Free Cisco 300-445 ENNA Practice Questions

Shared by Gill on 13-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Users on remote sites are reporting voice issues, can you identify possible causes and next steps from the following exhibits?

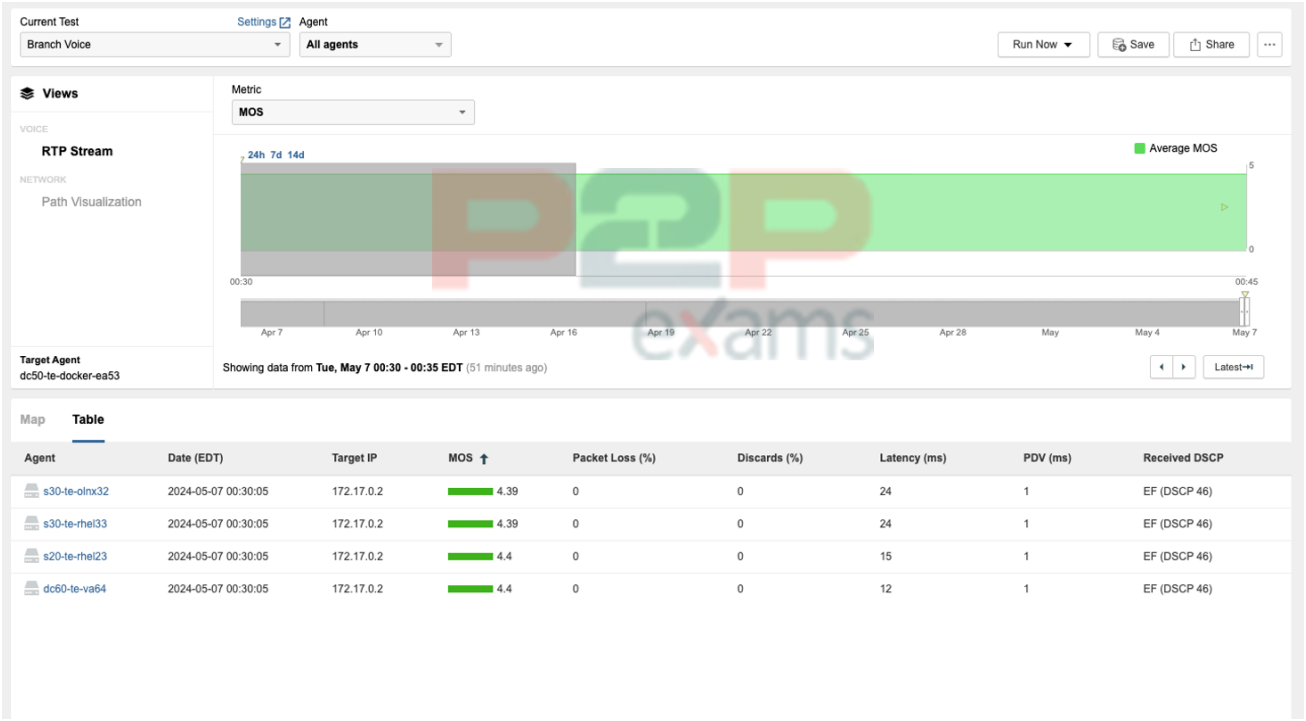

Agent	Date (EDT)	Target IP	MOS	Packet Loss (%)	Discards (%)	Latency (ms)	PDV (ms)	Received DSCP
s30-te-olnx32	2024-05-07 00:30:05	172.17.0.2	4.39	0	0	24	1	EF (DSCP 46)
s30-te-rhel33	2024-05-07 00:30:05	172.17.0.2	4.39	0	0	24	1	EF (DSCP 46)
s20-te-rhel23	2024-05-07 00:30:05	172.17.0.2	4.4	0	0	15	1	EF (DSCP 46)
dc60-te-va64	2024-05-07 00:30:05	172.17.0.2	4.4	0	0	12	1	EF (DSCP 46)

