PassQueen
a.k.a dumpsqueen.com

Riverbed Certified Solutions Professional - Network Performance Management

Riverbed 299-01

Version Demo

Total Demo Questions: 27

Total Premium Questions: 278

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Network Performance Management Solution Architecture	8
Topic 2, Network Performance Management Solution Deployment	8
Topic 3, AppResponse	81
Topic 4, NetProfiler	173
Topic 5, Portal	6
Topic 6, Mix Questions	2
Total	278

QUESTION NO: 1

Within Cascade Profiler you can set the recipient for policy notifications under Advanced configuration to:

A. Log Only, SNMP Only, Email Only, SNMP and Email

B. Log Only, Default, Owner, or any defined recipient

C. Log Only, Default, Owner or any User Account

D. Log Only, Default, Owner or any User Name obtained via the Identity Integration

ANSWER: B

Explanation:

In Cascade Profiler, the policy-notification recipient setting is designed to use the product's configured notification-recipient model rather than requiring a notification transport choice at each policy. The available recipient selections are *Log Only*, *Default*, *Owner*, or a specifically defined recipient. *Log Only* records the notification in the system without addressing it to a recipient. *Default* uses the globally configured default recipient, while *Owner* directs the notification according to the policy or object ownership context. Selecting a defined recipient allows the administrator to target a recipient that has already been configured in Cascade Profiler. This approach lets notification behavior be administered centrally and then reused consistently across policies. Riverbed's documentation resources for legacy Cascade/SteelCentral products are available through the Riverbed Support portal and the Riverbed Documentation portal. Therefore, "Log Only, Default, Owner, or any defined recipient" correctly identifies the Advanced configuration recipient choices.

QUESTION NO: 2

Hosts that do not match group definitions within a group type are:

A. Placed in an unassigned group.

B. Not placed into a group within the group type.

C. Not allowed

D. Flagged as an error.

E. Likely to cause errors in the future.

ANSWER: B

Explanation:

Not placed into a group within the group type. is correct. In Riverbed NetProfiler, a group type is a classification framework containing one or more group definitions that identify hosts according to configured matching criteria, such as addresses, host names, or other available host attributes. Membership is evaluated independently for each group type. A host is included in a group only when it satisfies that group's definition; if it does not satisfy any definition in that particular type, it simply has no membership in that group type. This does not prevent the host from being monitored, recorded, or classified under other group types. It means that reports, filters, and rollups scoped to that group type will not associate the unmatched host with one of its defined groups. Administrators can address this by creating or adjusting a group definition when they want those hosts represented within the type. This behavior supports precise, administrator-controlled categorization rather than assigning unmatched hosts to an implied catch-all classification. Riverbed NetProfiler provides the host and traffic visibility on which these logical groupings and reporting views are based. See the Riverbed NetProfiler product overview and Riverbed's product documentation portal for product documentation and configuration guidance.

QUESTION NO: 3

What integrations are available when an enterprise already owns an SNMP monitoring technology (such as Solarwinds or HPOV)? (Select 2)

A. Cascade Profiler can use switch port discovery to retrieve ARP and CAM tables directly from the NMS.

B. Most NMS' can be configured to SNMP Poll Cascade Profiler to retrieve conversation information.

C. Network API (from NMS via URL strings to get Cascade Reports)

D. There are no integrations; NMS' are competitors to Cascade Profiler.

E. External Links (from Cascade Profiler via URL strings to NMS)

ANSWER: C E

Explanation:

Network API (from NMS via URL strings to get Cascade Reports) is a supported integration approach because Cascade Profiler exposes report content through URL-based requests. An existing network-management system can use those URLs to open or retrieve the relevant Cascade report views in the context of an alert, device, application, or troubleshooting workflow. This lets operations teams retain their established NMS as a primary event and infrastructure-monitoring console while adding flow-based performance and conversation analysis from Cascade Profiler.

External Links (from Cascade Profiler via URL strings to NMS) is also supported. External-link configuration provides contextual navigation in the opposite direction: an administrator can build URL templates in Cascade Profiler that open the corresponding managed object, graph, event, or other NMS view. Values known to Cascade Profiler, such as an address or interface-related context, can be passed in the generated URL. Together, these two URL integrations support practical bidirectional workflow navigation rather than requiring replacement of the enterprise's existing SNMP-management platform. Riverbed documentation resources are available through the Riverbed Documentation Portal and the Riverbed Support Portal.

QUESTION NO: 4

How many Group Types Can you create on your Cascade Profiler?

