also confirms that recovery-point and recovery-time objectives can be met in practice. Used together, cloud storage integration provides stronger copy resilience and retention flexibility, while recurring integrity tests provide confidence that those protected copies are usable for recovery. See the Hitachi Ops Center Protector documentation and the Hitachi Ops Center Protector product overview.

QUESTION NO: 2

In the context of Hitachi Ops Center Protector, what characteristic is true about a node? (Choose Two)

A. Nodes can operate independently of each other

B. Nodes can be dynamically added or removed without significant downtime

C. Each node requires manual configuration and setup

D. Nodes are interdependent and require manual synchronization

ANSWER: A B

Explanation:

“Nodes can operate independently of each other” is correct because Ops Center Protector uses a distributed architecture in which registered nodes perform their assigned data-protection roles, such as acting as application hosts, media servers, or repository-related components. The central management services coordinate protection policies, jobs, inventory, and monitoring, but a node remains a separately installed and managed system that performs its own local protection processing. This design permits protection infrastructure to be distributed across hosts and locations according to workload and storage requirements.

“Nodes can be dynamically added or removed without significant downtime” is also correct. Protector is designed to scale by registering additional nodes when protected workloads, media-server capacity, or repository resources increase. Nodes can subsequently be decommissioned when they are no longer required, provided that associated policies, data movements, and protection dependencies are handled appropriately. These lifecycle operations are administrative changes rather than a requirement to stop the overall Protector environment, allowing the deployment to evolve while other appropriately configured protection activities continue. Hitachi’s product documentation describes the Protector architecture and its node-management workflow in the Ops Center Protector overview and the Managing nodes documentation.

QUESTION NO: 3

After a change to network connectivity between Protector nodes and a backup destination, several scheduled jobs fail. Which two sources of information should the administrator review first to isolate the failure? (Choose two)

A. System logs and error messages

B. Backup-operation status and performance metrics

C. User-interface preference settings

D. The source LUN identifier by itself

ANSWER: A B

Explanation:

System logs and error messages provide time-stamped details about authentication, connectivity, configuration, and job-processing failures. **Backup-operation status and performance metrics** show which jobs failed, when the behavior changed, and whether throughput or duration degraded before failure. A report of user-interface preferences does not diagnose a data path, and the source LUN identifier alone does not establish why connectivity to a backup destination failed.

QUESTION NO: 4

What is the advantage of using policy-driven replication in file-based data protection?

A. It reduces the complexity of managing individual file backups

B. It increases the speed of data transfers

C. It eliminates the need for encryption

D. It allows for direct manipulation of file metadata

ANSWER: A

Explanation:

It reduces the complexity of managing individual file backups. In policy-driven file protection, an administrator defines the required protection behavior once—for example, the source data set, target repository, schedule, retention requirements, and replication settings—and then applies that policy consistently to the appropriate protected objects. The protection platform uses the policy to orchestrate and run the data flow, rather than requiring the administrator to configure, monitor, and maintain separate backup or replication settings for every individual file. This model improves operational consistency as the environment grows, makes recurring protection tasks easier to administer, and helps ensure that files assigned to the same protection policy receive the same treatment. Changes can also be made centrally by updating the policy, which is much more manageable than revising many individual backup definitions. Hitachi Ops Center Protector is designed around policy-based data-protection operations and automated data flows; its documentation describes policies as the mechanism used to define and control protection activities. See the Hitachi Ops Center Protector documentation and the Ops Center Protector product overview.

QUESTION NO: 5

What role does the control layer play in the Hitachi Ops Center Protector architecture?

