PassQueen
a.k.a dumpsqueen.com

RCPE Certified Professional WAN Optimization

Riverbed 830-01

Version Demo

Total Demo Questions: 52

Total Premium Questions: 520

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

Using an Interceptor appliance is an alternative to what network designs? (Select 2)

A. WCCP

B. PBR

C. Connection Forwarding

D. Fixed-target rules

E. Peering rules

ANSWER: A B

Explanation:

Riverbed Interceptor appliances provide transparent traffic interception and load distribution for SteelHead appliances without requiring the network to use router- or switch-based redirection designs. In particular, an Interceptor can identify eligible TCP flows, select an appropriate SteelHead, and redirect the flow for WAN optimization while maintaining the transparency of the original client-to-server connection. This makes an Interceptor-based deployment an alternative to a Web Cache Communication Protocol deployment, where network devices use WCCP to redirect selected traffic to optimization appliances. It is also an alternative to a policy-based-routing deployment, where routing policy changes steer traffic toward a SteelHead or other optimization device. Both approaches otherwise depend on configuration and operational coordination in the routing/switching infrastructure. Interceptor deployment centralizes the traffic-discovery and appliance-selection role in Riverbed infrastructure, which is particularly useful where changing router policy or enabling WCCP is impractical. Riverbed documentation describes Interceptor as a platform for transparent interception and distribution of optimized traffic across SteelHead resources; see the Riverbed SteelHead product information and Riverbed's product documentation portal for deployment and configuration documentation.

QUESTION NO: 2

Which of these deployments is likely to have the shortest failover times?

A. Virtual In-path WCCP deployment

B. Virtual In-path PBR deployment

C. Server-side Out-of-Path deployment

D. Serial cluster deployment

ANSWER: A

Explanation:

Virtual In-path WCCP deployment is correct because Web Cache Communication Protocol provides a router-driven mechanism for directing traffic to, monitoring, and withdrawing traffic from available optimization devices. In a virtual in-path design, the router maintains WCCP control-plane communication with the SteelHead appliance. When an appliance becomes unavailable, the router can remove its redirect assignment and stop sending new traffic to it, allowing traffic to resume through the normal forwarding path or to be distributed to another available device. This avoids dependence on a physical bypass card and makes failover primarily a function of the WCCP failure-detection and redirect timers, which can be configured to meet the environment's availability requirements. WCCP is specifically designed to support redirection to a set of service devices and to adapt the assignment when the service group membership changes. Riverbed documents WCCP as a supported traffic-redirection method for SteelHead deployments, and Cisco documents the protocol's router-to-service-device control and assignment behavior. See the Riverbed documentation portal and Cisco's WCCP configuration guide.

QUESTION NO: 3

When configuring Path Selection on Interceptor Appliance Clusters, which of the following guidelines does NOT apply?

A. Static routes must be defined on Interceptors

B. SteelHead appliances must have Subnet Side rules configured

C. Local Site Subnets need to be configured in the Topology

D. Fair Peering V2 disabled on the Interceptor appliances

ANSWER: B

Explanation:

SteelHead appliances must have Subnet Side rules configured is the guideline that does not apply to Path Selection on Interceptor appliance clusters. Path Selection is an Interceptor-cluster function that uses topology information and the Interceptor appliances' routing and peering configuration to select an appropriate SteelHead cluster for optimized traffic. It is not dependent on configuring Subnet Side rules on the SteelHead appliances.

In this design, the Interceptor cluster identifies traffic and applies the Path Selection configuration based on site and subnet information maintained in SteelCentral Controller. The relevant preparation is therefore centered on making the Interceptor cluster able to correctly recognize local and remote sites and to forward traffic according to the defined network paths. SteelHead configuration remains important for normal optimization operation, but Subnet Side rules are not a stated Path Selection prerequisite for an Interceptor appliance cluster.

Riverbed's SteelCentral Controller documentation describes the configuration requirements and workflow for Interceptor-cluster Path Selection. See the SteelCentral Controller 9.9 User Guide and Riverbed's product documentation portal.

QUESTION NO: 4

-- Exhibit –

-- Exhibit --

Refer to the exhibit. According to the log information, what does that message represent? (Select 2)

A. This Steelhead appliance is closing its out-of-band Splice with the peer Steelhead appliance 172.16.10.101.

B. This Steelhead appliance is closing its out-of-band Splice with the peer Steelhead appliance 172.16.11.101.

C. The local Steelhead appliance is 172.16.11.101.

D. The local Steelhead appliance is 172.16.12.101.

ANSWER: B D

Explanation:

The message identifies an out-of-band (OOB) splice teardown involving the peer SteelHead appliance at 172.16.11.101. Therefore, “This Steelhead appliance is closing its out-of-band Splice with the peer Steelhead appliance 172.16.11.101.” accurately describes the peer endpoint named by the log entry. An OOB splice is part of SteelHead’s optimized connection handling; logging its closure indicates that the appliance is ending that optimized splice state with the specified remote peer.

The log also identifies the reporting, local SteelHead appliance as 172.16.12.101. Thus, “The local Steelhead appliance is 172.16.12.101.” is also correct. In SteelHead logging, the local address is the appliance generating the entry, while the peer address identifies the remote SteelHead participating in the optimized connection. Correctly distinguishing local and peer addresses is important when tracing optimization-path behavior and validating connectivity between SteelHead appliances. Riverbed’s SteelHead documentation and support resources provide the operational context for connection optimization, logging, and troubleshooting: Riverbed Product Documentation and Riverbed Support.

QUESTION NO: 5

-- Exhibit –

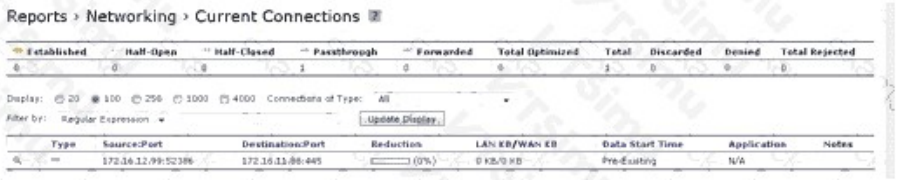

-- Exhibit --

Refer to the exhibit. What does the display represent? (Select 2)

A. This Steelhead appliance did not intercept the SYN packet for this TCP connection.

B. This Steelhead appliance did not intercept the SYN+ packet for this TCP connection.

C. This connection is not being optimized.

D. This connection will be optimized after the TCP handshake.

ANSWER: A C

Explanation:

The display indicates that this Steelhead appliance did not intercept the SYN packet for this TCP connection and that this connection is not being optimized. SteelHead optimization depends on the appliance observing and processing the initial TCP connection establishment, beginning with the client SYN. Interception enables the appliance to apply its transparent proxy behavior, identify whether a peer SteelHead is available, and establish the optimized transport mechanisms used for WAN optimization. When the initial SYN is not intercepted, the appliance has not taken ownership of the TCP setup for that flow. As a result, it cannot retroactively insert itself into the already-established session and negotiate the normal SteelHead optimization protocol for that connection. The traffic therefore continues as a pass-through TCP session rather than receiving optimization for that specific flow. This behavior is consistent with Riverbed's description of SteelHead as a transparent optimization platform that accelerates eligible TCP-based application traffic only when the flow is correctly intercepted and processed from connection establishment. See the Riverbed Documentation portal for SteelHead deployment and configuration documentation, and Riverbed's SteelHead product overview for an overview of its WAN-optimization operation.

QUESTION NO: 6

With regard to Secure Transport, what statement is true about the Maximum Segment Size (MSS)?

A. MSS is negotiated between the client and the server and is not altered in any way by the SteelHead appliances

B. MSS must be correctly configured manually on each SteelHead appliance

C. MSS is automatically negotiated by the SteelHead appliances for optimized and pass-through traffic

D. MSS must be set centrally by the SCC when configuring the Sites & Networks page

E. MSS is automatically negotiated by the SteelHead appliances, but only for optimized traffic

ANSWER: C

Explanation:

MSS is automatically negotiated by the SteelHead appliances for optimized and pass-through traffic. Secure Transport protects traffic between SteelHead appliances by encapsulating it, which introduces additional packet overhead. To prevent packets from exceeding the effective path MTU after this overhead is added, the SteelHead appliances automatically account for the encapsulation and negotiate an appropriate TCP maximum segment size. This behavior allows endpoints and applications to send traffic without requiring administrators to calculate and deploy a special MSS value manually on every appliance or centrally through SCC configuration.

The automatic handling applies whether traffic is optimized or is passed through the appliance. This is important because both traffic types can traverse the Secure Transport tunnel and therefore must use a segment size that accommodates the tunnel headers while avoiding fragmentation. By negotiating MSS dynamically, SteelHead appliances maintain reliable TCP delivery and preserve effective use of the available WAN path without imposing a separate operational configuration task. Riverbed documents this Secure Transport behavior in the SteelCentral Controller documentation: Secure Transport documentation. Additional product documentation is available through the Riverbed SteelHead support portal.