A. 5

B. 20

C. 15

D. 10

E. Unlimited

ANSWER: D

Explanation:

10 is correct. In Cascade Profiler, a group type is a top-level organizational category used to classify related groups for reporting, filtering, and analysis. For example, an administrator can define a group type for a particular operational classification and then create the individual groups that belong within that classification. Cascade Profiler supports creation of up to ten such group types. This limit applies to the number of group-type categories, rather than to the number of individual hosts, applications, or traffic records that can be viewed in the system. Group types help make profiler views consistent across teams by providing a structured way to organize entities and apply group-based analysis. This is especially useful when reporting on traffic by business-relevant categories rather than relying solely on raw IP addresses or ports. Legacy Cascade Profiler documentation is available through Riverbed's customer support and documentation resources; access can require an entitled support account. See the Riverbed Support Portal and the Riverbed Documentation portal for product documentation and available legacy resources.

QUESTION NO: 5

When you create an index within Cascade Pilot for a trace filE. (Select 2)

A. The index will be removed when Cascade Pilot is closed.

B. The operation can take a few minutes.

C. You cannot apply Views that do not support indexing.

D. The trace file icon changes to a lightning bolt to indicate that it is indexed.

ANSWER: B D

Explanation:

“The operation can take a few minutes.” is correct because creating an index requires Cascade Pilot to read and process the trace data, building information that can be used to accelerate subsequent analysis. The duration depends on practical factors such as the trace’s size, packet count, and the available workstation resources. Index creation is therefore a processing task rather than an instantaneous change to the trace.

“The trace file icon changes to a lightning bolt to indicate that it is indexed.” is also correct. Cascade Pilot uses the lightning-bolt visual indicator in its trace-file interface to show that an index is available for that trace. This provides an immediate way for an analyst to identify indexed files before opening or working with them further. Indexing is intended to improve the responsiveness of supported analysis operations by making relevant trace information easier for the application to retrieve. Riverbed product documentation and software resources are available through the Riverbed Support Portal and the Riverbed Support site.

QUESTION NO: 6

How is the Cascade Pilot license activated?

A. The license key is entered during installation. The license is activated either through an Internet connection or by phone.

B. The license key is entered during installation. The license must be activated through an Internet connection.

C. The license key is entered during installation. There is no activation but Cascade Pilot will detect if the license is already in use on the network.

D. There is no license key activation. The software image is shipped with the license key installed.

ANSWER: A

Explanation:

The license key is entered during installation. The license is activated either through an Internet connection or by phone. Cascade Pilot uses a product-key licensing workflow in which the installer accepts the supplied license key and then completes activation with Riverbed. When the installation host has Internet access, the software can submit the activation request electronically. Riverbed also supports a phone-based activation path, which accommodates installations where direct Internet access is unavailable or restricted by network policy. This two-path process is important for network-monitoring deployments, because the appliance or host may be installed in a secured management network without outbound connectivity. The key identifies the purchased entitlement, while activation registers that entitlement for use on the installed system. Riverbed’s support and licensing resources remain the appropriate authoritative channels for entitlement, license-key, and activation assistance: Riverbed Support and Riverbed Documentation.

QUESTION NO: 7

In order for the Cascade solution to track an application based on its fingerprint, data can be used from which devices? (Select 3)

A. Packeteer Device

B. Cascade Gateway

C. Cascade Sensor

D. Cisco NBAR device

ANSWER: A C D

Explanation:

Packeteer Device, **Cascade Sensor**, and **Cisco NBAR device** are valid sources of application-fingerprint information for the Cascade solution. Cascade can use application-identification data supplied by these devices to associate observed traffic with known applications and report on that traffic using application-aware views. A Packeteer device contributes application-recognition information based on its traffic-classification capabilities. A Cascade Sensor supplies detailed packet-derived visibility and can identify applications from the traffic it observes. A Cisco NBAR device provides Cisco's network-based application-recognition classifications, which can be incorporated as application information for monitoring and reporting.

Using these data sources lets Cascade identify traffic by application characteristics rather than relying only on addresses, ports, or interfaces. This is particularly useful when applications use dynamic ports, encrypted transports, or shared infrastructure, because the reporting model can retain an application-oriented context. Riverbed's network-performance-management portfolio is designed to consolidate flow, packet, and application visibility from network instrumentation, while Cisco describes NBAR as a technology for recognizing and classifying network applications. See Riverbed Network Performance Management and Cisco Network-Based Application Recognition.

QUESTION NO: 8

Which of the following data is required for Cascade Profiler to correctly calculate RoundTrip-Time (RTT) between a branch and the datacenter for Steelhead optimized traffic?

