DUMPSARENA

Riverbed Certified Solutions Associate - Network Performance Management

Riverbed 201-01

Version Demo

Total Demo Questions: 28

Total Premium Questions: 280

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which mode should be enabled to monitor traffic in high-throughput environments such as server farms or data centers?

A. Database Performance mode

B. Application Stream Analysis boost mode

C. Server Boost mode

D. Web Transaction Analysis mode

ANSWER: B

Explanation:

Application Stream Analysis boost mode is the appropriate setting for high-throughput monitoring deployments, including server farms and data centers. In these environments, the monitoring appliance must handle a large volume of concurrent application conversations without sacrificing the visibility needed to identify application and network performance issues. Boost mode is intended to increase the system's stream-analysis capacity so that it can process traffic efficiently when connection counts and aggregate throughput are high.

Application Stream Analysis provides application-centric visibility by examining traffic streams and associating observed behavior with applications, clients, servers, and network services. Enabling its boost capability aligns the appliance's analysis behavior with the scale typical of consolidated data-center workloads, where numerous servers and clients generate substantial east-west and north-south traffic. This allows operations teams to maintain useful application-performance insight while accommodating demanding capture and analysis rates. Riverbed describes AppResponse as a platform for monitoring application and network performance and directing investigation toward observed performance conditions. See the Riverbed AppResponse product information and the Riverbed documentation portal for product documentation and configuration guidance.

QUESTION NO: 2

How do you configure a new user for the Cascade Shark appliance?

A. From the User Management section of the Web Interface

B. From the Basic Settings page of the Web Interface

C. From the Users menu of Cascade Pilot

D. From the Advanced Settings page of the Web Interface

ANSWER: A

Explanation:

From the User Management section of the Web Interface is correct because user accounts are administered locally on the Cascade Shark appliance through its browser-based management interface. The User Management area is the appliance configuration location intended for creating accounts and assigning the access privileges that determine what an authenticated user can do on the appliance. This separates identity and authorization administration from general appliance setup, so an administrator can add a new account without changing network, capture, or other operational settings. After signing in with an account that has administrative privileges, the administrator opens User Management, creates the user, supplies the required account details and credentials, and saves the configuration. The new account can then authenticate to the Shark appliance according to the privileges assigned during setup. Cascade Shark is part of Riverbed's network-performance-management and packet-analysis portfolio; Riverbed's product and support resources provide the relevant appliance documentation and software support materials. See the Riverbed packet analysis product page and the Riverbed Support portal for product documentation and support resources.

QUESTION NO: 3

Protocols that are typically monitored in TCP/IP networks by visibility tools include (Select 3):

A. SAN HBA Protocol

B. TCP Protocols such as HTTP and HTTPS

C. UDP Protocols such as DNS and Video streaming

D. IPv4 Protocols including TCP, UDP, and ICMP

E. IPv4 Protocols such as x.25, TTY and TTL

ANSWER: B C D

Explanation:

TCP/IP network visibility tools commonly inspect traffic across both transport and network layers, then associate that traffic with applications and services. TCP-based application traffic such as HTTP and HTTPS is a core monitoring target because it supports analysis of sessions, transactions, response times, retransmissions, and end-user experience. UDP traffic is also important because many services use it for latency-sensitive or connectionless communications; DNS and video streaming are representative examples that visibility platforms can identify and measure. At the network layer, IPv4 is fundamental to packet-based monitoring, and IPv4 traffic commonly carries TCP, UDP, and ICMP. Monitoring these protocols together lets an operations team correlate application behavior with transport conditions and underlying IP connectivity. This layered approach is central to network performance management: packet and flow visibility can expose application use, traffic composition, delay, packet loss, and protocol-level communication patterns. The Internet standards define IPv4 as the network-layer protocol that transports higher-layer payloads, including TCP, UDP, and ICMP, while UDP defines the connectionless transport behavior commonly observed in real-time and infrastructure traffic. See RFC 791: Internet Protocol and RFC 768: User Datagram Protocol.

QUESTION NO: 4

To which devices is it possible to send data directly to (not including flows from other Cascade devices)? (Select 2)

A. Sensor

B. Sensor-VE

C. Gateway

D. Profiler

E. Enterprise Profiler

F. Express

ANSWER: C F

Explanation:

Gateway and **Express** are the Cascade components that can be configured as direct recipients of flow data from non-Cascade sources. A Gateway is designed to collect flow records directly from network infrastructure and flow-exporting devices, providing a collection point before the information is made available for analysis in the Cascade environment. This permits organizations to collect telemetry from routers, switches, firewalls, and other exporters without requiring that the flows first traverse another Cascade appliance.