A. Manages data storage and retrieval

B. Coordinates and enforces protection policies

C. Encrypts data before transmission

D. Monitors network traffic

ANSWER: B

Explanation:

The control layer **coordinates and enforces protection policies**. In Hitachi Ops Center Protector, this layer provides the centralized management functions that define how workloads and data are protected. Administrators use it to create and manage policies, associate policies with protected resources, establish schedules and retention behavior, and coordinate the protection operations performed by the underlying data movers and storage targets. This centralized policy-driven approach ensures that protection is applied consistently across the environment rather than requiring each backup or recovery activity to be configured independently. The control layer also maintains the operational context needed to initiate, track, and manage protection jobs, enabling governed backup, replication, and recovery workflows. In other words, it is the architectural layer responsible for orchestration and policy control, connecting administrative intent to the services that carry out data-protection operations. Hitachi describes Ops Center Protector as software for centralized data-protection management and orchestration; see the Hitachi Ops Center Protector product page and the Hitachi Vantara documentation portal for product documentation and architecture guidance.

QUESTION NO: 6

A remote replication relationship is healthy, but monitoring shows that the amount of data waiting to be transferred continues to increase during peak activity. The protected application has not changed its recovery-point objective. What should the administrator investigate first?

A. The available replication-path bandwidth compared with the workload change rate

B. The retention period for the remote copy

C. The role assigned to the backup administrator

D. The naming convention used for the source LUNs

ANSWER: A

Explanation:

The available replication-path bandwidth compared with the workload change rate is correct. A growing backlog indicates that changed blocks are being produced faster than the replication path can transfer them. If this persists, the remote copy can fall behind the required recovery-point objective. Retention affects how long copies are kept, not the rate of replication. RBAC and LUN naming do not address data-transfer throughput.

QUESTION NO: 7

What requirements are critical for block-based data protection?(Choosetwo)

A. Compatibility between storage systems

B. High network bandwidth

C. Manual configuration of storage arrays

D. Regular software updates

ANSWER: A B

Explanation:

Block-based data protection relies on replication technologies that operate on storage volumes and track changed blocks between a primary and secondary copy. Therefore, **Compatibility between storage systems** is essential: the participating systems, their replication features, and their supported software or microcode levels must meet the relevant interoperability requirements so that volume pairing, copy operations, and recovery workflows can be created and managed reliably. **High network bandwidth** is also critical because changed blocks must be transferred across the replication path. Bandwidth must

be sufficient for the protected workload's change rate and the required recovery-point objective; otherwise, replication can lag and the secondary copy may not meet its intended protection target. Hitachi Ops Center Protection orchestrates such storage-based data-protection operations, while the underlying storage replication configuration must satisfy the platform's supported topology and connectivity requirements. See the Hitachi Ops Center Protector documentation and the Hitachi Ops Center Protector product overview for product and replication context.

QUESTION NO: 8

What does integration of Hitachi Ops Center Protector with Amazon S3 enable for cloud backups? (Choose Two)

A. It allows for immediate restoration of operational systems

B. It supports the archiving of older backups to S3 Glacier

C. It facilitates the synchronization of data across multiple S3 regions

D. It reduces the necessity for on-premise backup infrastructures

ANSWER: B D

Explanation:

Integration with Amazon S3 provides Ops Center Protector with cloud object storage that can be used as a backup destination, extending protection beyond locally deployed storage. This reduces the necessity for on-premise backup infrastructures because organizations can retain backup copies in AWS rather than continually expanding local disk, tape, and associated facilities. The cloud target can support off-site retention strategies while Protector continues to manage backup policies and recovery operations.

It also supports the archiving of older backups to S3 Glacier. Amazon S3 Lifecycle rules can transition eligible objects from S3 storage classes to Amazon S3 Glacier storage classes, providing lower-cost long-term retention for backup data that is less likely to require rapid access. This approach lets organizations keep operational backups in an appropriate S3 tier while applying lifecycle-based archival to older backup objects. Hitachi identifies Ops Center Protector as supporting cloud-based data protection, and AWS documents the lifecycle transition capability used for archival. See Hitachi Ops Center Protector and Amazon S3 Lifecycle transition considerations.

QUESTION NO: 9

If a PUT request is made to the Hitachi Ops Center Protector REST API to update a policy and it returns a 204 status code, what does this imply?