(Select 2)

A. NetFlow from router on WAN side of branch Steelhead

B. NetFlow from router on LAN side of branch Steelhead

C. CascadeFlow from branch Steelhead

D. CascadeFlow from Data center Steelhead

E. Sensor at Data center

F. Sensor-VE sending Cascade flow data

ANSWER: C D

Explanation:

For Steelhead-optimized traffic, Cascade Profiler requires **CascadeFlow from branch Steelhead** and **CascadeFlow from Data center Steelhead** to calculate RTT correctly. Steelhead appliances export CascadeFlow records containing application-flow and WAN-optimization telemetry observed at each endpoint of the optimized connection. Receiving records from both the branch-side and data-center-side Steelhead gives Profiler the paired endpoint perspective it needs to correlate the same optimized flow and derive the round-trip timing across the WAN path. This is particularly important because the optimized WAN-side connection is not equivalent to the original client-server connection: Steelhead creates and manages the optimized transport relationship between peer appliances. Endpoint reporting from both peers therefore supplies the timing context associated with that relationship, allowing Profiler to represent RTT for the optimized traffic accurately. Riverbed documentation describes SteelHead as an appliance that optimizes traffic between peer SteelHead systems and provides product documentation through its documentation portal. See the Riverbed Documentation portal and Riverbed's SteelHead product overview for the architecture and telemetry context.

QUESTION NO: 9

Which of the following metrics are monitored in an Application Performance Policy? (Select

4)

A. Increase in Server Delay

B. Decrease in Average Connection Application-level Throughput

C. Increases in the number of TCP retransmissions

D. Decreases in the number of new connections to the application servers

E. Increase in the number of Active Connections

ANSWER: B C D E

Explanation:

An Application Performance Policy evaluates measurable changes in application traffic behavior that can indicate a developing service-quality or capacity issue. **Decrease in Average Connection Application-level Throughput** is a key performance indicator because it measures reduced application payload delivery per connection. **Increases in the number of TCP retransmissions** are monitored because retransmissions expose loss or impaired TCP delivery that can degrade user-visible application performance. Connection-volume measures are also important policy inputs: **Decreases in the number of new connections to the application servers** can reveal a reduction in demand or a failure affecting clients' ability to establish application sessions, while **Increase in the number of Active Connections** can identify an abnormal rise in concurrent application sessions and associated resource pressure. Together, these metrics let an administrator alert on throughput degradation, transport reliability symptoms, and meaningful changes in application connection behavior. Riverbed's documentation resources for NetProfiler and its policy-based network and application monitoring capabilities are available through the Riverbed Documentation Portal and the Riverbed NetProfiler product page.

QUESTION NO: 10

Which of the following is false in regards to Cascade Profiler reporting and alerting? (Select

2)

A. Any dependency (connection graph) can be exported to SVG (Visio) format.

B. Reports are real time and can not be scheduled.

C. Reports can be sent via PDF, HTML, CSV, or DOC format.

D. Events can be sent via email or SNMP v1 or v3 traps.

E. Events can be sent to both an email list AND SNMP target.

ANSWER: B C

Explanation:

The false statements are "Reports are real time and can not be scheduled." and "Reports can be sent via PDF, HTML, CSV, or DOC format." Cascade Profiler reporting is not limited to an operator viewing live data at the instant a report is requested. Its reporting workflow supports saved report definitions and scheduled delivery, enabling recurring operational, capacity, and troubleshooting reports to be generated without requiring a user to manually run each one. Scheduled report output can then be distributed using the reporting system's supported delivery mechanisms.

The listed set of report formats is also false because DOC is not a Cascade Profiler report output format. Profiler reporting supports standard web, document, and data-oriented outputs, but Microsoft Word DOC is not included as a supported report-delivery type. This distinction matters when integrating scheduled reports into an organization's reporting process: document requirements should be based on supported formats rather than assuming that an Office Word file can be generated directly. Riverbed's product documentation and support resources provide the authoritative version-specific guidance for report scheduling, output formats, and alert delivery configuration: Riverbed Documentation Portal and Riverbed Support Portal.

QUESTION NO: 11

The Discovery Wizard has been used to map a complex application to be monitored by Cascade's service monitoring for a company with 25 remote locations and one data center, each defined as a group within the group type used. There are two front-end segments to signify incoming end user connections, and 15 back-end connections among all of the servers. The only metric/policy activated for each segment is the 'Response Time' metric.

How many policies on the Cascade Profiler will this service consume?