Express also supports direct flow ingestion. As an integrated, simplified Cascade deployment option, it can receive flow data directly from external exporters and make that data available for monitoring and analysis. Therefore, both Gateway and Express satisfy the condition that the data is sent directly to the device and is not flow information relayed from another Cascade device. Riverbed's product and documentation resources describe the flow-collection and network-performance-monitoring capabilities underlying these deployment roles: Riverbed SteelCentral NetProfiler and Riverbed Documentation.

QUESTION NO: 5

Authentication mechanisms supported by the Cascade Profiler arE. (Select 2)

A. Local user accounts

B. Active Directory

C. TACACS+

D. RADIUS

E. LDAP

ANSWER: A D

Explanation:

Cascade Profiler supports authentication through **Local user accounts** and **RADIUS**. Local accounts are maintained directly on the Profiler appliance and provide a self-contained way to create users and assign the appropriate Profiler roles and permissions. This is useful for administration, initial access, and environments where authentication is not delegated to a centralized identity service.

RADIUS provides centralized authentication for Cascade Profiler by allowing the appliance to submit user login requests to a configured RADIUS server. Organizations commonly use this approach to apply a consistent access-control policy and reuse centrally managed credentials for network-management systems. The Profiler configuration identifies the RADIUS service and shared-secret settings required for these authentication exchanges, while access within the product remains governed by the user authorization model.

These two mechanisms reflect the authentication choices documented for the Cascade Profiler generation covered by this exam objective. Riverbed's product documentation resources are available through the Riverbed Documentation portal; additional legacy product materials and downloads can be accessed through the Riverbed Support portal.

QUESTION NO: 6

Within NetProfile, how do User-Defined Policies differ from policies such as the Application Performance Policy, Link Congestion Policy, and Link Outage Policy?

A. User-Defined Policies are used only to compare current traffic against fixed values while other policies may allow you to compare traffic against both fixed values and a baseline, depending on the type of policy.

B. User-Defined Policies can be used to compare current traffic against a baseline of traffic built over the past several days, weeks, and months, while other policies compare current traffic against fixed values only.

C. User-Defined Policies allow you to set a minimum amount that a policy must change by before an alert will be created, whereas other policies do not.

D. There are no specific differences between User-Defined Policies and other policies, except the names of the policies.

ANSWER: B

Explanation:

User-Defined Policies can be used to compare current traffic against a baseline of traffic built over the past several days, weeks, and months, while other policies compare current traffic against fixed values only. This distinction makes User-Defined Policies appropriate when expected traffic behavior is relative to an organization's own historical pattern rather than to a universally defined threshold. NetProfiler establishes a baseline from previously observed traffic over multiple time periods, then evaluates the current observation in relation to that learned normal behavior. This lets a policy identify meaningful deviations, such as an unusual increase or decrease in traffic, even where the normal volume differs by application, location, time of day, or business cycle. In contrast, the named predefined policy types are designed around explicit configured limits for their respective conditions. Their alerts are therefore evaluated against fixed values selected by the administrator. Riverbed's NetProfiler product information describes its role in collecting and analyzing network-flow data

for performance visibility, while the Riverbed documentation portal provides product configuration and policy guidance: Riverbed NetProfiler and Riverbed Documentation.

QUESTION NO: 7

User Interface Preferences give users the ability to: (Select 3)

A. change the Navigation Bar Position.

B. change were reports are saved on Profiler.

C. change data unit from Bytes to Bits.

D. change user Time Zone.

E. search for lost templates.

ANSWER: A C D

Explanation:

User Interface Preferences in Riverbed SteelCentral NetProfiler are per-user settings intended to tailor how the product interface is displayed and interpreted for an individual operator. “change the Navigation Bar Position.” is correct because the interface preferences include presentation controls that let a user place the navigation bar in the preferred screen location, supporting different workflows and display layouts. “change data unit from Bytes to Bits.” is also correct: users can select the units used when presenting traffic and throughput values, allowing reports and charts to align with the conventions used by their operations team. “change user Time Zone.” is correct because the time-zone preference controls how timestamps and time-based data are rendered in the user interface. This is especially important when a NetProfiler appliance is accessed by users in different geographic locations, since each user can view report periods, alarms, and traffic events in an appropriate local time context. These settings personalize the UI display without changing the collected network-flow data or the appliance-wide data-retention configuration. Riverbed product documentation and support resources are available through the Riverbed Documentation portal and the Riverbed Support portal.