A. The request was incorrect and nothing was updated.

B. The policy was successfully updated but there is no content in the response.

C. The policy update request created a new policy instead.

D. There was a server error while processing the request.

ANSWER: B

Explanation:

The policy was successfully updated but there is no content in the response. An HTTP `204 No Content` response is a successful status code: it confirms that the server received the PUT request, processed it successfully, and has no response body to return. For a REST API policy-update operation, this means the requested update completed without the API needing to return a representation of the updated policy. Clients should therefore treat the operation as successful based on the status code, rather than attempt to parse a response payload. If verification of the final policy configuration is required, the client can subsequently retrieve the policy through the appropriate GET endpoint. This behavior follows the standard HTTP definition of 204, which specifies that a request has succeeded and that the response contains no additional content.

See the HTTP Semantics specification for 204 No Content and Hitachi's Ops Center Protector documentation for product documentation and API resources.

QUESTION NO: 10

When creating reports in Hitachi Ops Center Protector, which of the following options is available?

A. Real-time reporting only

B. Scheduled reporting

C. Manual report generation

D. Event-triggered reporting

ANSWER: B

Explanation:

Scheduled reporting is available in Hitachi Ops Center Protector. Reporting is intended to provide administrators with recurring visibility into protection operations, such as backup activity, copy status, capacity-related information, and the overall state of protected resources. A report schedule lets an administrator define when a report is generated and distributed, supporting regular operational review without requiring an administrator to recreate the report each time it is needed. This is useful for daily checks, weekly operational summaries, and periodic audit or management reporting.

The scheduling capability is configured as part of the reporting workflow, where the report definition identifies the information to be collected and the schedule defines the recurring execution time. This enables consistent reporting over time and helps teams monitor data-protection service levels and trends. Hitachi describes Ops Center Protector as a centralized platform for managing and monitoring data-protection operations; its reporting functions support that operational-management role. See the Hitachi Ops Center Protector product page and the Hitachi Vantara documentation portal for product documentation and release-specific administration guidance.

QUESTION NO: 11

Backup jobs are completing successfully, but their duration has increased each week and may soon exceed the available protection window. Which two uses of Hitachi Ops Center Protector reporting are most relevant to the investigation? (Choose two)

A. Review job-duration and throughput trends

B. Review backup-operation status and history

C. Review user-interface display preferences

D. Review physical inventory labels on storage hardware

E. Review individual users' personal notification preferences

ANSWER: A B

Explanation:

Reviewing job-duration and throughput trends helps identify whether the protection window is being consumed by a gradual performance decline. **Reviewing backup-operation status and history** provides the job-level information needed to correlate the trend with affected policies, sources, or destinations.

User-interface preferences and physical inventory labels do not provide the operational evidence needed to analyze a growing backup duration.

QUESTION NO: 12

An organization is preparing to recover the Hitachi Ops Center Protector management environment after loss of its server. Which two actions are essential to make the recovery practical? (Choose two)

A. Exporting and retaining the Protector configuration

B. Storing the exported recovery material independently of the Protector server

C. Backing up every protected production volume to the Protector server

D. Recreating all policies manually after the outage

ANSWER: A B

Explanation:

Exporting and retaining the Protector configuration preserves the policies and operational settings needed to rebuild the management environment. **Storing the exported recovery material independently of the Protector server** prevents a failure of that server or its local storage from destroying both the product and its recovery information. Backing up every protected production volume is part of workload protection, but it does not replace recovery of the Protector configuration. Recreating policies manually only after an outage is slower and more error-prone than restoring the saved configuration.

QUESTION NO: 13

In the context of cloud backup strategies, what is a significant advantage of using Unified BackupInfrastructure (UBI) with Ops Center Protector?