A. 40

B. 65

C. 17

D. 55

E. 425

ANSWER: B

Explanation:

65 is correct. In Cascade Service Monitoring, a service-monitoring policy is consumed for each enabled metric on each applicable service segment instance. Because Response Time is the only enabled metric, the calculation is based on one policy per segment instance rather than multiple policies per segment. The two front-end segments are instantiated for each of the 25 remote-location groups, producing 50 Response Time policies: two segments multiplied by 25 locations. The 15 back-end server connections each contribute one additional Response Time policy, producing 15 more policies. The service therefore consumes 65 policies in total: 50 for the front-end monitoring instances plus 15 for the back-end connections. This sizing approach is important when planning Profiler service-monitoring capacity, since group expansion creates distinct monitored instances and enabled metrics determine the policy count for each instance. Riverbed's product documentation portal provides the documentation set for SteelCentral and legacy Cascade products at Riverbed Documentation. Riverbed's Network Performance Management product information is available at Riverbed Network Performance Management.

QUESTION NO: 12

The Cascade Profiler authenticates users before logging them in. What are the possible choices for authentication? (Select 3)

A. Cascade Profiler local database

B. User account specifies authentication by RADIUS

C. TACACS

D. RADIUS server to authenticate the user

ANSWER: A B D

Explanation:

Cascade Profiler supports authentication using its own local user database as well as an external RADIUS infrastructure. "Cascade Profiler local database" is therefore a supported choice when users are created and maintained directly in the Profiler. For external authentication, "User account specifies authentication by RADIUS" represents the account-level selection that directs authentication requests to RADIUS rather than the local credential store. "RADIUS server to authenticate the user" is also required as the external authentication source: the Profiler must be configured with the RADIUS server information so it can send the user's credentials for validation during login. Together, these selections describe the supported local and RADIUS-based authentication configuration paths available in Cascade Profiler. Riverbed documentation and product-support resources for the Cascade/SteelCentral NetProfiler product family are available through the Riverbed SteelCentral NetProfiler support page and the Riverbed documentation portal.

QUESTION NO: 13

Which of the following are considerations when configuring a SPAN session to feed a Cascade Shark appliance? (Select 3)

A. The SPAN destination port must have sufficient capacity for the mirrored traffic.

B. Packet de-duplication should be considered when the SPAN configuration can create duplicate packets.

C. The SPAN destination port should be reserved for the monitoring connection.

D. SPAN guarantees delivery of every mirrored packet regardless of destination-port bandwidth.

E. The Cascade Shark must be configured as the default gateway for the SPAN source VLAN.

ANSWER: A B C

Explanation:

The SPAN destination port must provide sufficient capacity for the mirrored traffic. When aggregate source traffic exceeds the destination-port capacity, the switch can drop mirrored packets even though production forwarding is unaffected. A VLAN SPAN configuration can produce more than one copy of a packet for some intra-VLAN traffic patterns, so packet de-duplication should be considered for the Cascade Shark capture interface. Also, a SPAN destination port is a monitoring output and should not be used as an ordinary production switch port. These practices help preserve capture fidelity and avoid misleading packet counts during analysis. Riverbed documentation is available through the Riverbed Documentation portal.

QUESTION NO: 14

What are two differences between NetFlow version 5 and NetFlow version 9 (select 2)

A. NetFlow version 5 generally support ingress flow export only; NetFlow version 9 supports both ingress and egress export.

B. NetFlow version 5 is used for Switches, NetFlow version 9 is used for Routers.

C. NetFlow version 9 includes information about CPU, Power-status and other router performance characteristics; NetFlow version 5 does not.

D. NetFlow version 9 includes the ability to export the Time-To-Live (TTL); NetFlow version 5 does not.

E. NetFlow version 9 includes the ability to export the packet latency, NetFlow version 5 does not.

ANSWER: A D

Explanation:

“NetFlow version 5 generally support ingress flow export only; NetFlow version 9 supports both ingress and egress export.” is correct in the context of Cisco NetFlow deployment. Version 5 is a fixed-record export format commonly associated with traditional ingress accounting, whereas version 9 is template-based and is used with the more flexible flow-export capabilities that can report flows observed on ingress or egress interfaces. This flexibility lets collectors distinguish the observation direction and use records appropriate to the configured export direction.