QUESTION NO: 8

Under Layer 4 Mappings, the available Application List settings arE. (Select 3)

A. Application Name

B. IP Address

C. Host or Groups

D. Ports

ANSWER: A C D

Explanation:

The available Application List settings for a Layer 4 mapping are **Application Name**, **Host or Groups**, and **Ports**. Together, these settings define a Layer 4-based traffic classification rule. Application Name provides the label that NetProfiler displays for traffic matching the mapping, making the resulting application category recognizable in reports, dashboards, and investigations. Host or Groups identifies the relevant endpoint scope, allowing the mapping to target traffic associated with specified hosts or configured host groups. Ports supplies the Layer 4 service criteria used to recognize the traffic, such as TCP or UDP port values and ranges. This combination is appropriate for Layer 4 classification because it uses network and transport attributes rather than application-layer payload inspection. Riverbed NetProfiler uses application definitions and mappings to turn observed flow data into meaningful application-level reporting, so a clear name plus endpoint and port scope produces useful, consistent visibility. See Riverbed’s NetProfiler product information and the Riverbed documentation portal for product documentation and configuration guidance.

QUESTION NO: 9

Network Round-Trip-Time is measured by Cascade sensors as:

A. Time from 1st ACK to 1st DAT

B. Time from (SYNACK to 1st DATA) – client delay.

C. Time from SYN to 1st DAT

D. Time from SYN to 1st ACK.

ANSWER: D

Explanation:

Time from SYN to 1st ACK is correct. Cascade sensors derive Network Round-Trip-Time from the TCP connection-establishment exchange observed in packet traffic. The client starts the exchange by transmitting a SYN, and the first acknowledgement returned by the server acknowledges that initial SYN (normally carried in the SYN/ACK segment). The elapsed time between those observed events represents the network path delay for the request-and-return portion of the handshake, without incorporating application processing or later data-transfer timing. Because the calculation uses TCP control packets that are visible at the sensor, it provides a consistent network-latency measure even before an application payload transaction is underway. This aligns with TCP's defined three-way connection establishment, in which a SYN is acknowledged with a segment containing an ACK field. See the TCP specification in RFC 793 and Riverbed's Network Performance Management product information at Riverbed NPM. In practical analysis, this metric helps separate connectivity and path-latency effects from delays attributable to the client, server, or application behavior.

QUESTION NO: 10

Which of the following are differences between the Cascade Sensor and the Cascade Shark appliances? (Select 2)

A. Cascade Shark appliance can monitor 10GE interface traffic but Cascade Sensor cannot.

B. Cascade Sensor can provide flow data to Cascade Profiler but Cascade Shark appliance cannot.

C. Cascade Shark appliance has much higher write-to-disk rates than the Cascade Sensor and is the appropriate appliance when packet capture is the primary requirement.

D. Cascade Sensor can recognize the application type by its DPI signature database but Cascade Shark appliance cannot.

ANSWER: C D

Explanation:

Cascade Shark appliances are purpose-built for high-volume, persistent packet capture and forensic packet analysis. Their storage architecture provides substantially greater sustained write-to-disk throughput than a Cascade Sensor, making a Shark the appropriate deployment choice where retaining packets for later investigation is the primary objective. This distinction is central to packet-centric troubleshooting workflows, in which analysts need to retrieve and inspect the actual packets associated with an event rather than rely only on summarized telemetry.

Cascade Sensor appliances, in contrast, perform application-aware monitoring through deep packet inspection. The Sensor uses its application-signature database to classify traffic by application, enabling application-level visibility and reporting in the Cascade environment. This DPI-based classification provides context that is useful for network performance management, such as associating observed traffic behavior with the application responsible for it. Riverbed describes its network-performance portfolio as combining packet analysis with application-aware visibility; see the Riverbed SteelCentral product information and the Riverbed Support portal for product documentation and archived technical resources.

QUESTION NO: 11

What is a benefit of Cascade's Deep-Packet Inspection? (Select 4)

A. Cascade can determine every URL that a client goes to out on the internet.

B. Cascade can identify applications at Layer7 which will aid in determining business versus non-business use-age of resources.

C. Cascade can identify applications that use multiple common ports based upon Layer7 to more readily track application dependencies.

D. Cascade can be configured with URLs, match those URL's to application names and, in order to understand which applications are utilized on web servers that host multiple web applications over the same TCP port.

E. Cascade can identify unencrypted applications such as FTP and TELNET regardless of port number, and other applications (such as Peer to Peer) that may violate a company's security policies.

ANSWER: B C D E

Explanation:

Cascade Deep-Packet Inspection provides Layer 7 application awareness, allowing traffic to be classified from protocol behavior and application signatures rather than relying only on TCP or UDP port numbers. This visibility helps operations teams distinguish business applications from recreational or otherwise non-business resource use, so application usage can be understood in meaningful business terms. It also recognizes applications that use several commonly assigned ports, improving visibility of the related traffic flows and dependencies that a port-only view can fragment.