QUESTION NO: 7

Which are true about the Steelhead appliance backups taken and stored by the CMC appliance? (Select 3)

A. They are text file snapshots of the current running configuration taken at 3 a.m. (by default)

B. The CMC appliance creates hyperlinks of the backups for easy viewing with a Web browser

C. The configurations can be copied from the CMC appliance back to the Steelhead appliance via SCP

D. There is a limit of five configurations per Steelhead appliance that can be stored at any one time in a rolling queue

E. The Steelhead configuration files can be edited on the CMC appliance

ANSWER: A B C

Explanation:

The correct statements describe the CMC's scheduled configuration-backup workflow for managed SteelHead appliances. A backup is a text-based snapshot of the appliance's running configuration, and the standard scheduled-backup time is 3:00 a.m. unless the administrator changes the schedule. Because the saved configuration is textual, it provides a practical record of the settings active on the managed SteelHead at the time that the backup ran. The CMC also presents stored backups as hyperlinks in its web interface, enabling an administrator to open and review a backup directly in a browser rather than needing to retrieve it first through a command-line session. In addition, a stored configuration can be copied from the CMC back to its SteelHead appliance through SCP, which supports configuration recovery or reuse through the appliance's secure-copy mechanisms. These functions make the CMC repository useful both for routine configuration protection and for accessing a prior configuration when operational recovery is required. Riverbed product documentation and legacy software resources are available through the Riverbed Support portal; current Riverbed technical documentation is published at docs.riverbed.com.

QUESTION NO: 8

A customer would like to optimize an in-house application running over UDP. What is required in order to do this? (Choose two.)

A. Packet-mode optimization enabled in the General Service Settings.

B. Auto-Discover UDP rules correctly configured on the client-side SteelHead.

C. Fixed target rules correctly configured on the client-side and server-side SteelHeads.

D. An UDP license purchased and installed on all SteelHeads.

E. Only TCP is supported for optimization on SteelHead appliances.

ANSWER: A C

Explanation:

Packet-mode optimization enabled in the General Service Settings. is required because RiOS handles UDP optimization through packet-mode processing. Unlike TCP optimization, which can use TCP connection setup and related session information, UDP is connectionless and must be processed using the packet-mode optimization capability. Enabling this service setting permits the SteelHead to apply supported UDP optimization techniques to matching traffic.

Fixed target rules correctly configured on the client-side and server-side SteelHeads. is also required. UDP has no TCP handshake through which SteelHeads can automatically discover and establish an optimized peer relationship. Fixed-target configuration explicitly identifies the remote peer SteelHead for the UDP flow, allowing the client-side SteelHead to direct optimized traffic to the intended server-side SteelHead and ensuring that both appliances have the corresponding flow configuration. The rules must match the application's UDP traffic and identify the correct peer relationship so that packet-mode optimization can be used end to end.

Riverbed documents packet-mode and fixed-target configuration in the SteelHead User Guide. Additional product documentation and version-specific guidance are available through the Riverbed Documentation Portal.

QUESTION NO: 9

What are an Interceptor appliance ' s main functions? (Select 2)

A. Provides a way to intercept and redirect interesting traffic for optimization

B. Monitors multiple links and redirect interesting traffic over a cluster of Steelhead appliances in a logical in-path deployment

C. Replaces any existing Layer 4 switch in a given environment

D. Used to configure and work with WCCP and PBR

E. Required component Steelhead appliances in all Steelhead appliance deployments

ANSWER: A B

Explanation:

Riverbed Interceptor appliances provide scalable traffic-steering capabilities for SteelHead WAN optimization deployments. "Provides a way to intercept and redirect interesting traffic for optimization" is correct because the Interceptor identifies traffic that is eligible for optimization and transparently redirects that traffic to an appropriate SteelHead appliance. This enables optimization without requiring every connection to be physically routed through a single SteelHead in-path device.

"Monitors multiple links and redirect interesting traffic over a cluster of Steelhead appliances in a logical in-path deployment" is also correct. An Interceptor can be deployed as part of a scalable, logically in-path design, where it observes traffic across links and distributes redirected flows among multiple SteelHeads. This architecture supports load sharing and scale-out optimization while maintaining transparent handling of client-server traffic. The Interceptor's role is therefore traffic interception, selection, redirection, and balancing toward available optimization resources; the SteelHead appliances perform the WAN optimization itself. Riverbed describes SteelHead as its WAN optimization platform and provides product and documentation resources at Riverbed SteelHead and Riverbed Documentation Portal.

QUESTION NO: 10

When deploying policy-based routing for multiple Steelhead appliances in a cluster, the router should be configured to: (Select 2)

A. Set IP next-hop address to the virtual IP address of the Steelhead appliance cluster

B. Set IP next-hop addresses for each Steelhead appliance

C. Check the state level information on each Steelhead appliance via SLA or CDP

D. Auto-detect the Steelhead appliance clustering protocol

E. Check the state level information on the virtual IP address of the Steelhead appliance cluster

ANSWER: B C

Explanation:

When a deployment uses policy-based routing to direct traffic through a Steelhead cluster, the router must identify each cluster member as an individual forwarding target. Therefore, *Set IP next-hop addresses for each Steelhead appliance* is required. Providing the individual appliance addresses allows the routing policy to distribute or select traffic among the available Steelheads and preserves a usable forwarding path if a member becomes unavailable. This model does not depend on a single shared forwarding address for the cluster.

The router should also monitor the availability or health state of those individual next hops. Accordingly, *Check the state level information on each Steelhead appliance via SLA or CDP* supports resilient policy-based routing: the router can avoid steering traffic to an appliance that is no longer reachable or able to process optimized flows. Tracking must be associated with each appliance because the routing decision is made against the appliance-specific next-hop addresses. This approach aligns next-hop selection with member availability and allows traffic to continue through healthy cluster members. Riverbed's product documentation is available through the Riverbed Documentation Portal; router-side reachability tracking concepts are documented in Cisco's IP SLA documentation.

QUESTION NO: 11

-- Exhibit –

-- Exhibit --

Refer to the exhibit. A Steelhead administrator has recently decided to implement two Interceptor appliances and one Steelhead appliance on his data center. After configuring the Interceptor appliances, and the Steelhead appliance, he decided to check if the devices will work appropriately. According to the information shown, what statements are true? (Select 2)

A. There are two Interceptor appliances: 172.16.19.111 and 172.16.19.112 in parallel mode, and one Steelhed appliance. 172.16.10.101.

B. There are two Interceptor appliances: 172.16.19.111 and 172.16.19.112 in failover mode, and one Steelhead appliance. 172.16.10.101.

C. This scenario should work appropriately and the devices are configured correctly.

D. This scenario will only work if you add one more Steelhead appliance.

ANSWER: B C

Explanation:

There are two Interceptor appliances: 172.16.19.111 and 172.16.19.112 in failover mode, and one Steelhead appliance. 172.16.10.101. This identifies the intended high-availability design: the two SteelHead Interceptors form a failover pair, so one device can continue traffic-interception duties if its peer becomes unavailable. The SteelHead appliance at 172.16.10.101 is the optimization peer to which the Interceptors direct eligible traffic. An Interceptor failover deployment

provides resiliency at the traffic-discovery and redirection layer; it does not require that every deployment contain an additional SteelHead appliance solely because two Interceptors are present.

This scenario should work appropriately and the devices are configured correctly. With the Interceptor pair configured for failover and the SteelHead appliance defined as an available optimization target, the design supports normal discovery, load direction to the SteelHead, and continuity of interception services following loss of the active Interceptor. Riverbed's product documentation describes SteelHead Interceptor as the appliance used to discover and direct traffic to SteelHead appliances, including highly available deployment designs. See the Riverbed Documentation portal and Riverbed SteelHead product information.

QUESTION NO: 12

Which of the following are Steelhead Mobile Controller alarms? (Select 2)

A. Memory paging

B. Datastore errors

C. Peer version mismatch

D. SSL errors

E. File system full

ANSWER: A E

Explanation:

Memory paging and **File system full** are Steelhead Mobile Controller alarms because they identify controller-host operating-system resource conditions that can affect the availability and reliable operation of the Mobile Controller service. Memory paging indicates that the controller system is under memory pressure and is using disk-backed paging rather than keeping needed data in physical memory. This condition is operationally significant because sustained paging can reduce responsiveness of management and mobile-optimization services. File-system capacity is similarly monitored: a full filesystem can prevent the controller from writing required operational data, including logs, configuration-related data, reports, and other persistent service information. These alerts enable an administrator to investigate resource consumption and restore sufficient capacity before normal controller functions are affected. Riverbed documentation is published through the company's documentation portal, which provides product documentation and administrative guidance at Riverbed Documentation. Riverbed's support portal also provides access to product documentation, software resources, and knowledge articles for Steelhead and related products at Riverbed Support. Monitoring these host-level conditions is a core administrative practice for maintaining a healthy Steelhead Mobile Controller deployment.

QUESTION NO: 13

In Amazon Web Services (AWS) terminology, what is meant by an Instance Type?

A. Various configurations of CPU, memory, storage, and networking capacity for your instances.

B. Virtual computing environment

C. A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances

D. Preconfigured template for your instances that package your server (including the operating system and additional software)

ANSWER: A

Explanation:

Various configurations of CPU, memory, storage, and networking capacity for your instances is correct. In Amazon EC2, an instance type defines the hardware profile allocated to a virtual server. Instance types are organized into families designed for different workload needs, such as general-purpose, compute-optimized, memory-optimized, storage-optimized,

accelerated computing, and high-performance computing. The specific instance type selected determines the available vCPUs, memory capacity, networking performance, local instance storage where applicable, processor architecture, and supported features. For example, a workload requiring substantial in-memory processing may need an instance type with a larger memory-to-vCPU ratio, while a CPU-intensive application may benefit from a compute-optimized type. Selecting an appropriate instance type is therefore a core capacity and performance decision when deploying EC2 workloads. AWS provides instance-type specifications and comparison tools so architects can match the virtual machine's resources to application requirements and cost objectives. See the AWS EC2 instance types overview and the Amazon EC2 User Guide: instance types.