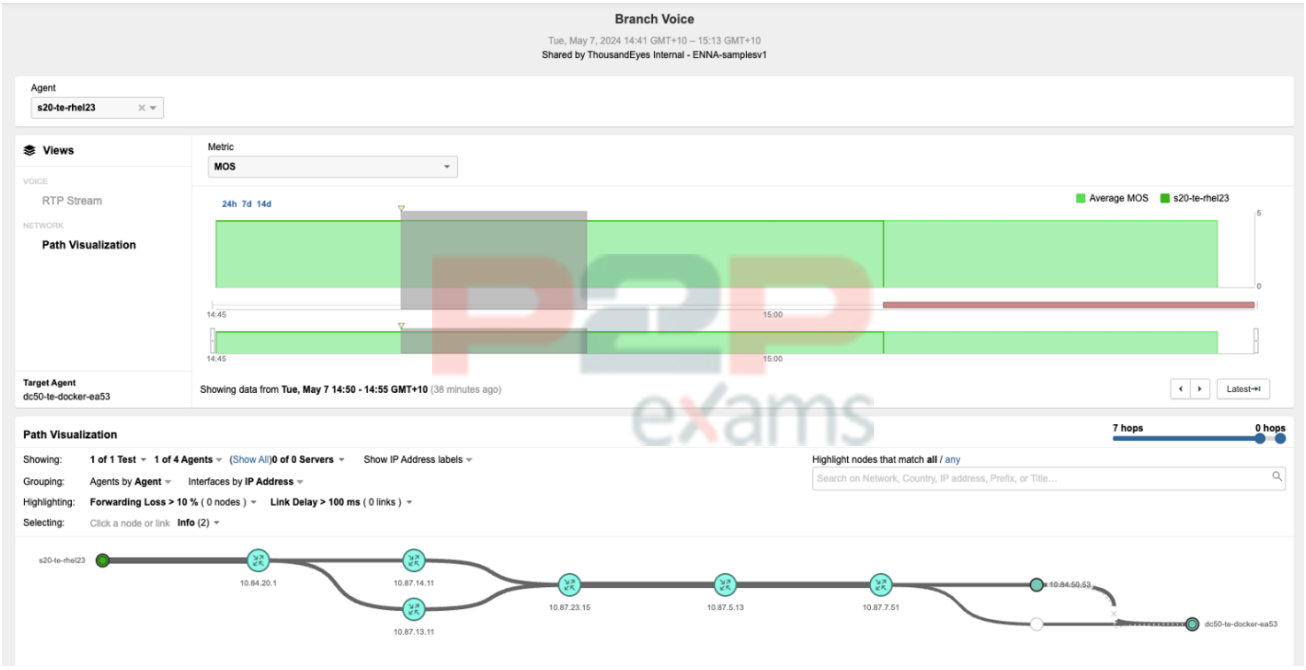

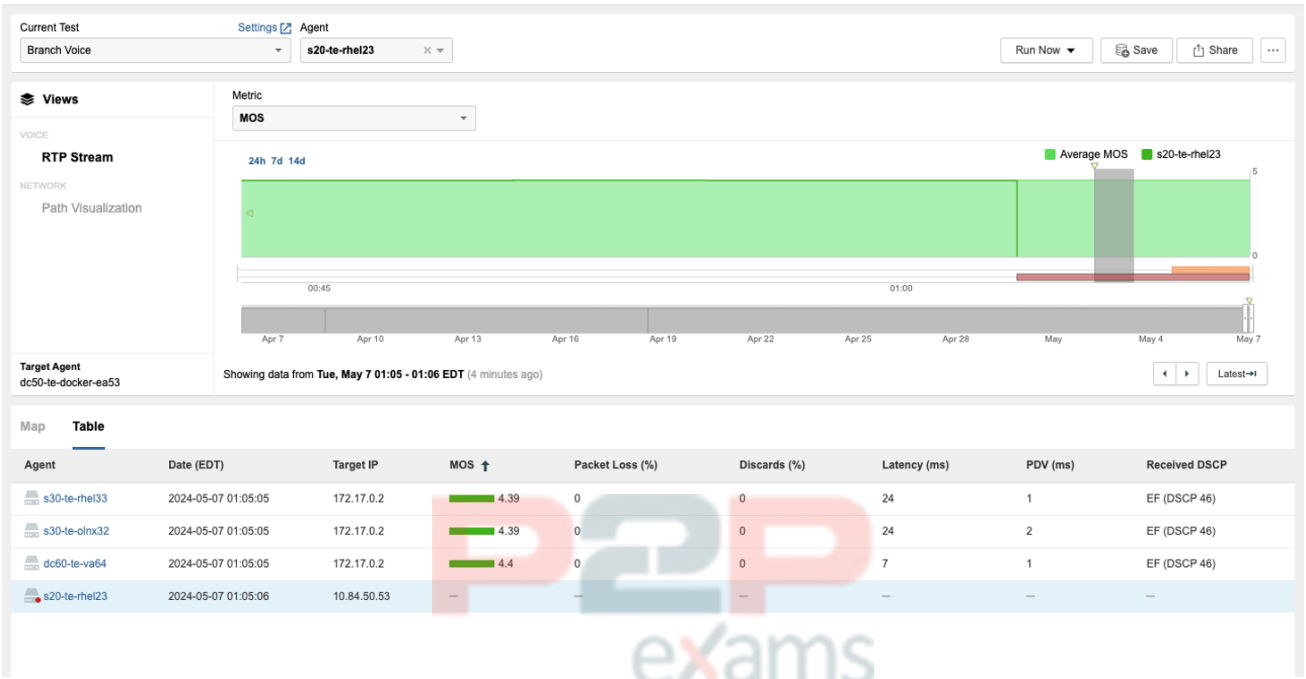

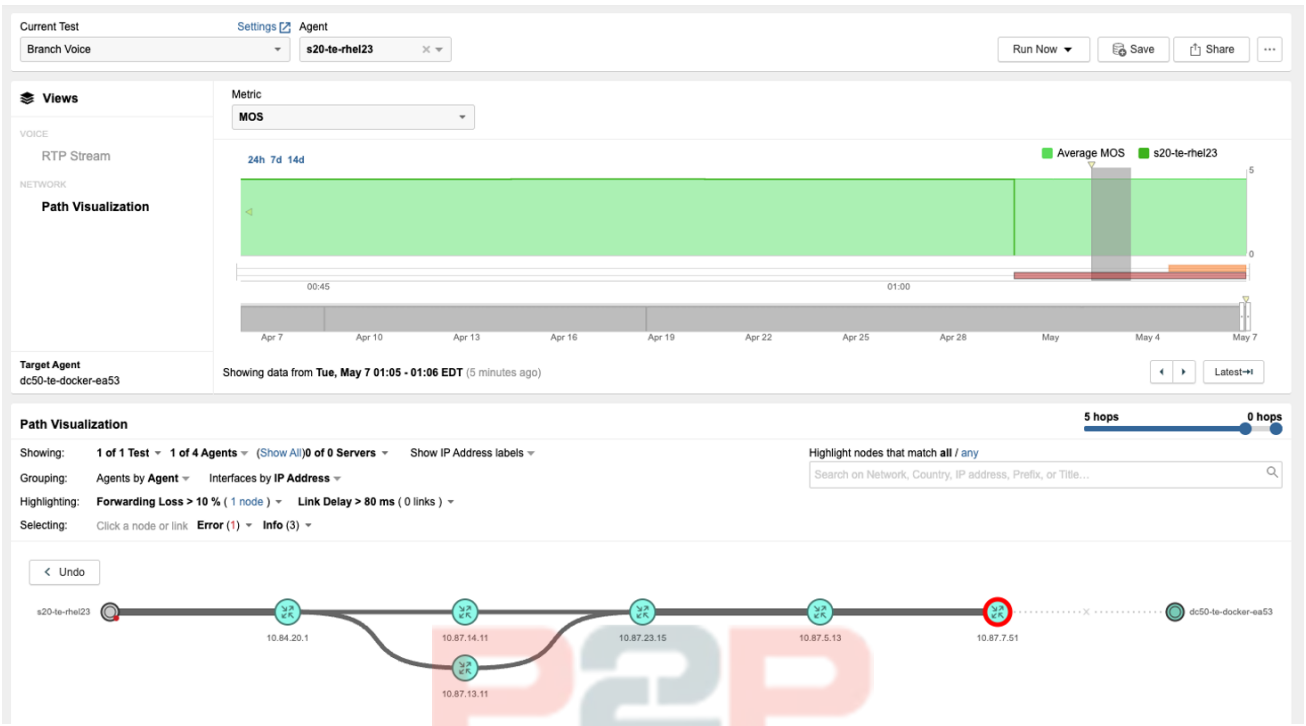

Options:

A- Involve the Voice team as the RTP test does not return any relevant results for the agent located at site 20 (identified as s20 in the exhibit)

B- Verify the routing changes on device 10.87.7.51

C- Verify the docker host 10.84.50.53 and ensure the agent container is running.

D- Analyze the jitter and latency trends on the affected voice paths to identify potential network congestion.

Answer:

C

Question 2

Question Type: MultipleChoice

Employees and customers of a retail company are experiencing performance issues with the store website, such as slowness during the login process or failure when adding items to the cart. Which test type is the most useful for identifying the root cause of these problems?

Options:

A- HTTP Server test type
B- Page Load test type
C- Transaction test type