For web-hosting environments, Cascade can use configured URL information and application naming to distinguish web applications sharing the same server and TCP port. This is particularly valuable where a single web server hosts multiple services that cannot be separated simply by IP address and port. DPI also identifies clear-text protocols such as FTP and TELNET regardless of the port on which they run, and can expose peer-to-peer application traffic that may conflict with an organization's security or acceptable-use policy. These capabilities reflect the general purpose of Riverbed's packet and application analysis: associating observed traffic with applications and delivering actionable network and application visibility. See Riverbed's Network Performance Management overview and packet analysis resources.

QUESTION NO: 12

How does Cascade performance analytics assist with Performance Monitoring?

A. By setting static thresholds for Application metrics and Interface metrics, tolerance can be determined. Cascade will use these thresholds and tolerances to report on deviations indicative of Performance problems.

B. The Customer only needs to identify their critical interfaces and applications, and Cascade will automatically baseline their behavior and report on deviations indicative of Performance problems.

C. The Performance Analytics are able to detect security threats such as host scans and worms.

D. After base-lining, Cascade can re-route congested traffic to avoid congested application delivery paths.

ANSWER: B

Explanation:

The Customer only needs to identify their critical interfaces and applications, and Cascade will automatically baseline their behavior and report on deviations indicative of Performance problems. This describes the essential performance-monitoring value of Cascade Performance Analytics: it learns normal operating behavior from observed network and application traffic, creating dynamic baselines rather than requiring an administrator to manually define a fixed acceptable value for every metric. Once the monitored interfaces and business-critical applications are identified, the platform can compare current behavior with the learned baseline and highlight meaningful deviations. This helps operations teams recognize abnormal response time, throughput, traffic-volume, or dependency behavior in the context of what is normal for that application or interface. Automatic baselining is particularly useful because normal performance can vary by time of day, day of week, application usage patterns, and location. Cascade's analytics therefore supports faster detection and investigation of performance conditions by presenting deviations that deserve attention, while allowing teams to focus monitoring on services that matter most to the organization. Riverbed's Network Performance Management overview describes its emphasis on visibility, monitoring, and faster troubleshooting across network and application environments: Riverbed Network

Performance Management. Riverbed product documentation and resources are available through the Riverbed Documentation Portal.

QUESTION NO: 13

One of the theories behind NetProfiler Services is:

A. Many organizations do not know all the aspects of a given application or service they provide, which makes it difficult to fully map that service and monitor it from end to end. NetProfiler has visibility into many different areas and can often see connections used by different parts of a service that may be unknown to different administrators. This visibility allows NetProfiler to identify all the different components of a service and provide an end-to-end view.

B. When monitoring an application, it is often necessary to think of that application as a service instead of an application. By utilizing NetProfiler Services, you are able to take your application and easily convert it to a service. This conversion allows you to treat the application differently so it can be more effectively monitored.

C. Applications are large and complex and often traverse different groups within a single organization. Without the concept of a service in NetProfiler, it would not be possible to monitor the traffic from a single application since each group would have to purchase their own NetProfiler to monitor their individual aspect of the application.

D. None of the above.

ANSWER: A

Explanation:

“Many organizations do not know all the aspects of a given application or service they provide, which makes it difficult to fully map that service and monitor it from end to end. NetProfiler has visibility into many different areas and can often see connections used by different parts of a service that may be unknown to different administrators. This visibility allows NetProfiler to identify all the different components of a service and provide an end-to-end view” correctly expresses the purpose behind NetProfiler Services. A business service commonly depends on multiple communicating tiers, hosts, networks, and protocols, often owned by separate teams. Individual administrators may understand their own portion of the environment but not every dependency involved in delivering the overall service. NetProfiler’s flow-based visibility can correlate the relevant traffic and infrastructure elements into a service-oriented view. This lets operations teams examine the health and performance of the complete delivered service rather than limiting analysis to isolated devices or a single application component. That service context helps reveal dependencies and communication paths that would otherwise be difficult to identify across organizational and technical boundaries. Riverbed positions NetProfiler as a network-performance monitoring and visibility platform, including visibility into application traffic and dependencies. See Riverbed Network Performance Management and Riverbed Support for Riverbed product and documentation resources.