QUESTION NO: 14

The CMC appliance version 6.1 supports which SNMP versions? (Select 3)

A. Version 1

B. Version 2

C. Version 2c

D. Version 3

ANSWER: A C D

Explanation:

CMC appliance version 6.1 supports SNMP Version 1, Version 2c, and Version 3 for management monitoring and notification integration. SNMPv1 provides the original, basic community-string-based management protocol support. SNMPv2c retains community-based authentication while adding the operational improvements associated with the SNMPv2 protocol, making it a common choice for compatibility with established network-management systems. SNMPv3 is also supported and is the preferred choice where security is important because it introduces a user-based security model that can provide authentication and privacy (encryption) for SNMP management traffic. Supporting these three versions allows a CMC deployment to interoperate with older monitoring platforms, widely deployed SNMPv2c collectors, and security-conscious SNMPv3 management environments. This matches Riverbed's CMC documentation set for the 6.1 release and the standard SNMP version model described by the Internet Engineering Task Force. See the Riverbed Support portal for Riverbed product documentation and the IETF's SNMP architecture specification for the SNMPv3 framework.

QUESTION NO: 15

Which command shows the last boot partition and the next boot partition? (Select 2)

A. show bootvar

B. show images

C. show info

D. show bootpar

E. show system

F. show version

ANSWER: A B

Explanation:

`show bootvar` displays the appliance boot-variable information directly, including the partition from which the SteelHead most recently booted and the partition selected for the next restart. This makes it the primary operational command for confirming pending boot-partition changes, such as those made before activating a newly installed software image.

`show images` provides the installed image and partition inventory and includes boot-related partition status. It can therefore be used to identify both the partition used by the prior boot and the partition designated for the subsequent boot, while also associating those partitions with their available software images. Using these two commands together is a sound operational practice: boot variables show the configured boot selection, and the image display provides the installed-image and partition context needed to validate the selection before rebooting. Riverbed publishes product documentation and command-reference material through its Documentation Portal; authenticated software and support resources are available through the Riverbed Support Portal.

QUESTION NO: 16

During normal operation of a WCCP cluster, Hash or Mask prevent asymmetry; however, in periods of failover this may not be the case. What should you do to prevent asymmetry during a failover?

A. Disable “Allow Neighbor Failure”

B. Apply QoS to the OOB connections

C. Configure a backup Designated Web Cache (DWC)

D. Ensure fully meshed Connection Forwarding is configured

ANSWER: D

Explanation:

Ensure fully meshed Connection Forwarding is configured. WCCP Hash and Mask assignments ordinarily provide deterministic redirection, so packets for a given flow are sent to the SteelHead that owns that flow. During a failure event, however, a WCCP service group can be reassigned before all in-flight connections have ended. Packets belonging to an existing optimized connection can consequently be redirected to a different SteelHead from the one that originally established and maintains the connection state.

Riverbed Connection Forwarding allows a SteelHead that receives a packet for a connection owned by another cluster member to forward that packet to the owning SteelHead. A fully meshed configuration is required so that every participating SteelHead can reach every other potential connection owner. This preserves connection affinity while membership and bucket assignments converge, avoiding the asymmetric delivery that can interrupt optimized sessions during failover. The configuration must use the appropriate peer reachability and forwarding settings for all members of the WCCP cluster, including members that might become active after a failure.

See Riverbed’s SteelHead documentation and software resources and the WCCP protocol specification, RFC 3748, for WCCP service-group assignment and failover behavior.

QUESTION NO: 17

Branch office users live streaming video URL fragment requests are not optimized even though Microsoft Silverlight stream splitting is enabled. What could be the causing this? (Select 3)

A. The client-side Steelhead appliance optimization service wasn't restarted.

B. HTTP optimization is disabled on the client-side Steelhead appliance.

C. The server-side Steelhead appliance optimization service wasn't restarted.

D. HTTP optimization is disabled on the server-side Steelhead appliance.

ANSWER: A B D

Explanation:

Microsoft Silverlight stream splitting depends on the HTTP optimization feature being active across the optimized path. The client-side SteelHead must have HTTP optimization enabled so that it can recognize the branch client’s HTTP requests and

apply the appropriate Silverlight stream-splitting behavior. The server-side SteelHead must also have HTTP optimization enabled because it participates in the optimized HTTP flow toward the content server and supports the end-to-end application optimization exchange.

In addition, enabling or changing the relevant client-side HTTP/Silverlight optimization settings requires the client-side SteelHead optimization service to be restarted before the configuration is applied to new optimized connections. If that service restart was not performed, URL fragment requests can continue to be handled without Silverlight stream splitting despite the feature appearing enabled in the configuration. These requirements ensure the SteelHead can identify the fragmented live-stream requests and apply HTTP optimization consistently between the branch and data-center appliances. Riverbed product documentation and software resources are available through the Riverbed Documentation portal and the Riverbed Support portal.

QUESTION NO: 18

You have a SteelCentral Controller for SteelHead Mobile with a policy that includes auto-discovery rules, enabled Location Awareness and set a minimum latency of 10ms. Your client is now in the branch office, has received this policy and the nearest SteelHead is 2 ms away. What will happen to connections?

A. The branch steelhead will pass through the traffic because of the RVBD in-path rule.

B. The SteelHead Mobile client will bypass the branch SteelHead using Enhanced Auto Discovery.

C. The SteelHead Mobile client will not optimize, letting the branch SteelHead do the optimization.

D. The branch SteelHead will shut-down the SteelHead mobile client services as they are not needed in branch.

E. The SteelHead Mobile client will optimize as normal, sharing information with the branch SteelHead.

ANSWER: C

Explanation:

The SteelHead Mobile client will not optimize, letting the branch SteelHead do the optimization. Location Awareness is intended to prevent SteelHead Mobile from attempting WAN optimization when the endpoint is connected at a location already served by a nearby SteelHead appliance. The configured minimum latency establishes the proximity threshold used for this decision. Because the client's measured latency to the branch SteelHead is 2 ms, which is below the configured 10 ms minimum latency, SteelHead Mobile recognizes that it is effectively local to that appliance. It therefore suppresses its own optimization activity for applicable connections and relies on the branch SteelHead to provide WAN optimization for traffic traversing the WAN from that office. This avoids unnecessary or conflicting optimization by both the endpoint client and the local branch appliance, while retaining the benefit of the branch SteelHead's optimization path. Riverbed's SteelHead documentation and product material describe SteelHead's role in optimizing WAN traffic and provide the documentation resources used to configure SteelHead and SteelHead Mobile deployments: Riverbed Documentation Portal and Riverbed SteelHead product information.

QUESTION NO: 19

Which of the following control messages are NOT used by WCCP? (Select 2)

A. PEEK_A_BOO

B. I_SEE_YOU

C. REDIRECT_ASSIGN

D. REMOVAL_QUERY

E. KEEPALIVE

ANSWER: A E

Explanation:

PEEK_A_BOO and KEEPALIVE are not WCCP control-message types. WCCP version 2 uses a defined set of control messages to establish communication between web caches and WCCP routers, maintain service-group membership, distribute redirect assignments, and remove a cache from service. The protocol's named messages include HERE_I_AM, sent by a cache to announce its presence; I_SEE_YOU, sent by a router in response; REDIRECT_ASSIGN, which communicates assignment information for redirected traffic; and REMOVAL_QUERY, used in the cache-removal process. Although periodic protocol exchanges provide ongoing membership and state awareness, WCCP does not define a control message literally named KEEPALIVE. Likewise, PEEK_A_BOO is not part of the WCCP protocol vocabulary. These selections are therefore the two message names that do not belong to WCCP. The complete WCCP version 2 message definitions and control-message formats are specified in RFC 3040: Web Cache Communication Protocol Version 2.0. Riverbed's SteelHead documentation also describes WCCP as an interception and redirection integration mechanism with network infrastructure: Riverbed Documentation Portal.

QUESTION NO: 20

You are troubleshooting a connection which is being passed-through. An in-path rule to optimize it is present, both SteelHeads are working and not in admission control. You examine the logs and find:

[intercept.WARN] asymmetric routing between 10.11.111.19:33262 and 10.11.25.23:5001 detected (no SYN/ACK)

What is the reason for the asymmetry?

A. Complete asymmetry

B. SYN Retransmit

C. Server-side asymmetry

D. Client-side asymmetry

E. Multi-SYN Retransmit

ANSWER: C

Explanation:

Server-side asymmetry is indicated by the no SYN/ACK message. For a TCP connection to be intercepted and optimized by a SteelHead pair, the handshake traffic must be visible in both directions: the client-originated SYN travels toward the server, and the server-originated SYN/ACK returns toward the client. This warning shows that the SteelHead observed the initial SYN but did not observe the corresponding SYN/ACK on the expected return path. Therefore, traffic is asymmetric on the server-side portion of the optimized path: the return packet from the server is taking a path that bypasses the relevant in-path SteelHead, or otherwise does not return through the expected interception path.

This is a routing and path-visibility condition, rather than an admission-control or optimization-rule issue. Investigation should focus on the server-side default gateway, return routes, policy-based routing, load balancers, firewalls, and any redundant WAN paths that could direct the server-to-client response away from the SteelHead. Riverbed's documentation portal provides product configuration and troubleshooting guidance at Riverbed Documentation. The TCP three-way handshake behavior underlying the SYN/SYN-ACK observation is specified in RFC 793.

