PassQueen
a.k.a dumpsqueen.com

Riverbed Certified Solutions Associate – Storage Delivery

Riverbed 501-01

Version Demo

Total Demo Questions: 9

Total Premium Questions: 90

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Storage Delivery Solutions	6
Topic 2, Whitewater	30
Topic 3, Granite	54
Total	90

QUESTION NO: 1

The delay between taking a snapshot at the edge and flushing the snapshot to the storage array is called what?

A. Snapshot Target

B. Snapshot Lag

C. Snapshot Window

D. Snapshot Integration

ANSWER: B

Explanation:

Snapshot Lag is the term for the elapsed delay between creation of a snapshot at the edge appliance and the point at which that snapshot is flushed, or committed, to the storage array. In a SteelFusion deployment, edge-side operations can continue locally while snapshot data and associated changes are transferred back to centralized storage. This creates a measurable interval during which the edge has taken the snapshot but the storage array has not yet received the completed flush. That interval is important operationally because it indicates how far the centrally stored recovery point may trail the snapshot state available at the edge. Administrators use snapshot-lag information when monitoring connectivity, available bandwidth, replication progress, and the expected recovery point for a protected workload. A growing lag can indicate that transfer capacity is insufficient for the rate of changed data, or that there is a connectivity or storage-side condition preventing timely completion. Riverbed's product documentation and support resources provide the authoritative terminology and operational guidance for SteelFusion snapshot and storage workflows: Riverbed Documentation Portal and Riverbed Support Portal.

QUESTION NO: 2

Granite Edge appliance supports which of the following protocols? (Select 2)

A. Fibre Channel

B. iSCSI

C. MPIO

D. NFS

ANSWER: A B

Explanation:

Granite Edge supports both Fibre Channel and iSCSI for presenting centralized Granite storage to applications at the branch. These are storage-area-network block protocols: Fibre Channel uses a dedicated FC fabric, while iSCSI transports SCSI block commands across an IP network. This allows a Granite Edge appliance to fit into sites using either an existing Fibre Channel SAN design or Ethernet-based iSCSI connectivity, while Granite Core centrally manages and protects the underlying data. The protocol choice is therefore driven by the host and switching infrastructure deployed at the branch, not by a different data-delivery model. Riverbed's Granite documentation describes Edge as the branch component that presents storage services locally while data is centralized through the Granite architecture. Refer to Riverbed's product documentation portal for Granite deployment and administration materials, and the Riverbed Support portal for product documentation and lifecycle resources. Fibre Channel and iSCSI are the two applicable storage protocols in this question.

QUESTION NO: 3

If a user is considering having backup data sent to Amazon Glacier, select which recommendations should be reviewed to ensure Whitewater appliance is appropriately designed for the environment? (Select 2)

A. Whitewater appliances should be sized larger to reduce the amount of potential access activity from Amazon Glacier.

B. Whitewater appliances should only deduplicate against local cache storage and not against local and Amazon Glacier storage.

C. Whitewater appliances should only accept archive data if using Amazon Glacier storage.

ANSWER: A B

Explanation:

“Whitewater appliances should be sized larger to reduce the amount of potential access activity from Amazon Glacier.” is appropriate because Amazon Glacier is designed for low-cost archival storage rather than frequent, immediate restores. Retrieval can require a restore request and a waiting period before archived objects are available. Sizing the Whitewater appliance with sufficient local cache capacity helps retain a larger working set of recently protected data locally, reducing the likelihood that backup software or restore operations must retrieve data from Glacier. This design consideration is particularly important where recovery-time objectives require prompt access to recent backup sets.

“Whitewater appliances should only deduplicate against local cache storage and not against local and Amazon Glacier storage.” is also appropriate for a Glacier-based archive design. Deduplication and backup processing should use the appliance’s active local cache and its associated metadata rather than depend on data held in the archival tier. Keeping operational deduplication activity local avoids making archive retrieval a prerequisite for normal backup processing and preserves Glacier’s intended role as a durable, infrequently accessed archive destination. Amazon describes Glacier storage classes as archival services with retrieval workflows and timing considerations, reinforcing the need to minimize routine access to archived objects. See Amazon S3 Glacier documentation and AWS documentation on restoring archived objects.

QUESTION NO: 4

When configuring high availability for Granite Edge, which of the following is true?

A. MPIO is enabled on a Steelhead EX appliance to achieve high availability.

B. A pair of Steelhead CX appliances act as active/standby peers.

C. Granite Edge does not support high availability.

D. A pair of Steelhead EX appliances act as active/active peers.

E. A pair of Steelhead EX appliances act as active/standby peers.

ANSWER: E

Explanation:

A pair of Steelhead EX appliances act as active/standby peers. Granite Edge high availability is implemented with two compatible SteelHead EX appliances deployed as an HA pair at the branch. One appliance is the active peer and provides the Granite Edge services, including access to the projected branch storage and associated optimization functions. The second appliance remains in standby and maintains the state needed to take over if the active peer becomes unavailable. This design supplies service continuity while avoiding simultaneous, independent servicing of the same Granite Edge workload by both peers. Correct HA deployment therefore requires configuring the appliances as a peer pair and following the platform’s HA cabling, network, software-version, and configuration prerequisites. After a failure, the standby peer assumes the active role so branch users can continue to access the Granite-projected storage with minimal interruption. Riverbed’s product documentation portal provides the SteelHead EX and Granite configuration documentation, including deployment and high-availability guidance: Riverbed Documentation. Riverbed’s support portal also provides software documentation and release-specific implementation information for SteelHead EX appliances: Riverbed Support.