QUESTION NO: 14

Which fields are commonly used to identify a TCP flow? (Select 4)

A. Source IP address

B. Destination IP address

C. Source port

D. Destination port

E. DSCP value

F. Source MAC address

ANSWER: A B C D

Explanation:

A TCP flow is commonly identified from the endpoint and transport information present in the flow tuple. **Source IP address** and **Destination IP address** identify the two communicating hosts. **Source port** and **Destination port** identify the transport-layer endpoints used by the client and server. These values, together with the IP protocol value, allow flow collectors such as Cascade Profiler to distinguish one conversation from another. DSCP is a traffic-treatment marking, while a MAC address is a Layer 2 attribute and is not normally part of the standard IP flow tuple. See the NetFlow Services Export Version 9 RFC for flow-record concepts.

QUESTION NO: 15

To drill down and analyze a file within Cascade Pilot, you should. (Select 2)

A. Select a graphic object in a View chart, and drag a View on the object.

B. Select a graphic object in a View chart, and drag a View on the trace file.

C. Select the View, right-click, and apply it with the filter associated with a graphic object.

D. Right-click in the View chart, and select "Drill-down" to select the View.

ANSWER: A D

Explanation:

Cascade Pilot supports interactive drill-down from a View chart so that an analyst can move from a summarized visualization to a more focused analysis of the traffic represented by a selected chart element. Selecting a graphic object in a View chart and dragging a View onto that object applies the selected object's context as the starting point for the new analysis. This lets the analyst open an appropriate View while retaining the relevant traffic selection, such as the time interval or the metric represented by the object.

Right-clicking in the View chart and choosing *Drill-down* is also a supported interactive workflow. The Drill-down command presents the available View choices for the selected chart context, enabling the analyst to choose the next visualization or analysis view without manually rebuilding the selection. Together, these actions provide the normal Cascade Pilot workflow for progressively narrowing an investigation within an opened trace file: select the visual evidence of interest, then open a context-aware drill-down View. Riverbed documentation resources and product materials are available through the Riverbed Documentation portal and the Riverbed Support portal.

QUESTION NO: 16

Deployment of Virtual NetProfiler, Virtual NetExpress, and Virtual Gateway is supported on which version of ESXi? (Choose three.)

A. 5.5

B. 5.0

C. 5.1

D. 4.0

ANSWER: A B C

Explanation:

Virtual NetProfiler, Virtual NetExpress, and Virtual Gateway support VMware vSphere ESXi 5.0, 5.1, and 5.5 for the product release context reflected by this question. These virtual Riverbed appliances are supplied as virtual-machine deployments and require a supported ESXi hypervisor version to provide the virtual CPU, memory, disk, and virtual-network resources needed for packet and flow processing. ESXi 5.0 established the baseline supported VMware platform for these deployments, while ESXi 5.1 and ESXi 5.5 continued that supported vSphere compatibility. Accordingly, an administrator can deploy each of the named virtual appliances on any of those three ESXi releases, subject to meeting the appliance's sizing, licensing, and virtual-network-interface requirements. Riverbed's product documentation and software resources are

available through the Riverbed NetProfiler support page. VMware also retains product documentation for the relevant platform family in its vSphere 5.5 installation documentation, which identifies ESXi as the hypervisor on which virtual appliances are deployed.

QUESTION NO: 17

A list of live interfaces within the Cascade Pilot interface:

A. Appear listed under the devices tab of the sources panel.

B. Appear in the main workspace after selecting “choose an interface” from the main menu.

C. Appear listed under the associated Cascade Shark capture job trace in the sources panel files tab.

D. Must be selected from the “Capture. Interfaces” menu.

ANSWER: A

Explanation:

Appear listed under the devices tab of the sources panel. In Cascade Pilot, the Sources panel is the navigation area used to select the data source for analysis. Its Devices tab organizes available appliances and their live data sources, including the interfaces that can be opened for real-time monitoring and analysis. Selecting a listed live interface establishes the context for the workspace and lets the analyst view traffic and performance information from that interface. This design distinguishes live device/interface sources from stored artifacts such as packet captures, trace files, or other files, which are organized separately in the Sources panel. The wording is specifically about a *list* of live interfaces, so the relevant location is the Devices tab rather than an action in a workspace or capture-selection menu. Riverbed’s product documentation resources are available through the Riverbed Documentation portal, and Riverbed identifies Cascade as part of its network-performance-management product family in its Network Performance Management overview.