QUESTION NO: 21

A customer with a new physical in-path installation decides to turn on the optimization service during lunch time on a working day. This is done so that he can see if the users perceive any performance improvement. When starting the optimization service, he does not enable the in-path kickoff option. Assuming there were several already active connections when he starts the optimization service, how will the Steelhead appliances handle these existing connections?

A. The Steelhead appliance will optimize all connections, including pre-existing

B. The Steelhead appliance will pass-through all pre-existing connections, leaving them un-optimized

C. The Steelhead appliance will reset all pre-existing connections, thus automatically changing their optimization status from un-optimized to optimized

D. The Steelhead appliance will list the connections as “pre-existing” on the Current Connections page

E. Both B and D

F. Both C and D

ANSWER: E

Explanation:

The combined outcome of passing through all pre-existing connections without optimizing them and listing them as “pre-existing” on the Current Connections page is correct. A SteelHead appliance determines whether to optimize a flow when that TCP connection is established and the appliance observes its initial connection setup. Connections that were already active before the optimization service was enabled have not gone through that initial optimization decision process. Therefore, they continue as pass-through traffic for the lifetime of those existing sessions rather than becoming optimized in the middle of the connection.

In-path kickoff is the mechanism intended to interrupt established connections so that endpoints reconnect and the newly created flows can be evaluated for optimization. Because kickoff was not enabled in this scenario, SteelHead preserves the active user sessions instead of resetting them. The Current Connections page identifies these sessions as pre-existing, giving the administrator visibility into why they are not yet receiving optimization. New connections created after the optimization service starts can be intercepted and optimized normally, provided they meet the applicable peering, policy, and optimization criteria. Riverbed product documentation and support resources are available through the Riverbed Documentation portal and the Riverbed Support portal.

QUESTION NO: 22

RiOS v7.0 and later, the packet-mode-optimization feature can: (Select 2)

A. optimize TCP IPv6

B. optimize UDP IPv6

C. optimize UDP IPv4

D. optimize TCP IPv4

ANSWER: A C

Explanation:

In RiOS 7.0 and later, packet-mode optimization extends optimized transport handling to two specific protocol/address-family combinations: TCP carried over IPv6 and UDP carried over IPv4. Therefore, *optimize TCP IPv6* and *optimize UDP IPv4* are the applicable capabilities. Packet-mode optimization is designed for traffic that is handled as packets rather than through the standard full TCP proxy optimization path. This enables SteelHead appliances to apply relevant WAN-optimization behavior where supported while preserving the characteristics required by the applicable protocol and IP version. The feature’s support matrix is intentionally specific: it is not a blanket statement that every combination of TCP or UDP with IPv4 or IPv6 receives packet-mode optimization. When configuring an in-path rule or validating an optimization policy, the traffic’s transport protocol and IP addressing family must therefore be considered together. Riverbed’s SteelHead documentation and software resources provide the authoritative feature documentation and release-specific behavior: Riverbed SteelHead software documentation and Riverbed SteelHead product information.

QUESTION NO: 23

You need to perform a tcpdump of traffic entering the Steelhead appliance ' s lan0_0 interface. Which of the following demonstrate the proper command format to perform the tcpdump and capture to a file?

A. tcpdump –p lan0_0 –S 300 –w lancap.cap

B. tcpdump –i lan0_0 –s 300 –w lancap.cap

C. tcpdump –i lan0_0

D. tcpdump –I lan0_0 –s 300

E. tcpdump –i lan0_0 -3 300 –r lancap.cap

ANSWER: B

Explanation:

`tcpdump -i lan0_0 -s 300 -w lancap.cap` is the correct command format because it supplies all elements required for an interface-specific packet capture written to a file. The `-i lan0_0` argument selects the SteelHead appliance's `lan0_0` network interface as the capture source. The `-s 300` argument sets the snapshot length, limiting each captured packet to its first 300 bytes; this helps control capture-file size while retaining packet headers and an initial portion of payload data useful for troubleshooting. The `-w lancap.cap` argument writes packets in tcpdump/pcap format to the named capture file rather than displaying decoded packet output in the terminal. Such a pcap file can subsequently be examined with tcpdump or opened in packet-analysis tools such as Wireshark. Tcpdump documents `-i` as interface selection, `-s` as the snapshot-length setting, and `-w` as writing raw packets to a savefile. Riverbed's documentation portal provides SteelHead product documentation and operational guidance for appliance troubleshooting.

References: tcpdump manual page; Riverbed Documentation.

QUESTION NO: 24

Refer to the exhibit.

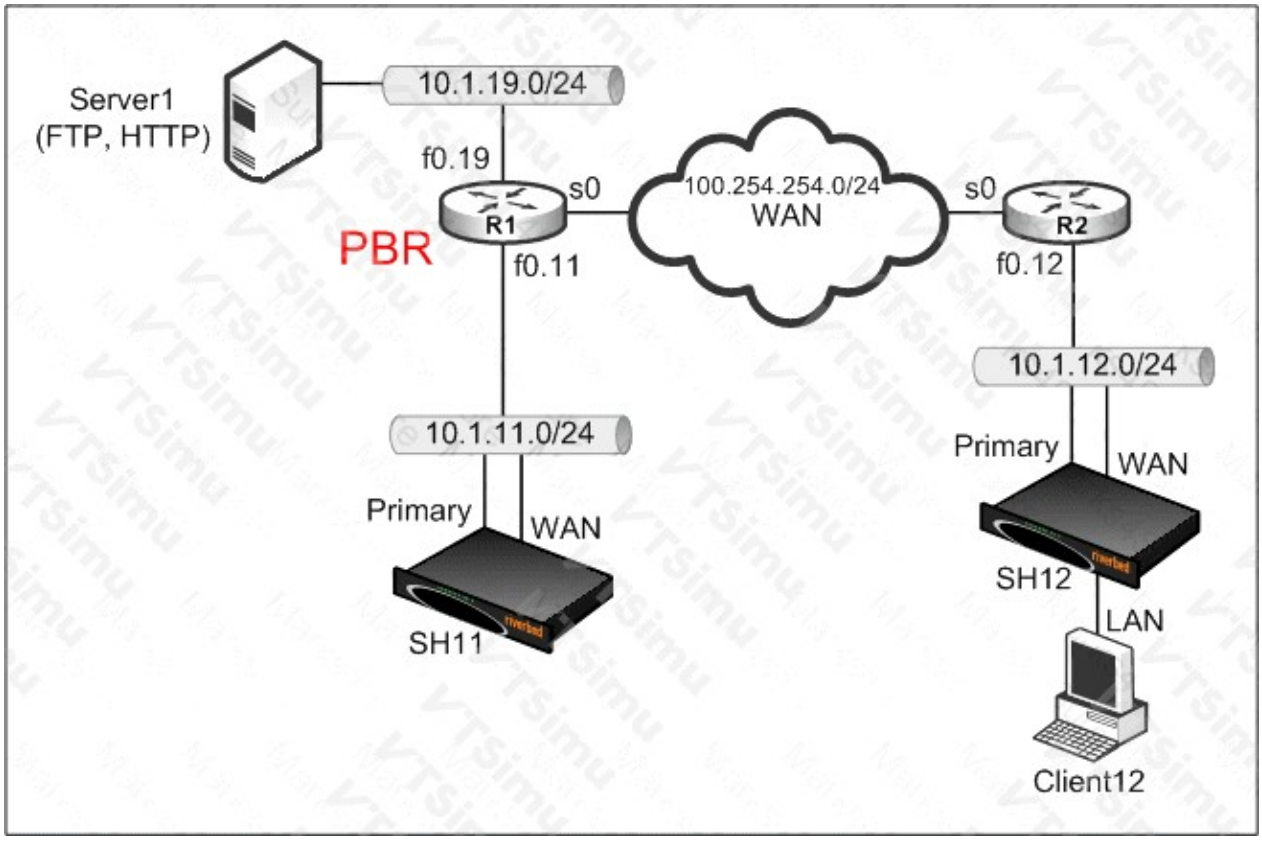

The customer would like to prevent a "black-holing" situation when using PBR when SH11 goes down. The Steelhead appliances are directly connected to the routers. What options are available to determine state level status of the Steelhead appliance? (Select 3)

A. CDP

B. NBAR

C. Interface Tracking

D. HTTP

E. ICMP

ANSWER: A C E

Explanation:

CDP, Interface Tracking, and ICMP provide practical ways for the directly connected router to determine whether the Steelhead is available and to make policy-based routing decisions that avoid sending traffic to an unavailable optimization device. CDP can validate the presence of the directly attached neighbor at Layer 2 and is useful where Cisco Discovery Protocol is enabled on both router and Steelhead interfaces. Interface Tracking lets the router associate PBR behavior with the operational state of the physical or logical interface toward the Steelhead; if that link loses carrier or otherwise transitions down, tracked routing behavior can withdraw the next-hop path. ICMP supports active reachability testing, commonly through an IP SLA operation tied to an object track. This detects a Steelhead or path failure even when the Ethernet interface remains electrically up, such as during an appliance, software, or forwarding failure. Combining link-state tracking with an active ICMP probe gives PBR a dependable fail-open mechanism, allowing traffic to use its normal routing path when the Steelhead is not usable. Cisco documents object tracking and IP SLA integration for conditional policy-routing decisions in its IP SLA Configuration Guide; Riverbed product documentation is available through the Riverbed Documentation Portal.