A. It allows for the discontinuation of all local backups.

B. It simplifies backup architecture by centralizing control.

C. It reduces the reliance on encryption for data security.

D. It mandates the use of a single cloud provider.

ANSWER: B

Explanation:

It simplifies backup architecture by centralizing control. Unified Backup Infrastructure (UBI), used with Hitachi Ops Center Protector, is intended to provide a consistent, centrally managed framework for protecting data across the infrastructure. Rather than requiring separate backup-management approaches for each workload, storage platform, or target location, administrators can use a common control plane to define policies, schedule and monitor data-protection operations, manage copy targets, and review protection status. This operational consistency is especially valuable in cloud-oriented strategies, where an organization may combine on-premises systems with cloud or object-storage targets. Centralized control helps standardize protection processes, reduces administrative complexity, and improves visibility into whether required backup and recovery copies are being created successfully. It also supports a more scalable model: as protected environments and backup targets grow, teams can apply policy-driven management instead of building isolated backup processes. Hitachi describes Ops Center Protector as software for centrally managing data protection and copy-data operations across supported environments. See the Hitachi Ops Center Protector product page and the Ops Center Protector documentation.

QUESTION NO: 14

What are the key features of Role-Based Access Control (RBAC) in Hitachi Ops Center Protector? (Choose two)

A. Granular control over user permissions

B. Automatic encryption of all user data

C. Simplified user management across different roles

D. Real-time monitoring of user activities

ANSWER: A C

Explanation:

Role-Based Access Control in Hitachi Ops Center Protector provides **granular control over user permissions** by associating privileges with defined roles. Administrators can assign users only the authority needed for their operational responsibilities, such as viewing resources, managing protection policies, or performing recovery-related tasks. This supports the principle of least privilege and helps ensure that administrative actions are appropriately controlled within the data-protection environment.

RBAC also delivers **simplified user management across different roles**. Rather than configuring permissions separately for each individual account, an administrator defines or uses suitable roles and assigns users to them. When job responsibilities change, access can be updated efficiently through role assignments, creating a more consistent and scalable approach to access administration. This role-oriented model is especially useful when multiple teams require different levels of access to Protector functions and resources. Hitachi documentation describes the product's user and access-management capabilities in the Ops Center Protector documentation; product information is also available from Hitachi Vantara Ops Center Protector.

QUESTION NO: 15

What does integration of Hitachi Ops Center Protector with Amazon S3 enable for cloud backups?(ChooseTwo)

A. It allows for immediate restoration of operational systems

B. It supports the archiving of older backups to S3 Glacier

C. It facilitates the synchronization of data across multiple S3 regions

D. It reduces the necessity for on-premise backup infrastructures

ANSWER: B D

Explanation:

Integration with Amazon S3 provides Hitachi Ops Center Protector with cloud object storage that can serve as a scalable backup destination. This reduces the need to build and maintain equivalent on-premise backup-capacity infrastructure, while allowing backup data to be retained in AWS according to organizational data-protection and retention requirements. Hitachi describes Ops Center Protector as providing policy-based data protection with support for cloud targets, enabling organizations to use cloud storage as part of a broader backup and recovery strategy.

Amazon S3 also supports lifecycle-based archival of retained objects into lower-cost Amazon S3 Glacier storage classes. This makes it practical to keep older backup copies for long retention periods at a lower storage cost while preserving the ability to retrieve them when required. Lifecycle transitions are configured in AWS for the S3 bucket used as the backup target, aligning the storage class with the age and access requirements of the backup data. See the Hitachi Ops Center Protector product page and the Amazon S3 lifecycle transition documentation.

QUESTION NO: 16

What is the recommended practice for backing up Hitachi Ops Center Protector configuration settings and data?

A. Manual copy to external storage daily

B. Use of third-party backup solutions

C. Automated, scheduled backups within the software

D. On-demand backups before major changes

ANSWER: C

Explanation:

Automated, scheduled backups within the software is the recommended practice because Ops Center Protector's protection configuration is operationally important: it includes the definitions and relationships needed to manage data protection activities consistently. Using the product's built-in configuration-backup capability on a schedule makes the process repeatable, reduces the risk that an administrator forgets to create a backup, and provides recoverable configuration data after a server failure, rebuild, or other recovery event. The backup schedule should be aligned with the organization's change frequency and retention requirements, and the resulting backup data should be retained in a location that remains available if the Protector server itself is lost. Scheduled backups also establish a dependable recovery baseline, allowing teams to restore the Protector environment with its protection policies and associated settings intact. This approach is consistent with the product's administrative guidance for managing and protecting the Ops Center Protector environment and with Hitachi's positioning of Protector as centrally managed data-protection software. See the Hitachi Ops Center Protector documentation and the Ops Center Protector product page for product and administration information.