QUESTION NO: 18

What are the features available in Cascade Pilot that are NOT available in Cascade Pilot Personal Edition (PE)? (Select 3)

A. Drill down from Cascade Profiler

B. Indexing files for fast analysis

C. Visual troubleshooting and analysis

D. Distributed capture and analysis (Cascade Shark appliance)

E. Application Transaction diagrams

ANSWER: A D E

Explanation:

Cascade Pilot’s full edition was designed to operate as part of the broader Cascade performance-management workflow, whereas Pilot Personal Edition was a standalone, limited packet-trace analysis tool. **Drill down from Cascade Profiler** is therefore a full-product capability: it enables an analyst to move from Profiler’s network and application visibility into focused packet-level investigation in Pilot. **Distributed capture and analysis (Cascade Shark appliance)** is also a full-edition feature because it relies on integration with a Cascade Shark appliance to obtain and analyze packet captures collected remotely across the network. **Application Transaction diagrams** are an advanced visual-analysis capability of the complete Cascade Pilot product, allowing packet activity to be interpreted as application transactions and exchanges rather than only as individual frames. These capabilities depend on the integrated Cascade ecosystem and the full Pilot feature set, beyond the local, personal-use analysis scope of PE. Riverbed’s packet-analysis product information describes the platform’s packet-level troubleshooting and integration-oriented workflow: Riverbed Packet Analyzer. Riverbed’s support portal also provides product documentation and historical software resources for Riverbed performance-management products: Riverbed Support.

QUESTION NO: 19

Which of the following statements are true about NetProfiler Security Profiles? (Choose three.)

A. A security profile can be created for “business hours” or “weekends” or any other time periods you want to specify.

B. Security Profiles utilize the SNORT Threat Signature Database.

C. Security Profile data is aged out after one month.

D. Security Profiles negate the need for IDS/IPS type systems.

E. The appliance collects traffic data from the monitored network and aggregate it into security profiles.

F. Recent statistics play a larger role in the profile than older statistics.

ANSWER: A E F

Explanation:

NetProfiler Security Profiles establish a behavioral baseline for traffic observed on the monitored network. Consequently, “The appliance collects traffic data from the monitored network and aggregate it into security profiles” describes the fundamental purpose of the feature: it derives profile data from observed flow records rather than requiring manually supplied traffic baselines. Profiles can also reflect expected differences in network behavior over defined periods. Therefore, “A security profile can be created for “business hours” or “weekends” or any other time periods you want to specify” is correct, because organizations can define time-based profile periods that match their operational schedules and recurring traffic patterns.

The baseline is adaptive rather than an equal-weighted historical snapshot. “Recent statistics play a larger role in the profile than older statistics” is correct because the profiling mechanism weights newer observations more heavily, enabling the expected behavior model to adjust as applications, users, and network usage naturally change. This makes security analysis more useful for identifying departures from current normal behavior while retaining historical context. Riverbed describes NetProfiler as providing flow-based network visibility and analytics; its product information is available at Riverbed NetProfiler. Riverbed documentation resources, including product documentation and release-specific material, are available through the Riverbed Documentation Portal.

QUESTION NO: 20

What is a benefit that Cascade might provide to Network Engineering? (Select 3)

A. Ability to view link usage over the last few months for capacity planning purposes.

B. Ability to consolidate and view all Steelhead and Router configurations across the network.

C. Ability to view applications and protocols used across the enterprise.

D. Ability to view WAN use by top hosts, top servers, top application, top ports, top protocols, and QOS markings.

ANSWER: A C D

Explanation:

Cascade provides network engineering teams with visibility into traffic behavior and long-term utilization, supporting both operational analysis and planning. The ability to view link usage over the last few months for capacity planning purposes is a core benefit because historical traffic and utilization trends let engineers identify sustained growth, validate WAN upgrades, and make capacity decisions based on measured demand rather than isolated snapshots. Cascade also identifies applications and protocols used across the enterprise, enabling engineers to understand what is consuming network resources and to distinguish business-critical traffic from other traffic categories. Its flow-based analytics further support detailed WAN analysis by exposing top hosts, servers, applications, ports, protocols, and QoS markings. This level of drill-down helps network teams locate the endpoints and traffic classes behind utilization, validate QoS treatment, and prioritize remediation or policy changes. These capabilities align with Riverbed network performance management’s emphasis on comprehensive traffic visibility, application-aware analysis, and historical reporting for troubleshooting and capacity

management. See Riverbed Network Performance Management and the Riverbed Documentation Portal for Riverbed product and documentation resources.

QUESTION NO: 21

To organize hosts into logical arrangements what is important to configure on Cascade Profiler?