QUESTION NO: 25

Where in the Steelhead appliance can you verify IP address settings? (Select 2)

A. Interface Statistics

B. View the configuration

C. Neighbor Statistics

D. QoS Statistics

E. Log messages

ANSWER: A B

Explanation:

Interface Statistics and **View the configuration** are the appropriate places to verify SteelHead IP addressing. Interface Statistics provides the operational, interface-level view needed to confirm the currently active network configuration for a physical or logical interface, including its address-related status. This is useful when validating that the management, primary, auxiliary, or in-path interface is up and using the expected network settings. View the configuration provides the configured appliance settings, allowing an administrator to inspect the saved interface configuration and confirm the IP address, subnet mask or prefix, default gateway, and related network parameters that will govern appliance connectivity. Using both views is a sound operational practice: the configuration view verifies the intended persistent settings, while the interface status view validates what is active on the appliance. Riverbed's SteelHead documentation organizes appliance administration around network-interface configuration and monitoring, including the distinction between configured settings and interface operational information. See the Riverbed Documentation portal for SteelHead product documentation and the Riverbed Support documentation library for version-specific SteelHead Management Console administration guides.

QUESTION NO: 26

MAPI optimization configuration should be applied.

A. All SteelHeads except those configured for failover.

B. On both the client-side and server-side SteelHeads.

C. All SteelHeads except those configured for failover and data store sync.

D. On the server-side SteelHead

E. On the client-side SteelHead

ANSWER: B

Explanation:

MAPI optimization configuration should be applied **on both the client-side and server-side SteelHeads**. MAPI is a stateful Microsoft Exchange client/server protocol, and SteelHead optimization relies on the paired appliances participating in the optimized flow. The client-side SteelHead recognizes and processes the MAPI traffic originating from Outlook clients, while the server-side SteelHead participates at the Exchange-server end of the connection. Enabling the feature at both ends ensures that each appliance advertises and uses the required MAPI optimization capability for connections traversing the SteelHead pair.

This paired configuration is important because an optimized MAPI session is not merely a local traffic classification setting; it is a coordinated protocol optimization between the two peer SteelHeads. With MAPI optimization configured at both locations, the appliances can reduce the effects of WAN latency on Exchange communications through application-aware handling of the MAPI exchanges. Riverbed documentation and software guidance are available through the Riverbed Documentation portal and the Riverbed Support portal, where the applicable SteelHead/RiOS release documentation should be consulted for version-specific configuration details.

QUESTION NO: 27

WCCP protocol has a methodology of load balancing with two or more Steelhead appliances. What are the maximum number of buckets able to be used for the MASK assignment?

A. 8

B. 32

C. 64

D. 128

E. 256

ANSWER: D

Explanation:

128 is the maximum number of buckets usable with WCCP MASK assignment for SteelHead load balancing. WCCP distributes redirected traffic by calculating a bucket value from selected packet fields and assigning those bucket ranges to members of the WCCP service group. With MASK assignment, the appliance and router use the configured mask/value information to determine which member receives traffic, allowing multiple SteelHead appliances to share optimization work while maintaining a deterministic traffic distribution.

For the SteelHead WCCP implementation covered by this certification objective, MASK assignment is limited to 128 usable buckets. This limit is important when planning load distribution: the available bucket space is divided among participating SteelHeads, and the allocation should reflect appliance capacity and expected traffic volume. Keeping the assignment consistent across the WCCP router and all group members ensures that traffic is redirected predictably and that optimization flows remain associated with the appropriate appliance.

The WCCP version 2 protocol defines the assignment mechanisms and mask/value format used for this behavior; see RFC 3740: Web Cache Communication Protocol Version 2.0. Cisco's WCCP version 2 configuration documentation also describes the protocol's service-group assignment operation: Cisco IOS WCCP Version 2 documentation.

QUESTION NO: 28

Which of the following configuration policies are supported on the CMC appliance? (Select 4)

A. Security Policy

B. Optimization Policy

C. System Setting Policy

D. License Policy

E. Networking Policy

ANSWER: A B C E

Explanation:

The CMC appliance supports centrally defined configuration policies that can be assigned to managed SteelHead appliances, allowing common settings to be applied consistently across sites. **Security Policy** is supported for centrally managing security-related appliance configuration, including the controls used to protect management and data-path operation. **Optimization Policy** is supported because CMC is designed to distribute WAN-optimization settings, enabling standardized optimization behavior across groups of managed appliances. **System Setting Policy** is also supported for common system-level configuration that administrators want to maintain consistently throughout the managed deployment. **Networking Policy** is supported for centrally managed network configuration, helping administrators apply an approved networking baseline to multiple appliances instead of configuring each device independently.

These policy types are part of CMC's centralized configuration-management function: policies define reusable settings, and the administrator assigns them to selected devices or groups. This approach reduces configuration drift and makes large-scale SteelHead administration more predictable. Riverbed's product documentation and support resources provide the applicable appliance administration and configuration-policy guidance: Riverbed Documentation Portal and Riverbed Support Portal.

QUESTION NO: 29

When bi-directional network address translation (NAT) is used between the Steelhead appliances, optimization is not working. A workaround to this is: (Select 2)

A. Use fixed-target rules on the server-side Steelhead appliance

B. Enable server-side out-of-the-path on the server-side Steelhead appliance

C. Enable destination oobtransparency

D. Enable Full Transparency

E. There is no work around for a double NAT ' d situation

ANSWER: C D

Explanation:

Enable destination oobtransparency and Enable Full Transparency. Bi-directional NAT can cause the client-side and server-side Steelhead appliances to observe different address information for the same TCP flow. Since SteelHead optimization depends on correctly identifying and pairing the optimized connection across both appliances, address translation in both directions can prevent the peer appliances from matching the flow as expected. Full Transparency enables the SteelHead appliance to preserve the original endpoint addressing behavior needed for the optimized connection, rather than relying solely on the translated addresses visible at its local interface. Destination oobtransparency complements this behavior for out-of-band handling by ensuring that the destination information used by the SteelHead remains transparent and consistent when the optimized connection is established across translated network boundaries. Together, these settings provide the required NAT-awareness workaround so the SteelHead peers can establish and maintain optimized sessions despite double, or bi-directional, NAT. Riverbed's product documentation and support resources provide deployment and configuration guidance for SteelHead transparency and NAT scenarios; see the Riverbed Documentation Portal and the Riverbed Support Portal.

QUESTION NO: 30

What is the default In-path rule on an Interceptor appliance?

A. Optimize All

B. Pass-through All

C. Auto All

D. Redirect All

ANSWER: D

Explanation:

Redirect All is the default in-path rule on a Riverbed SteelHead Interceptor appliance. The Interceptor is designed to sit transparently in the network path, identify eligible TCP flows, and redirect those flows to an appropriate SteelHead appliance for WAN optimization. A default redirect-all policy provides the baseline behavior needed for the Interceptor to distribute traffic among SteelHead appliances or SteelHead groups after the appliance has been deployed in-path and its peering and redirection configuration has been established.

In-path rules are evaluated to determine how traffic matching specified source, destination, port, VLAN, and other criteria is handled. Administrators can add more-specific rules above the default rule when they need to exempt traffic, direct it to particular optimization resources, or apply a different forwarding behavior. If no more-specific rule matches, the default **Redirect All** rule ensures that traffic is handled by the Interceptor's normal redirection process rather than requiring an administrator to create an initial catch-all rule manually.

Riverbed's product documentation resources are available through the Riverbed Documentation Portal and the Riverbed Support Portal.

QUESTION NO: 31

Which deployment options are available to Virtual SteelHead on ESXi?

A. Physical In-path (with Riverbed bypass NIC)

B. Virtual In-path

C. Out-of-path

D. Server-side Out-of-Path

E. All of the above

ANSWER: A B C

Explanation:

Virtual SteelHead on VMware ESXi supports several network-placement models so that WAN optimization can be adapted to the traffic-forwarding design and availability requirements of the virtualized site. **Physical In-path (with Riverbed bypass NIC)** is supported where the ESXi host uses a compatible Riverbed bypass network interface. This placement puts optimized traffic directly through the Virtual SteelHead while allowing the bypass hardware to preserve network connectivity if the host or Virtual SteelHead is unavailable. **Virtual In-path** is also supported and places the appliance logically in the traffic path through the ESXi virtual networking configuration. This is useful when traffic can be directed through the appliance by virtual switches, port groups, or appropriate network routing. **Out-of-path** deployment is available when traffic is redirected to the Virtual SteelHead rather than physically or virtually traversing it inline. Redirection can be provided by network mechanisms appropriate to the deployment, enabling optimization without placing the appliance directly in the forwarding path. These supported models give administrators flexibility to select a design that fits the topology while retaining SteelHead optimization capabilities. Riverbed product documentation and deployment resources are available from Riverbed Documentation and Riverbed Support.

QUESTION NO: 32

In a Steelhead deployment, the customer decides to use Full Transparency and block all incoming/outgoing TCP connections on port 7800 on their firewall. However, they are finding all TCP connections are passing through the Steelheads (even ones that should be optimized via the in-path rule set).

If Full Transparency uses the original IP addresses and TCP ports for inner connections, what is needed to fix the problem?