“NetFlow version 9 includes the ability to export the Time-To-Live (TTL); NetFlow version 5 does not.” is also correct. Version 5 has a fixed set of fields and does not contain an IP TTL value. Version 9 introduces templates that describe the fields in subsequently exported data records. Its standardized field definitions include the IP TTL information element, enabling an exporter to include TTL when that information is available and selected in its template. This template-based extensibility is a central difference between version 9 and the fixed layout of version 5. RFC 3954 defines NetFlow version 9’s template mechanism and field types, including IP TTL: RFC 3954. Cisco’s overview of NetFlow operation and export behavior is available at Cisco: NetFlow—A Comprehensive Overview.

QUESTION NO: 15

NetFlow v9 uses the concept of templates. How do you tell the Cascade Gateway the proper template information?

A. You configure the NetFlow v9 template on the Gateway for each NetFlow v9 exporting device.

B. You manually import the NetFlow v9 template for each device.

C. The template is automatically learned by the Gateway.

D. All of the above.

E. None of the above.

ANSWER: C

Explanation:

The template is automatically learned by the Gateway. NetFlow Version 9 is a template-based export protocol: before an exporter sends flow records that use a particular record format, it exports a Template FlowSet describing the fields and field lengths associated with that template ID. The Cascade Gateway, acting as the NetFlow collector, receives and caches this template information and then uses it to decode subsequent data FlowSets from that exporter. This dynamic exchange is fundamental to NetFlow v9 because different exporters can use different templates, including vendor-specific fields or different sets of standard information elements. The exporter also periodically retransmits templates so that a collector can recover template state after startup or expiration. Consequently, the Gateway does not require an administrator to manually import a template or create a per-device template definition in normal operation; it derives the required format from the templates sent in the NetFlow v9 export stream. This behavior follows the NetFlow v9 protocol specification's definition of Template FlowSets and collector template caching. See the NetFlow Version 9 RFC for the protocol's template and data FlowSet model, and Riverbed's SteelCentral product information for the platform context.

QUESTION NO: 16

Common methods for receiving alerts from a Network Performance Management system includE. (select 2)

A. Traps sent via NetFlow, sFlow, or IPFX

B. Alerts sent via eMail

C. Traps sent via SNMP

D. Traps sent via SPAN

E. Traps sent via FIX Protocol

ANSWER: B C

Explanation:

Alerts sent via eMail and Traps sent via SNMP are common notification mechanisms in network performance management deployments. Email notifications provide an accessible, asynchronous way to notify operations staff, distribution lists, or incident-management mail gateways when an alert condition is triggered. This allows alert policies to communicate important details such as the affected device, application, threshold, severity, and time of occurrence without requiring an operator to be actively viewing the management console.

SNMP traps provide a standard machine-to-machine notification method. An NPM platform can send an unsolicited trap to a configured SNMP manager or event-management platform when it detects a relevant condition. The receiving system can then correlate the event with other infrastructure alarms, create or enrich an incident, and apply its established escalation workflow. This interoperability is particularly useful in environments using centralized monitoring and event-management systems. Riverbed documents its monitoring and performance-management products through its product documentation portal; SNMP notification behavior is based on the standardized SNMP framework described in RFC 3413.

QUESTION NO: 17

Filtering in Cascade Pilot: (Select 2)

A. Is either a BPF or Wireshark Display filter only.

B. Can be created from a chart.

C. Can be changed after applying a View.

D. Is the equivalent of the Wireshark "display filter input".

ANSWER: B C

Explanation:

“Can be created from a chart.” is correct because Cascade Pilot supports an investigation workflow in which an analyst can select traffic represented in a visualization and use that selection to focus the current analysis. Charts therefore provide a practical starting point for creating a traffic filter based on observed hosts, applications, conversations, protocols, or other displayed dimensions. This lets the analyst move from a broad traffic overview to the packets or flows responsible for a charted condition without manually constructing every criterion.

“Can be changed after applying a View.” is also correct. A View supplies a useful initial investigation context, but it does not make the filtering criteria permanently fixed. Analysts can refine, replace, or otherwise adjust filtering after the View is applied as their investigation develops. This is important in packet and network-performance analysis, where an initial View often identifies a general traffic category and subsequent filtering narrows the scope to the relevant communications. Riverbed’s documentation resources are available through the Riverbed Documentation portal. For background on the display-oriented filtering concept used in packet analysis, see the Wireshark filter reference.