D- Agent-to-server test type
E- DNS Server test type
F- Agent-to-agent test type

Answer:

C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, selecting the appropriate test type is essential for isolating performance bottlenecks in complex web applications. When users report issues with specific multi-step workflows---such as logging into a portal or interacting with a shopping cart---the Transaction test type (Option C) is the most effective tool.

A Transaction test is a specialized Web-layer test that utilizes a script (typically written in JavaScript/TypeScript) to mimic real user interactions with a website. Unlike a simple HTTP Server test (Option A) that only checks for a 200 OK response from a single URL, or a Page Load test (Option B) that measures the rendering of a single page, a Transaction test follows a predefined user journey. For this retail scenario, the test can be configured to navigate to the homepage, enter credentials, click the login button, search for a product, and add it to the cart.

The primary advantage of this approach is that it provides granular, step-specific timing. It allows the engineer to identify precisely where the latency or failure occurs---for example, if the backend database takes too long to process the login or if an API call fails during the 'add to cart' action. While Agent-to-server (Option D) and DNS Server (Option E) tests provide valuable network and infrastructure data, they cannot validate the functional logic of a web application. Agent-to-agent (Option F) is used for measuring throughput between two managed points and is irrelevant for public-facing website testing. Therefore, for troubleshooting interactive web

processes, the Transaction test type is the definitive choice for pinpointing the source of the issue.

Question 3

Question Type: MultipleChoice

Refer to the exhibit.