A. â€œno in-path enableâ€□; since we don't really need in-path in this case

B. â€œdestination oobtransparencyâ€□; since the OOB (out-of-band) splice needs transparency also

C. â€œinner connection max 50â€□; since the default inner connection pool count of 20 is insufficient

D. No command needed, the connections are really optimized even though we report connections are passed through

ANSWER: B

Explanation:

`destination oobtransparency`; since the OOB (out-of-band) splice needs transparency also is correct. SteelHead optimization normally requires the peer appliances to establish and maintain their inter-SteelHead communication, commonly associated with TCP port 7800. In this scenario, the firewall policy prevents ordinary port-7800 traffic from traversing the path. Although Full Transparency preserves the original client and server addressing and ports for the inner, optimized connection, the out-of-band splice must also be configured for transparent handling. Applying `destination oobtransparency` to the relevant in-path rule enables the required transparent OOB behavior so that the SteelHeads can form the splice without relying on an openly permitted, conventional TCP/7800 flow across the firewall. Once the peer relationship and splice can be established under this transparent mode, connections matching the in-path rule are eligible for optimization rather than being reported as passed through. The rule should be applied only where its traffic-match criteria identify the intended optimized flows, and the configuration should be validated with connection and peering status information on both SteelHeads. Riverbed product documentation and support resources are available through Riverbed Documentation and Riverbed Support.

QUESTION NO: 33

You are using Steelhead appliances in an Oracle Forms environment and are not seeing adequate acceleration. What could be the cause of this?

A. You need to install a separate Oracle Forms license

B. Oracle Forms optimization is enabled by default. However, you must set an in-path rule to enable this feature.

C. Oracle Forms optimization can only be configured from the command line (CLI)

D. Oracle Forms optimization only works with Oracle Applications 9i

E. Oracle Forms optimization only works with Oracle Applications 8i

ANSWER: B

Explanation:

Oracle Forms optimization is enabled by default. However, you must set an in-path rule to enable this feature. This is the relevant cause when a SteelHead deployment does not achieve the expected acceleration for Oracle Forms traffic. SteelHead uses application-specific optimization to recognize and handle Oracle Forms protocol behavior appropriately, but the optimization must be activated for the applicable traffic through an in-path rule. The rule provides the policy match that associates the observed Oracle Forms connections with the Oracle Forms optimization setting. Consequently, simply deploying SteelHead appliances and having the feature available is not sufficient if no in-path rule enables it for the traffic path, source and destination criteria, or relevant ports in use. Administrators should verify that the traffic is actually passing through the SteelHead appliances in-path and then review the in-path rule configuration to ensure Oracle Forms optimization is enabled for the relevant rule. Riverbed's product information for SteelHead is available at Riverbed SteelHead, and product documentation and version-specific configuration guidance can be accessed through the Riverbed Support Portal.

QUESTION NO: 34

-- Exhibit –

-- Exhibit --

Refer to the exhibit. A Steelhead administrator has decided to implement two Interceptor appliances and one Steelhead appliance on his data center. After configuring the Interceptor appliances, and the Steelhead appliance, he decided to check if the connections are being forwarded appropriately. What happened? (Select 2)

A. The connection from client 172.16.12.99 to server 172.19.88 was optimized by Steelhead appliance 172.16.10.101.

B. The connection from client 172.16.19.88 to server 172.16.12.99 was optimized by Steelhead appliance 172.16.10.101.

C. The connection was not optimized because of the GRE tunnel.

D. The connection was optimized because of the GRE tunnel.

ANSWER: A D

Explanation:

The connection from client 172.16.12.99 to server 172.19.88 was optimized by Steelhead appliance 172.16.10.101. In an Interceptor deployment, the Interceptors identify eligible TCP flows and steer those flows to an available SteelHead appliance for WAN optimization. The flow shown is therefore assigned to SteelHead 172.16.10.101, which performs the optimization processing for that connection.

The connection was optimized because of the GRE tunnel. Interceptors can use GRE encapsulation to forward a detected connection to the selected SteelHead while preserving the information needed for the SteelHead and peer-side optimization process. The GRE tunnel is a traffic-steering mechanism in this topology: it transports the redirected flow to the optimization appliance rather than preventing optimization. Successful use of GRE forwarding confirms that the Interceptor pair can direct the relevant connection through the SteelHead selected for the flow. Riverbed's product documentation is available through the Riverbed Documentation portal; GRE encapsulation behavior is standardized in RFC 2784.

QUESTION NO: 35

You are configuring your SteelHead appliances for end-to-end Kerberos authentication. You have already successfully joined your SteelHead to the domain.

What is the next step?

A. Configure Delegation Mode

B. Configure a Replication User on the server-side SteelHead appliance

C. Configure Transparent Mode

D. Troubleshoot domain joins and authentication

ANSWER: B

Explanation:

After the SteelHead has joined the Active Directory domain, the next step for end-to-end Kerberos authentication is to configure a replication user on the server-side SteelHead appliance. This account is used by the server-side appliance to obtain and replicate the Kerberos authentication information needed to support optimized SMB connections while preserving end-to-end Kerberos behavior. The replication user must be a valid domain account and must be configured in the server-side SteelHead's SMB/Kerberos settings so that the appliance can communicate appropriately with Active Directory and the SMB server during authentication processing. Domain membership alone establishes the SteelHead's identity and trust relationship with Active Directory, but it does not complete the configuration required for Kerberos ticket replication. Once the replication user is configured, the remaining SMB optimization and authentication settings can be validated according to the deployment requirements. Riverbed's SteelHead protocol configuration guidance describes the SMB Kerberos workflow,

including the domain-join prerequisite and the replication-user configuration that follows it. See the Riverbed SteelHead SMB configuration documentation and the Riverbed Support portal.

QUESTION NO: 36

When a WCCP router uses GRE redirection, how is redirected traffic delivered to a SteelHead appliance?

A. The router encapsulates the packet in GRE and sends it to the SteelHead

B. The router changes the destination MAC address only

C. The router converts the packet to an IPsec tunnel packet

D. The router sends only the TCP payload to the SteelHead

ANSWER: A

Explanation:

With GRE redirection, the WCCP router encapsulates the redirected packet in a GRE header and sends the encapsulated packet to the SteelHead appliance. GRE encapsulation allows the router and the cache engine or SteelHead to communicate across a routed Layer 3 network, rather than requiring them to be directly connected on the same Layer 2 segment.

The SteelHead decapsulates the redirected packet, processes it according to the configured deployment and in-path rules, and returns traffic through the applicable forwarding path. WCCP supports GRE encapsulation as a packet-forwarding method. See RFC 3748, Web Cache Communication Protocol Version 2.0 and the Riverbed Documentation portal.

QUESTION NO: 37

-- Exhibit –

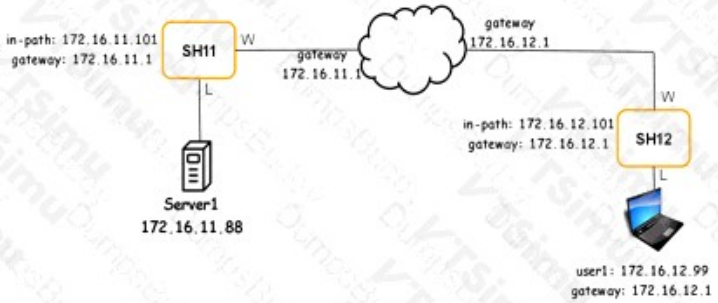

-- Exhibit --

Refer to the exhibit. On the Current Connections report, what does the number under "WAN KB" represent?

A. Segments

B. Datagrams

C. References

D. Packets

ANSWER: C

Explanation:

References is correct. SteelHead WAN optimization uses data reduction to avoid repeatedly transmitting byte sequences that have already been learned by the peer appliance. When a sequence is available in the shared data store, the sending SteelHead can transmit a compact reference rather than the original application data. The receiving SteelHead uses that reference to locate the matching stored data and reconstruct the original byte stream before delivering it locally. Consequently, the WAN-side amount shown in the Current Connections report reflects the optimized WAN representation of the traffic, including these data-store references, rather than a direct copy of all LAN-side application payload. This is why the WAN figure can be substantially lower than the corresponding LAN data figure for a connection with effective data reduction. Riverbed describes this behavior as its Scalable Data Referencing technology, which identifies redundant data and replaces it with references across the WAN. See Riverbed's SteelHead product overview and the SteelHead documentation and software support page for product documentation on optimization and reporting behavior.

QUESTION NO: 38

For a new connection, a Load Balancing rule matches but none of the target Steelhead appliances in the rule are available. What will happen to the new connection?

A. The connection is passed through

B. The next rule in the list is examined

C. The Interceptor appliance triggers an alarm

D. The Interceptor appliance discontinues redirecting any new connections and places this rule in admissions control

ANSWER: B

Explanation:

The next rule in the list is examined is correct. Riverbed Interceptor evaluates load-balancing rules in their configured order for each new connection. A match alone does not necessarily finalize redirection: the Interceptor must also be able to select an available target SteelHead appliance from the matching rule's target set. Availability is determined through the Interceptor's target monitoring and connection-admission logic, so targets that cannot accept or service the connection are not selected. When every target associated with a matching load-balancing rule is unavailable, the Interceptor continues rule processing and evaluates the next rule. This lets administrators create ordered fallback policies, such as directing a particular traffic class to a preferred SteelHead group first and then to an alternate group if the preferred group has no available members. Rule ordering is therefore important: subsequent rules can provide the intended contingency behavior for unavailable target sets. If processing eventually reaches the end without a usable redirection decision, the applicable default handling determines the connection's final treatment. Riverbed documentation and product resources for SteelHead and Interceptor deployment concepts are available from the Riverbed Documentation portal and the Riverbed SteelHead product page.

QUESTION NO: 39

Refer to the exhibit.