QUESTION NO: 18

When working with capture job traces on a Cascade Shark appliance from Cascade Pilot:

A. Views should always be applied directly to the job trace.

B. Trace clips should first be created, and views applied to them.

C. Packets should be exported from the capture job trace before beginning analysis.

D. The capture job trace should first be indexed by right-clicking on it and selecting "Index".

ANSWER: B

Explanation:

Trace clips should first be created, and views applied to them. A capture job trace can be a large, continuously collected packet-data source, whereas a trace clip isolates the precise time interval and traffic of interest for an investigation. Creating a clip gives the analyst a focused, manageable working data set before applying Cascade Pilot views and analysis operations. This supports an efficient workflow: identify the relevant period in the capture job, create a clip containing that evidence, and then use views to inspect conversations, protocols, performance behavior, or packet-level details within the selected data. The original capture job trace remains available as the broader source, allowing additional clips to be created later if the troubleshooting scope changes. This approach also avoids repeatedly applying potentially expensive analysis operations across an unnecessarily large capture-job data set. Riverbed’s packet-analysis tooling is designed around narrowing packet data to the traffic and interval relevant to the question being investigated; see Riverbed’s Packet Analyzer Plus product information and the Riverbed SteelCentral support and documentation portal for product documentation resources.

QUESTION NO: 19

How does Cascade track response time using its Sensor?

A. The Sensor watches TCP-three-way handshakes to measure Network RTT, Server Delay, and Response Time.

B. The Profiler will measure the delta between packets reported by each Sensor in the path.

C. The Sensor will measure HTTP page load times.

D. The Sensor will measure transaction times of over 100 applications using application signatures.

ANSWER: A

Explanation:

The Sensor passively analyzes TCP packet timing—including connection-establishment exchanges and the timing of subsequent request and response packets—to derive response-time components from observed network traffic. This lets Cascade distinguish the elapsed time attributable to network round-trip time from delay associated with the server and the application response. Because the calculation is based on packets seen by the Sensor, it does not require an agent on the application server or synthetic browser transactions. TCP's three-way handshake is particularly useful for establishing baseline network round-trip characteristics, while the observed timing of payload packets supplies the information needed to characterize server delay and overall response time for a TCP conversation. This passive, packet-based approach is consistent with Riverbed network-performance monitoring: Riverbed describes its observability portfolio as providing visibility into application and network performance from collected traffic and telemetry. See Riverbed Network Observability and Riverbed AppResponse. Thus, measuring Network RTT, Server Delay, and Response Time through TCP traffic observation is the appropriate description of how the Cascade Sensor tracks response time.

QUESTION NO: 20

Trace clips that no longer have data because their packet data and trend data is no longer available in the parent capture job trace:

A. Appear with a red lightning bolt icon within Cascade Pilot.

B. Appear with a lock icon within Cascade Pilot.

C. Appear with a grey lightning bolt icon within Cascade Pilot.

D. Appear with a yellow lightning bolt icon within Cascade Pilot.

E. Appear with the trace clip name in red text within Cascade Pilot.

ANSWER: E

Explanation:

Appear with the trace clip name in red text within Cascade Pilot. A trace clip is a saved subset or view derived from a parent capture-job trace, rather than a separate, independently retained packet capture. Its usability therefore depends on the continued availability of the associated packet and trend data in that parent trace. When the parent capture has aged out, been removed, or otherwise no longer retains the required data, Cascade Pilot identifies the affected clip as unavailable by displaying its name in red. This visual state lets an operator distinguish an unavailable historical clip from a normal saved clip while reviewing trace-clip entries in the application. The red name is specifically an availability indication: the clip definition can still be listed, but the underlying data needed to open or analyze it is no longer present. This behavior reflects the dependency of packet-analysis artifacts on retained capture data, a central concept in Riverbed packet-analysis workflows. Riverbed maintains product documentation and support resources through its documentation portal and provides product and lifecycle information through its product pages.

QUESTION NO: 21

For DNS reverse lookup, Cascade Profiler caches as follows:

A. Cache the most recent 500 IPs.

B. Obey DNS TTLs.

C. Cascade does not cache DNS responses.

D. For 24 hours.

ANSWER: D

Explanation:

For 24 hours. is correct. Cascade Profiler performs reverse DNS lookups to translate observed IP addresses into host names for display and analysis. To avoid repeatedly querying DNS servers for the same addresses, which would add lookup latency and unnecessary DNS-server load, the product retains reverse-lookup results in its DNS cache for a fixed 24-hour interval. During that interval, later views, reports, and drill-downs that need the same IP-to-name mapping can use the cached result rather than issuing another reverse query. This behavior provides consistent name presentation while limiting the operational overhead of resolving potentially large numbers of addresses seen in flow and packet-derived traffic data. Administrators should therefore account for the cache interval when troubleshooting recently changed PTR records or host-name mappings: a changed DNS reverse record may not appear in Cascade Profiler until the cached mapping expires or the relevant cache is refreshed. Riverbed product documentation and software documentation resources are available through the Riverbed Documentation portal; product-specific manuals and knowledge-base material can also be accessed through the Riverbed Support portal.