A. Host Baseline Profiles

B. Custom Port Definitions

C. Host Groupings

D. Mitigation Configuration

ANSWER: C

Explanation:

Host Groupings is correct because Cascade Profiler uses host groups to place individual hosts into meaningful logical arrangements, such as business units, geographic locations, server tiers, application environments, or other organizational categories. This configuration makes it possible to view, filter, compare, and report network and application performance data in a way that reflects how the organization actually operates, rather than requiring an administrator to work only with a flat list of individual IP-addressed hosts. Logical host organization also helps analysts quickly scope investigations and understand which systems belong to a particular service or operational domain. In Riverbed's network performance monitoring approach, contextual organization of monitored entities is essential for turning collected flow and performance information into usable operational visibility. Host Groupings therefore directly supplies the configuration requested: a mechanism for arranging hosts logically within the Profiler interface and its analysis workflows. Riverbed describes its network performance monitoring capabilities and the operational visibility they provide at Riverbed Network Performance Management. Additional Riverbed product and support resources are available through the Riverbed Support Portal.

QUESTION NO: 22

In the Cascade Profiler GUI, where can you map a name to the combination of IP Address and TCP Port Number?

A. Definitions->Host Groups

B. Definitions->Applications

C. Definitions->Port Groups

D. Definitions->Port Names

E. Integration->Switch Port Discovery

ANSWER: B

Explanation:

Definitions->Applications is correct because Cascade Profiler application definitions associate a meaningful application name with traffic-identification criteria, including an IP address and TCP port combination. This lets the Profiler classify and present traffic using a business-relevant or service-relevant application label rather than showing only raw endpoint and port data. For example, an administrator can define a specific service hosted at a particular address and TCP port so that reporting, analysis, and drill-down views consistently identify that traffic as the intended application. These definitions are particularly useful where a port number alone is not sufficiently descriptive, such as when multiple services use common ports or when an organization needs to distinguish a particular server-hosted service. The Applications definition area is therefore the appropriate GUI location for creating and maintaining this address-and-port-based mapping. Riverbed's Network Performance Management material describes the role of application-aware visibility and analysis in identifying and reporting networked application traffic: Riverbed Network Performance Management. Riverbed product documentation and support resources are available through the Riverbed Documentation portal.

QUESTION NO: 23

What is one of the benefits Cascade provides compared to other vendors that have both flow and probe technology?

A. Cascade will de-duplicate/Coalesce data from all sources, both flow and sensor data. Other vendors have you look at each source's data separately

B. Cascade Profiler will provide nanosecond granularity from flow data to detect micro-bursts.

C. Cascade can also accept SNMP traps from other devices.

D. Cascade integrates with sflow but not NetFlow.

ANSWER: A

Explanation:

Cascade will de-duplicate/Coalesce data from all sources, both flow and sensor data. Other vendors have you look at each source's data separately is correct because the Cascade platform's value was its unified analysis model. Cascade Profiler correlated and coalesced records obtained from flow exporters with traffic and transaction information supplied by Cascade Sensor appliances. Rather than treating each collector as an isolated reporting silo, it presents related network activity in a common view, allowing an analyst to investigate the same hosts, applications, conversations, and time periods without manually reconciling duplicate observations from separate data sources. This consolidation improves operational efficiency and helps preserve context as an investigation moves from broad flow-based visibility to detailed packet-derived evidence. Riverbed describes its network-performance-management approach as combining flow-based visibility with packet-based analysis, enabling teams to identify and diagnose performance issues from a unified operational perspective. The Cascade product family is part of Riverbed's historical network performance management portfolio; current Riverbed network observability information is available at Riverbed Aternity, while Riverbed's support portal provides product documentation and knowledge resources at Riverbed Support. The key benefit in the stated answer is therefore the coalescing of flow and sensor information into one analytical data set and workflow.

QUESTION NO: 24

What version of Windows is/are supported by the Cascade Active Directory Connector? (Select 2)

A. NT 4.0

B. Windows Server 2003

C. Windows Server 2008

D. Windows Server 2005

E. A & D

ANSWER: B C

Explanation:

The supported Windows platforms are **Windows Server 2003** and **Windows Server 2008**. The Cascade Active Directory Connector was designed to query Active Directory and import directory identity information for use by the Cascade/NetProfiler platform. Its supported operating-system matrix includes these server releases, which provide the required Active Directory and Windows management interfaces for the connector's operation. In this context, the connector should be installed on a supported Windows Server system that can communicate with the domain controllers and with the Cascade appliance or NetProfiler deployment. This is a legacy Riverbed product component, so its compatibility should be interpreted according to the documentation and support matrix applicable to the deployed Cascade software release; later Windows versions may require a newer product version or an updated connector. Riverbed's current documentation resources are available through the Riverbed Documentation Portal, while archived product documentation and software lifecycle materials can be accessed through the Riverbed Support Portal.

QUESTION NO: 25

How many Cascade Shark appliances can be accessed by a single Cascade Pilot license?