QUESTION NO: 5

Which acceptable ways Whitewater appliances can be expanded to hold additional storage? (Select 2)

A. Use external USB drives connected to the rear USB ports of Whitewater appliance.

B. Add an additional (up to 2) shelf to the Whitewater appliance head unit.

C. Adjust the data store disk size on the virtual Whitewater VM (within the ESXi host) while the virtual Whitewater appliance VM is offline.

D. Connect a 3rd party SAS drive array shelf to the Whitewater appliance head unit.

ANSWER: B C

Explanation:

Adding an additional (up to 2) shelf to the Whitewater appliance head unit is a supported hardware-growth method for applicable physical Whitewater models. The expansion shelves provide additional data-store capacity through the appliance's qualified storage expansion configuration, rather than changing the appliance's core operating role or using nonqualified peripheral storage. Capacity planning must account for the specific Whitewater model and its supported shelf limits, but the stated maximum of two shelves reflects the documented expansion approach for the relevant appliance family.

Adjusting the data store disk size on the virtual Whitewater VM (within the ESXi host) while the virtual Whitewater appliance VM is offline is the corresponding supported virtual deployment method. Whitewater Virtual Edition uses virtual disks presented by the hypervisor for its data-store capacity. Powering down the appliance before changing the provisioned virtual-disk size ensures the virtual hardware change is made safely, after which the appliance can recognize and use the expanded capacity according to its documented configuration procedure. Riverbed's product documentation is available through Riverbed Documentation; the underlying virtual-disk reconfiguration workflow is described in VMware vSphere Virtual Machine Administration.

QUESTION NO: 6

Granite appliance has deep snapshot integration with which vendors? (Select 3)

A. Hitachi Data Systems (HDS)

B. EMC

C. NetApp

D. Dell EqualLogic

E. 3PAR

ANSWER: B C D

Explanation:

Granite's deep snapshot capability was designed to coordinate appliance-side operations with storage-array snapshots, allowing a Granite deployment to create and manage application-consistent recovery points without treating the array merely as generic block storage. The supported integrated-storage set includes EMC, specifically the supported EMC storage platforms in the Granite release documentation; NetApp, through the array snapshot capabilities exposed by its storage software; and Dell EqualLogic, using its supported snapshot and replication interfaces. This integration is important because Granite can align snapshot creation with branch-office data-delivery and recovery workflows, helping maintain a consistent copy of data while reducing the amount of data that must traverse the WAN. The applicable vendor support is release-dependent, so implementations should always confirm the exact Granite software version and storage-array model against the compatibility information before configuration. Riverbed's product documentation portal provides versioned manuals and release notes for Granite and related storage-delivery products: Riverbed Documentation. Riverbed's support portal is also the authoritative location for software documentation, compatibility guidance, and downloadable release materials: Riverbed Support.

QUESTION NO: 7

What is the minimum size of the data store disk for the Virtual Whitewater appliance?

A. 5% of the cloud license capacity.

B. 10% of the cloud license capacity.

C. 15% of the cloud license capacity.

D. 20% of the cloud license capacity.

E. 25% of the cloud license capacity.

ANSWER: D

Explanation:

20% of the cloud license capacity. is the required minimum sizing for the Virtual Whitewater data store disk. The data store is the local disk area used by the appliance to stage, organize, and manage data before it is transferred to, or recovered from, the configured cloud-storage target. Sizing it as a proportion of licensed cloud capacity ensures that the virtual appliance has sufficient local working space for its supported workload and for normal cloud-backup and recovery operations.

For example, a Virtual Whitewater appliance licensed for 10 TB of cloud capacity requires at least 2 TB allocated to its data store disk. This is a capacity-planning prerequisite rather than a recommendation to allocate only currently consumed cloud space; the calculation is based on the appliance's licensed cloud capacity. Administrators should also ensure the virtual-machine storage presented for the data store meets the applicable performance and deployment requirements before configuring the appliance. Riverbed product documentation is available through the Riverbed Documentation portal, and product support resources, including software documentation access, are available from Riverbed Support.

QUESTION NO: 8

Which iSCSI authentication method can be used to authenticate an initiator when it logs in to a storage target?

A. CHAP

B. MPIO

C. VMFS

D. LRU

ANSWER: A

Explanation:

CHAP is an iSCSI authentication method used to validate the identity of an initiator during login to a storage target. Challenge-Handshake Authentication Protocol can be configured as one-way authentication, where the target validates the initiator, or mutual authentication, where both endpoints validate each other. This is separate from network reachability, LUN masking, and iSCSI discovery; those controls address different aspects of storage access. The IETF iSCSI specification describes login authentication and CHAP use in RFC 3720.

QUESTION NO: 9

If Granite appliance experiences a WAN outage, and a user requests a block that *is* contained in the blockstore, what happens?

A. Blue screen

B. I/O error

C. The read request is satisfied

D. The read request is queued

ANSWER: C

Explanation:

The read request is satisfied because the Granite Edge appliance's blockstore holds locally cached copies of data blocks that have already been obtained from the Granite Core. During a WAN outage, Granite can continue servicing read operations for data that is available in that local blockstore, so the appliance does not need to contact the Core to retrieve the requested block. This behavior is a fundamental availability benefit of Granite's edge-based data delivery architecture: local users retain access to previously cached data even when connectivity to the data center is interrupted. The requested block is therefore read directly from the local blockstore and returned to the user or application. This applies specifically to reads for blocks present in the cache; the appliance can satisfy the operation using its local stored copy while WAN connectivity is unavailable. Riverbed describes Granite as providing local performance and access to centralized storage through its edge appliance architecture. See the Riverbed Granite product information and Riverbed's Support Portal documentation resources for Granite deployment and operation details.