QUESTION NO: 17

Which functionality does Hitachi Ops Center Protector provide to enhance disaster recovery readiness?

A. It limits data replication to once every 24 hours.

B. It offers continuous data protection.

C. It only supports manual snapshot management.

D. It enables automated compliance checks.

ANSWER: B

Explanation:

It offers continuous data protection. Continuous data protection strengthens disaster recovery readiness by maintaining frequent, policy-driven recovery points rather than relying solely on an infrequent backup cycle. Hitachi Ops Center Protector centrally orchestrates data-protection operations across supported environments, including the creation and management of copies used for recovery. This approach helps organizations meet recovery-point objectives by making more current data available when a disruption, corruption event, or site-level incident requires recovery.

In practical disaster-recovery planning, recovery readiness depends on having protected copies that are created consistently, are visible to administrators, and can be used through repeatable recovery procedures. Ops Center Protector provides policy-based protection management and automation for such copy and recovery operations, reducing reliance on ad hoc manual activity. Its integration with Hitachi storage data-protection capabilities supports the use of replicated and point-in-time copies as part of a broader recovery strategy. Product information is available from Hitachi Vantara Ops Center Protector, and supporting product documentation is published through the Hitachi Vantara Documentation Portal.

QUESTION NO: 18

When creating a data flow for VMware protection, which of the following is the first step?

A. Defining the replication policy

B. Creating the necessary nodes

C. Configuring data encryption

D. Allocating storage space

ANSWER: B

Explanation:

Creating the necessary nodes is the first step because a VMware protection data flow must have its participating infrastructure defined before protection processing can be configured. In Ops Center Protection, nodes represent the systems and resources that take part in the workflow, such as the VMware environment and the storage or protection targets used by the data flow. The product uses these registered nodes to discover available resources, establish management connectivity, and present valid source and destination selections during later configuration steps. Once the required nodes are available, the administrator can create the VMware data flow and select its virtual-machine resources, protection destination, schedule, retention settings, and any replication behavior supported by the selected workflow. Establishing the node foundation first also ensures that Ops Center Protection can validate credentials and communication paths needed to execute jobs reliably. This sequencing follows the product's configuration model: register and manage the environment, then define the protection data flow that uses that environment. Refer to the Hitachi Ops Center Protection documentation and the Ops Center Protection product page for VMware-protection workflow and node-management information.

QUESTION NO: 19

Which feature of Hitachi Ops Center Protector is most beneficial for block storage data protection?

A. Policy-based automation

B. Real-time file sharing

C. Data compression

D. Network traffic monitoring

ANSWER: A

Explanation:

Policy-based automation is most beneficial because Hitachi Ops Center Protector uses policies to define, schedule, and consistently execute data-protection operations for block-storage resources. An administrator can specify protection requirements such as the source volumes, copy destination, frequency, retention, and recovery-point objectives, then allow Protector to orchestrate the appropriate copy-data workflow. This makes protection repeatable at scale and reduces the operational risk associated with manually configuring and running individual backup or replication jobs.

For block storage, consistent policy execution is especially valuable because applications often depend on predictable recovery points and coordinated protection of related volumes. Protector policies provide a centralized way to apply those requirements across supported Hitachi storage environments, monitor the resulting operations, and maintain protection according to the organization's service-level expectations. This approach supports both routine operational recovery and broader disaster-recovery planning while making it easier to standardize protection across many workloads. Hitachi describes Ops Center Protector as software for managing and automating data-protection operations; its documentation also identifies policies as the mechanism used to configure and run protection workflows. See the Hitachi Ops Center Protector product page and the Ops Center Protector documentation.