A. 2

B. 10

C. 100

D. No limit

ANSWER: B

Explanation:

10 is correct. A Cascade Pilot license provides access to up to ten Cascade Shark appliances. Cascade Pilot is the centralized analysis and workflow interface used with Cascade Shark packet-capture appliances: it enables an analyst to select an available Shark, retrieve or inspect packet data, and use Pilot's analysis tools without working independently at each capture appliance. The ten-Shark entitlement is a licensing limit, not merely a practical recommendation for simultaneous use or an appliance-discovery limit. Therefore, when sizing a legacy Cascade deployment, the number of Shark appliances to be accessed through Pilot must be matched to the Pilot licensing capacity. If an environment needs centralized access beyond ten Sharks, it requires additional licensing or an appropriately designed management approach. Riverbed's legacy SteelCentral Network Performance Management support materials, including the Cascade Pilot and Cascade Shark documentation areas, provide the product documentation and release information relevant to these licensing and deployment considerations. See the Riverbed Cascade Pilot support page and the Riverbed Cascade Shark support page.

QUESTION NO: 26

What is a benefit that Cascade provides that is not supported by most other network performance management solutions?

A. Identity integration to identify which users are logged into which client machines.

B. Integration with synthetic transaction tools to obtain response time information.

C. A fully functional CLI that is easier to use than a Web based GUI.

D. A Web based GUI that does not require a password or HTTPS connection allowing for easier connection.

ANSWER: A

Explanation:

Identity integration to identify which users are logged into which client machines is correct because Cascade's network-performance analysis can associate observed network activity with the person using an endpoint, rather than limiting analysis to IP addresses, hosts, ports, and applications. This identity-aware context is particularly useful when an operations team needs to move from a network symptom—such as slow response time or excessive traffic—to the affected employee or group. It supports more meaningful investigation, reporting, and prioritization because dynamic addressing, shared systems, and changing client locations can otherwise make an IP address alone an unreliable indicator of the responsible or affected user.

Cascade was designed to correlate network and application performance data with business-relevant context. Identity integration supplies that endpoint-user context by obtaining user-to-machine associations from enterprise identity sources. This enables analysts to investigate activity by user identity and relate performance observations to the actual users experiencing an issue. Riverbed's performance-management portfolio and its supporting documentation describe the emphasis on visibility, correlation, and operational troubleshooting across network and application environments. See the Riverbed Network Observability product information and the Riverbed Support portal for product and documentation resources.

QUESTION NO: 27

Packet de-duplication in Cascade Shark appliance applies to:

A. Packets across all the ports of an appliance.

B. Packets on a single physical capture port.

C. Packets on an aggregating capture port.

D. Packets across different Cascade Shark appliances.

ANSWER: B

Explanation:

Packets on a single physical capture port. is correct. In a Cascade Shark appliance, packet de-duplication is performed within the traffic stream received by an individual physical capture interface. This scope is important because the appliance evaluates packet characteristics and timing in the context of that one capture-port input, suppressing duplicate observations before they affect packet storage and analysis. The result is a cleaner packet trace and more accurate downstream analysis when duplicated traffic is present in the monitored feed.

This behavior reflects the appliance's capture-oriented architecture: capture ports are independently defined sources of packet data, and de-duplication is associated with the packets arriving at a particular physical source rather than being a chassis-wide or multi-appliance correlation function. When designing a capture deployment, engineers should therefore consider how monitored links are delivered to each physical capture port and ensure that duplication scenarios are addressed at the appropriate capture input.

Riverbed's product and documentation resources for its network performance monitoring and packet-analysis platforms are available through the Riverbed Network Performance Management page and the Riverbed Support portal.

QUESTION NO: 28

There are the major characteristics of groups in SteelCentral Portal.

A. Provide structure in the application model

B. Manually created

C. Can be expanded or collapsed

D. All of the above

ANSWER: D

Explanation:

All of the above is correct because SteelCentral Portal groups combine each of the stated characteristics. Groups provide a logical hierarchy and structure within the application model, enabling related applications, services, and associated components to be organized in a way that reflects the monitored environment. This structure helps operators navigate monitoring data and understand relationships at appropriate levels of the model.

Groups are manually created by Portal users or administrators rather than being automatically inferred as a fixed part of the application model. This gives teams the flexibility to organize monitored entities according to business services, departments, locations, ownership, or another operational convention. The group hierarchy is also designed for practical navigation: users can expand a group to view contained items and collapse it to reduce visual clutter and focus on the relevant portion of the model.

Together, manual creation, hierarchical structure, and expandable/collapsible presentation describe the intended use of groups in SteelCentral Portal. Riverbed's SteelCentral Portal product documentation and software-support resources are available through the Riverbed SteelCentral Portal support page; additional product context is provided on Riverbed's SteelCentral product page.