A user from one of your branch offices reports connections to servers are failing after a new deployment. The exhibit shows how the Steelhead appliances have been deployed. Which of the following will most likely fix the issue? (Select 2)

A. Fixed-target rule is needed on the client-side Steelhead appliance

B. Change the WCCP to use Redirect Out instead of Redirect In

C. Remove Redirect-IN from the LAN at the Branch Office

D. Add Redirect-IN on the WAN at the HQ

E. Remove fixed-target rule from the server-side Steelhead appliance

ANSWER: A D

Explanation:

A fixed-target rule is needed on the client-side Steelhead appliance because the appliance must have an explicit mapping for traffic whose server-side peer cannot be selected through normal auto-discovery. A fixed-target rule identifies the appropriate peer Steelhead to which optimized connections should be directed, allowing the client-side appliance to establish the correct optimized path instead of attempting an unsuitable discovery or connection path. Riverbed describes SteelHead as using deployment-aware mechanisms to intercept, redirect, and optimize application traffic between appliances; the peer-selection configuration must therefore agree with the network topology. See the Riverbed product documentation portal and Riverbed SteelHead overview.

Adding Redirect-IN on the WAN at the HQ supplies the missing interception direction at the data-center side of the path. Redirect-IN causes eligible traffic arriving on that interface to be directed through the SteelHead, ensuring that the HQ-side appliance receives the flows required to complete the optimized connection handling. Together, the client-side fixed target and HQ WAN-side Redirect-IN create a consistent forwarding and interception path from the branch client toward the HQ servers. This restores connectivity while preserving the intended SteelHead optimization workflow.

QUESTION NO: 40

How many failover buddies can be configured for a single Interceptor?

A. 1

B. 2

C. 3

D. 4

E. No limit

ANSWER: C

Explanation:

3 is correct. Riverbed SteelHead Interceptor supports a failover configuration in which an individual Interceptor can identify as many as three other Interceptors as failover buddies. This permits the appliance to participate in a resilient, multi-appliance deployment rather than being limited to a simple two-device high-availability pair. The configured buddies exchange the information needed to recognize an appliance failure and to preserve interception availability when a peer is no longer operational. Consequently, one Interceptor plus its three configured buddies can form a four-member failover arrangement. This limit is important when planning Interceptor placement, capacity, and failure domains: the peer relationships must be configured consistently among the participating appliances so that failover behavior is predictable. Riverbed's product documentation and software resources are available through the Riverbed Documentation portal, while version-specific manuals and technical articles can be obtained from Riverbed Support. The three-buddy maximum is the relevant platform configuration limit, not a count of the total appliances including the local Interceptor.

QUESTION NO: 41

The Optimization service has been disabled on a Steelhead appliance. Which of the following events occur as a result?

A. LAN and WAN interfaces stop intercepting

B. Primary interface becomes disabled

C. AUX interface becomes disabled

D. The Steelhead appliance reboots

ANSWER: A

Explanation:

When the Optimization service is disabled on a SteelHead appliance, the appliance stops intercepting traffic on its configured LAN and WAN interfaces. Traffic can continue through the network path, but it is no longer diverted into the SteelHead optimization engine for TCP optimization, data reduction, transport-streamlining, or application-specific processing. This behavior is useful for operational troubleshooting and maintenance because it allows an administrator to remove WAN optimization from the traffic path without requiring a reboot or necessarily taking the appliance's management functions offline. The interface links themselves remain physically and administratively available for their normal network roles; the change specifically affects the interception and optimization service. Re-enabling the Optimization service restores interception so that eligible traffic can again be processed according to the appliance configuration, including in-path deployment settings and configured optimization policies. Riverbed's product documentation portal provides SteelHead deployment and administration material describing service operation and traffic processing: Riverbed Documentation. Riverbed's support portal also provides product documentation and technical resources for SteelHead appliances: Riverbed Support.

QUESTION NO: 42

Select the command needed to add support for virtual in-path to your Steelhead appliance so that you can use WCCP or PBR with a Steelhead appliance.

A. in-path inpath0_0 enable

B. in-path oop enable

C. ip virtual enable

D. in-path virtual enable

E. in-path wccp_pbr enable

ANSWER: B

Explanation:

The command **in-path oop enable** is the correct SteelHead CLI configuration command because it enables out-of-path (OOP) support, also called virtual in-path support. In an OOP deployment, the SteelHead is not transparently cabled inline between network devices. Instead, a router or switch redirects selected traffic to the appliance using policy-based routing (PBR) or Web Cache Communication Protocol (WCCP). Enabling OOP support prepares the SteelHead in-path subsystem to receive and process this redirected traffic through its virtual in-path interfaces, while preserving the appliance's ability to optimize TCP flows and return them correctly to the network path. This setting is therefore required before configuring the WCCP service or PBR redirection policy that identifies which traffic should be sent to the SteelHead. Riverbed documentation describes WCCP and PBR as interception mechanisms for out-of-path SteelHead deployments and documents the associated virtual in-path configuration. See the Riverbed documentation portal at docs.riverbed.com and Riverbed Support resources at support.riverbed.com.

QUESTION NO: 43

Which of the following authentication methods are supported by the Steelhead appliance?

A. RADIUS and Local

B. TACACS and Local

C. TACACS, RADIUS, and Local

D. TACACS, Kerberos, and Local

E. TACACS, RADIUS, Kerberos, and Local

ANSWER: C

Explanation:

TACACS, RADIUS, and Local is correct because Steelhead appliances support administrator authentication through the appliance's own local user database as well as centralized AAA services using RADIUS or TACACS+. Local authentication provides access through accounts defined directly on the Steelhead appliance, which is useful for initial setup, emergency access, and deployments without an external identity service. RADIUS and TACACS+ allow organizations to integrate administrative access with centralized authentication infrastructure, helping apply consistent credentials, access control, and audit practices across network devices. In SteelHead documentation and interface terminology, TACACS commonly refers to TACACS+, the protocol used for device-administration AAA integration. The appliance authentication configuration lets an administrator select the applicable authentication method and define the external server settings, while retaining local access behavior appropriate to the configured policy. Riverbed's product documentation portal is available at docs.riverbed.com, and SteelHead product information is available at Riverbed SteelHead.

QUESTION NO: 44

What enables SteelHead appliances to track the specific data segments (and their references) that have been sent to, or received from, the other SteelHead appliances?

A. IP Identifier

B. Segstore ID

C. TCP Sequence Number

D. Bandwidth Delay Product (BDP)

ANSWER: B

Explanation:

Segstore ID is correct because it is the identifier SteelHead uses in its segment-store mechanisms to associate a data segment with its stored representation and its peer-to-peer reference information. SteelHead WAN optimization uses data reduction by recognizing byte patterns that have already been observed, storing those patterns locally, and communicating compact references rather than repeatedly transmitting the full data. For this process to work reliably between two SteelHead appliances, each appliance must be able to identify the relevant stored segment and maintain the reference relationship for data sent to and received from its peer. The Segstore ID supports that tracking role.

This capability is part of SteelHead's data-reduction architecture, which maintains segment-store state and enables the appliances to reconstruct the original data stream from references and locally available segment data. It is therefore specifically concerned with identifying optimized data segments and their references, rather than merely describing packet-level addressing or transport sequencing. Riverbed describes SteelHead as accelerating applications through WAN optimization and data reduction in its SteelHead product information. Additional product documentation and technical resources are available through the Riverbed Support portal.

QUESTION NO: 45

Refer to the exhibit.

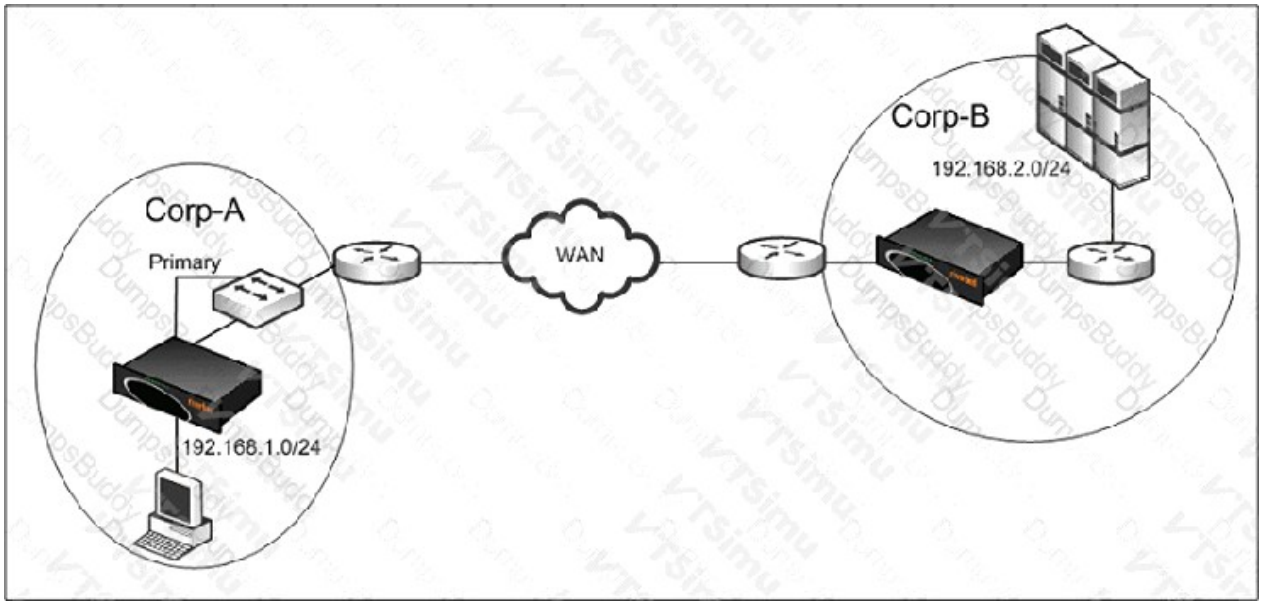

You configured CIFS prepopulation at Corp-A but it is not working. What is likely the cause of the issue?