QUESTION NO: 22

A user wants to use the link congestion analytic policy on Cascade Profiler to monitor five key WAN links for an increase in traffic outside of the expected traffic that includes TCP/80, TCP/443, TCP25, and UDP/53. Can the link congestion policy be used for this scenario?

A. No, because a link analytic can monitor only a single interface.

B. Yes, one can define separate link congestion analytic policies for each WAN interface.

C. No, because one can only define specific ports to include on a link congestion policy.

D. Yes, because one can define a port group that includes every UDP and TCP port except for those that include the expected traffic, and that port group can be used for the link congestion analytic policy.

E. No, and both A and C include valid reasons.

ANSWER: D

Explanation:

Yes, because one can define a port group that includes every UDP and TCP port except for those that include the expected traffic, and that port group can be used for the link congestion analytic policy. Cascade Profiler link-congestion analytics can evaluate traffic matching a defined traffic criterion across the selected monitored links. In this case, the relevant criterion is not the expected web, mail, and DNS traffic itself; it is the complement of that expected traffic. A port group provides a reusable way to express that condition by covering TCP and UDP ports other than TCP/80, TCP/443, TCP/25, and UDP/53. Associating that group with the link-congestion analytic policy lets the policy focus on unexpected traffic that might be consuming WAN capacity. The policy can then be applied to the set of five important WAN links, enabling a consistent threshold and traffic definition for all of them rather than requiring a distinct analytic definition per interface. This use of port-based traffic classification is appropriate when the goal is to isolate traffic that falls outside a known baseline of allowed or anticipated services. Riverbed product documentation and support resources for legacy Cascade/SteelCentral NetProfiler deployments are available through the Riverbed Documentation portal and the Riverbed Support portal.

QUESTION NO: 23

When using Network Performance Monitoring solutions the following statements are truE.

(select 2)

A. HTTP Object load times can be analyzed with full packet capture

B. SSL Decryption can be used without any special key imports provided the start of the TCP conversation is captured

C. Software Agents are required on Servers in order to measure the Network Round trip time of applications

D. ICMP (ping echo/response) must be used in order to measure Network Round Trip time

E. NetFlow records can be summarized to provide utilization on a router interface

ANSWER: A E

Explanation:

HTTP Object load times can be analyzed with full packet capture is correct because packet-level monitoring retains the protocol exchanges needed to reconstruct an HTTP transaction and distinguish the timings associated with individual web objects. Analysis can use the observed request and response traffic to identify object retrieval behavior and transaction timing, making packet capture especially valuable for diagnosing web-application performance from the network perspective. Riverbed Network Performance Management products are designed to combine network, application, and packet-derived visibility for this type of troubleshooting.

NetFlow records can be summarized to provide utilization on a router interface is also correct. Flow exporters report traffic attributes such as byte counts, packet counts, timestamps, interfaces, and endpoints. A monitoring collector can aggregate those records by ingress or egress interface over a selected time interval, calculate the traffic rate, and present interface-level utilization trends. This enables capacity analysis and identification of the conversations contributing to an interface's load. See Riverbed's Network Performance Management overview and NetProfiler product information for Riverbed's flow and network-performance monitoring capabilities.

QUESTION NO: 24

A packet capture (pcap) file that is loaded into Cascade Pilot for the first time is indexed:

A. When the file is opened in Cascade Pilot for analysis.

B. When the file is imported into Cascade Pilot for analysis.

C. When the first view is applied to the file.

D. When "Add Trend Index" is selected from a right-click menu within Cascade Pilot.

E. When the file is copied onto a Cascade Shark appliance.

ANSWER: D

Explanation:

When "Add Trend Index" is selected from a right-click menu within Cascade Pilot is correct. Cascade Pilot can open and analyze packet-capture files without immediately creating a persistent trend index. A trend index is an explicitly generated data structure that summarizes traffic over time and supports efficient trend-oriented investigation of the capture. This is particularly useful for large or repeatedly used trace files, because it allows Pilot to retrieve time-based information and populate applicable trend views without repeatedly processing the entire packet stream from the beginning.

The indexing action is therefore initiated deliberately from the capture file's context menu by selecting *Add Trend Index*. Once the index is built, it is associated with the capture and can be used in subsequent analysis sessions, improving responsiveness for supported views and time-range navigation. This behavior distinguishes ordinary capture loading from the optional preparation of a trace for indexed trend analysis. Riverbed documentation and product resources are available through the Riverbed Documentation portal and the Riverbed Support portal, where legacy Cascade Pilot product documentation may be accessed by customers with the appropriate support entitlement.