A. CIFS prepopulation requires WCCP

B. Primary interface needs to be on a different subnet

C. CIFS prepopulation requires a fixed-target rule

D. The Primary interface needs to be on the LAN side of the Steelhead appliance

E. CIFS prepopulation requires SMBv2

ANSWER: D

Explanation:

The Primary interface needs to be on the LAN side of the Steelhead appliance. CIFS prepopulation is initiated by the SteelHead using its Primary-interface address to establish the SMB/CIFS connection to the server and retrieve the selected files into the remote SteelHead's data store before users request them. Consequently, that source address must be reachable and appropriately routed through the optimized path in the same manner as the client-side LAN traffic. In the shown Corp-A topology, placing the Primary interface on the WAN side causes the prepopulation traffic to originate from the inappropriate side of the appliance. The connection therefore does not follow the expected LAN-to-WAN, in-path optimization flow needed for the appliance to populate data for local users. Configuring the LAN-facing interface as Primary ensures that prepopulation traffic uses a valid client-side source, traverses the SteelHead correctly, and can be optimized and stored for subsequent CIFS access. Riverbed's SteelHead documentation portal provides configuration and deployment guidance for optimization services, including CIFS/SMB behavior: Riverbed Documentation. For SteelHead platform and WAN-optimization architecture information, see Riverbed SteelHead.

QUESTION NO: 46

SteelHead appliances typically use the _____ interface as a source for all communication other than optimization.

A. In-path

B. LAN

C. Aux

D. Primary

E. WAN

ANSWER: D

Explanation:

The **Primary** interface is correct. On a SteelHead appliance, the primary interface is the management and non-optimization network interface. It is typically the source interface for appliance-originated traffic such as management access, secure shell sessions, web-based administration, monitoring, logging, Domain Name System lookups, Network Time Protocol synchronization, authentication, licensing-related communication, and communication with Riverbed management systems. This separation is intentional: optimization traffic is handled through the appliance's in-path interfaces, while operational and control-plane communications generally use the primary interface. Consequently, the network must provide appropriate routing and reachability from the primary interface to the administrative and infrastructure services that the SteelHead needs. Riverbed's product documentation describes SteelHead as a WAN-optimization platform and provides the deployment and configuration guidance governing its management and network-interface behavior. See the Riverbed Documentation portal for SteelHead product documentation and the Riverbed SteelHead product page for the platform overview.

QUESTION NO: 47

In an in-path deployment, if a naked SYN packet arrives on the WAN interface of Steelhead, what probe option is added to this SYN packet for auto-discovery?

A. Option-76

B. Option-78

C. SACK

D. WSCALE

E. None of the above

ANSWER: A

Explanation:

Option-76 is the TCP option used by SteelHead auto-discovery probing. In an in-path deployment, a SteelHead can observe a normal, unmodified ("naked") TCP SYN that is traversing the WAN. To determine whether a peer SteelHead is present and able to optimize the connection, it inserts its proprietary auto-discovery probe into the SYN before forwarding it across the WAN. The probe is carried in TCP option kind 76 and lets the receiving-side SteelHead recognize that the connection is eligible for peer discovery and subsequent WAN optimization negotiation.

This mechanism is deliberately embedded in the TCP handshake so that optimization can be discovered dynamically without requiring every application flow to be explicitly redirected to a configured peer. After a compatible peer recognizes the probe, the SteelHeads can establish the appropriate optimized data path while preserving the endpoint TCP connection semantics. TCP option kinds are extension fields in the TCP header; their standardized registration context is maintained by IANA, while the SteelHead-specific use of option 76 is part of Riverbed's auto-discovery behavior. See the Riverbed product documentation portal and the IANA TCP Parameters registry for TCP option reference context.

QUESTION NO: 48

Which of the following commands will show the log level configuration? (Select 2)

A. show configuration full

B. configuration show full

C. show logging

D. logging show

E. show log level

ANSWER: A C

Explanation:

`show logging` is the direct operational CLI command for displaying the appliance logging configuration, including the configured logging severity or level and related logging settings. It is the most focused command when an administrator needs to verify the current logging behavior without reviewing unrelated appliance configuration. `show configuration full` is also valid because it outputs the complete active configuration in CLI form. The full configuration includes the logging-related configuration statements, so it can be used to identify the configured log level as part of a broader configuration review, audit, or before-and-after change comparison. These commands are complementary in practice: use `show logging` for an immediate logging-focused status check, and use `show configuration full` when validating logging settings in the context of the entire SteelHead configuration. Riverbed's SteelHead documentation and product resources describe the appliance CLI and configuration-management workflows used for viewing operational and saved configuration details. See the Riverbed Documentation portal and the Riverbed Support portal for the CLI reference corresponding to the SteelHead software release in use.

QUESTION NO: 49

What Steelhead appliance models support the installation of Microsoft Windows Server 2008 R2 package on the RSP? (Select 2)

A. 1050

B. 250

C. 550

D. 3020

ANSWER: A D

Explanation:

1050 and **3020** support installation of the Microsoft Windows Server 2008 R2 package on the Riverbed Services Platform (RSP). RSP enables supported Steelhead appliances to host qualified auxiliary services directly on the appliance hardware, alongside WAN-optimization functionality. Microsoft Windows Server 2008 R2 was supplied as an RSP package for use on platforms that met the applicable generation, resource, and RSP compatibility requirements. The 1050 and 3020 are among the supported appliance models for this package, so they can be selected when deploying that Windows Server workload through the RSP environment. In practice, the package installation also depends on having the appropriate RiOS release, available RSP resources, and the relevant package image and licensing entitlements. Because these are legacy products and operating-system releases, administrators should consult the documentation associated with the exact RiOS version in use and obtain software through Riverbed's supported channels. Riverbed's product-documentation portal is available at Riverbed Documentation, and entitlement, software-download, and lifecycle-support resources are available through Riverbed Support.

QUESTION NO: 50

How can you provide updates to multiple Steelhead Mobile clients at the same time? (Select 2)

A. Using profiles

B. Using multiple Steelhead Mobile Controllers

C. Using Group IDs

D. Using policies

ANSWER: C D

Explanation:

Using Group IDs and using policies provide the scalable, centralized mechanisms for updating multiple Steelhead Mobile clients simultaneously. A Group ID associates individual Mobile clients with a defined logical group on the Steelhead Mobile Controller. After clients are enrolled in that group, an administrator can target the group rather than administering every endpoint independently. This is particularly useful when clients share a department, location, configuration requirement, or rollout schedule.

Policies are the controller-managed definitions that specify the operational settings delivered to clients. Assigning or revising a policy lets the administrator distribute the applicable configuration to all clients covered by that policy or group during their controller communication. Together, group membership identifies the intended collection of endpoints, while policy assignment supplies the settings that are consistently delivered to that collection. This model supports controlled, repeatable configuration management across a Mobile deployment and avoids requiring a separate manual update for each laptop or remote endpoint. Riverbed documentation and product resources are available through the Riverbed Documentation Portal and the Riverbed Support Portal.

QUESTION NO: 51

The default setting for "allow failure" in the Interceptor appliance version 2.0 is: (Select 2)

A. Is configurable

B. Is not configurable

C. Enabled by default

D. Disabled by default

ANSWER: A D

Explanation:

“Is configurable” and “Disabled by default” are correct. In SteelHead Interceptor version 2.0, Allow Failure is an administrator-controlled setting rather than an immutable appliance behavior. Its purpose is to determine whether the Interceptor can permit traffic-handling failure behavior when it cannot continue normal interception or redirection processing. Because network availability and deployment requirements differ, Riverbed exposes this behavior as a configurable operational setting.

The default is disabled, which provides a conservative initial configuration: the appliance does not automatically enable the Allow Failure behavior until an administrator explicitly selects it for the deployment. This default is important in environments where traffic steering, optimization policy, and failure handling must be deliberately validated before a less restrictive failure posture is used. Administrators should review the setting in the context of the Interceptor’s deployment topology, high-availability design, and the organization’s desired behavior during appliance or path failures.

Riverbed’s product documentation resources and current support documentation are available through the Riverbed Documentation Portal and the Riverbed Support Portal.

QUESTION NO: 52

In the exhibit,

where should the WCCP redirection take place for server return traffic, assuming redirect IN is used? (Select 2)

A. A

B. B

C. C

D. D

E. E

ANSWER: B C

Explanation:

For server return traffic using WCCP *redirect in*, interception must occur on the ingress path where packets sent by the server enter the routing infrastructure on their way back toward the client. This ensures that the return direction of an optimized connection is consistently delivered to the SteelHead appliance rather than bypassing it. The two marked locations represent the applicable inbound interception points on the server-return path. Applying redirection at both locations provides coverage for the relevant routing paths, including alternate or redundant paths shown in the topology, so that return packets remain symmetric with the optimized flow.

WCCP is designed for routers and switches to redirect selected traffic to a cache engine or optimization appliance. The direction is significant: *redirect in* examines traffic as it arrives at the configured interface, making the server-facing return-path ingress interfaces the appropriate placement for this traffic direction. Riverbed WCCP deployments use this mechanism to transparently steer traffic through SteelHead appliances while preserving the required bidirectional processing of optimized sessions. For WCCP protocol behavior, see RFC 2970: Web Cache Coordination Protocol Version 2.0; Riverbed product documentation is available through the SteelHead support portal.