QUESTION NO: 25

When editing a previously configured service policy, what options become available if you click the 'show advanced settings' checkbox? (Select 3)

A. Allows enabling/disabling the detection of dips in the metric.

B. Allows tuning of the tolerance range of the metric.

C. Allows setting of a noise floor for the metric.

D. Allows adjusting the notifications for the metric.

E. Allows enabling/disabling the detection of spikes in the metric.

ANSWER: A C E

Explanation:

When advanced settings are displayed for an existing service policy, the policy editor exposes controls that refine how metric anomalies are identified. **Allows enabling/disabling the detection of dips in the metric.** is correct because a policy can be configured to recognize meaningful downward deviations from the expected metric behavior. **Allows setting of a noise floor for the metric.** is also correct: the noise-floor setting prevents insignificant low-level variation from being treated as a meaningful condition, improving the practical signal quality of metric evaluation. **Allows enabling/disabling the detection of spikes in the metric.** is correct because the advanced policy controls separately support identifying upward deviations, allowing monitoring behavior to reflect the service characteristic being evaluated. Together, these settings let an administrator tailor anomaly detection to both positive and negative changes while avoiding alerts generated by normal background variation. This is consistent with Riverbed Network Performance Management's focus on using collected metric data to isolate material service-performance changes and support effective operational monitoring. Riverbed's Network Performance Management overview is available at Riverbed Network Performance Management, and Riverbed product documentation can be accessed through the Riverbed Documentation Portal.

QUESTION NO: 26

Layer 7 Fingerprints come from which of the following sources:

A. Cascade, User Defined, Packeteer, NBAR

B. Cascade, User Defined, Packeteer, NBAR, Steelhead appliance

C. Cascade, Packeteer, NBAR, Steelhead Appliance

D. Cascade, User Defined, Cisco, Netscreen, Steelhead appliance

ANSWER: A

Explanation:

Cascade, User Defined, Packeteer, NBAR is correct. In Riverbed Cascade/NetProfiler, Layer 7 fingerprints are application-identification definitions used to classify observed traffic beyond basic IP address, port, and protocol information. The fingerprint library can include Riverbed's built-in Cascade signatures, custom signatures created by an administrator for environment-specific applications, Packeteer-derived application signatures, and Cisco Network-Based Application Recognition (NBAR) definitions. Combining these sources gives NetProfiler a broad baseline for common enterprise applications while allowing organizations to extend identification for internally developed, specialized, or newly deployed services. This classification is important because application-level visibility drives meaningful traffic reporting, alerting, and troubleshooting: administrators can identify the application generating traffic and assess its use and performance rather than relying only on transport-layer characteristics. Riverbed's network-performance-management capabilities are designed around this kind of application-aware analysis, as described on the Riverbed Network Observability product page. Cisco documents NBAR as a deep-packet-inspection mechanism that recognizes applications and protocols, which is why NBAR signatures can serve as a Layer 7 fingerprint source; see the Cisco NBAR overview.

QUESTION NO: 27

When saving a new application performance analytic on Cascade Profiler, which of the following is true?

A. You must have at least three days of history of 15 minute rollups before saving the analytic.

B. You must have at least three days of history of 15 minute rollups before the analytic initializes.

C. If you include all possible metrics, but there is no possibility of measuring response time due to lack of Sensor coverage, the analytic will fail to initialize.

D. The analytic will not alert for a minimum of 1 week during baseline collection.

E. Both B and C.

ANSWER: B

Explanation:

You must have at least three days of history of 15 minute rollups before the analytic initializes. Saving the analytic and initializing it are separate stages. Cascade Profiler can save the analytic definition immediately, but the analytic needs a sufficient historical population of 15-minute rollup data before it can establish the expected behavior used for its analysis. The three-day requirement supplies enough time-series samples to create the initial model and begin evaluating application performance consistently. Consequently, an administrator should expect a newly saved analytic to remain in an initialization phase until the required historical rollups are available; this is not a restriction on saving the configuration itself. This behavior reflects the product's use of retained, aggregated flow and performance data to support application-level analysis rather than making a determination from an empty or very small dataset. Riverbed's Network Performance Management materials describe the platform's use of network and application performance data for analysis and troubleshooting; see Riverbed Network Performance Management and the Riverbed Support portal for product documentation and version-specific Cascade/SteelCentral resources.