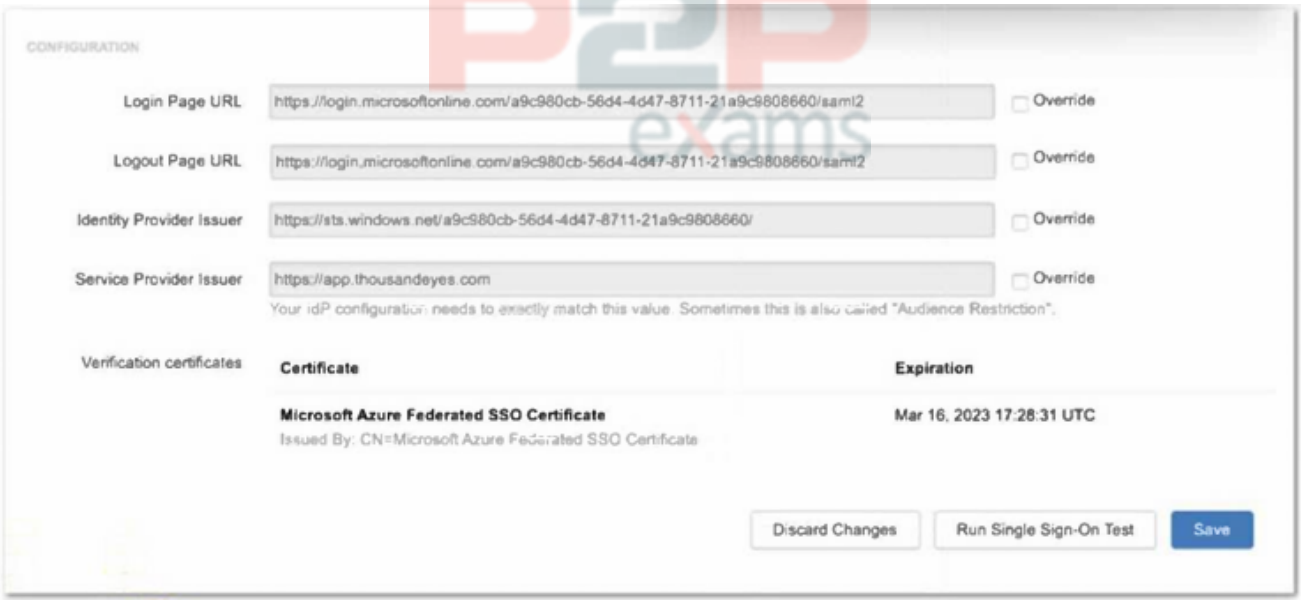

An engineer must configure Cisco ThousandEyes SSO to use Microsoft Entra ID using the configuration shown in the exhibit. Which feature must be set to override to complete the configuration?

Options:

A- Service Provider Issuer
B- Logout Page URL
C- Login Page URL
D- Identity Provider Issuer

Answer:

D

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA)

architecture, secure administrative access via Single Sign-On (SSO) is a critical component of platform governance. The exhibit illustrates the ThousandEyes SSO configuration panel being integrated with Microsoft Entra ID (formerly Azure AD). When configuring SAML-based authentication, the 'Identity Provider Issuer' (Option D) is a unique identifier provided by the IdP (Microsoft) that must match exactly between the two systems.

According to ENNA implementation guidelines, ThousandEyes populates default fields based on standard SAML metadata. However, Microsoft Entra ID often utilizes a specific GUID-based format for the Issuer URL (e.g., https://sts.windows.net/tenant-id/) that may differ from the generic URL format expected by the platform's initial auto-fill. To ensure a successful SAML handshake, the engineer must select the 'Override' checkbox next to the Identity Provider Issuer field. This action unlocks the field, allowing the engineer to manually paste the exact string provided in the Entra ID Federation Metadata document. If this value is not overridden and matched precisely, the SAML assertion will be rejected, resulting in a failed authentication attempt.

While the Login and Logout URLs (Options B and C) are also critical, they are typically correctly identified during the initial setup or metadata import; the Identity Provider Issuer is the most frequent point of mismatch requiring an manual override in Entra ID environments due to its strict 'Audience Restriction' requirements. The Service Provider Issuer (Option A) is generally a fixed value (https://app.thousandeyes.com) that rarely requires overriding as it defines ThousandEyes' own identity to the IdP.

Therefore, selecting the override for the Identity Provider Issuer is the necessary step to complete the integration and allow enterprise users to authenticate securely using their corporate credentials.

Question 4

Question Type: MultipleChoice

Which of the following metrics can be used to configure an alert rule for Endpoint Agent HTTP Server tests? (Choose two)

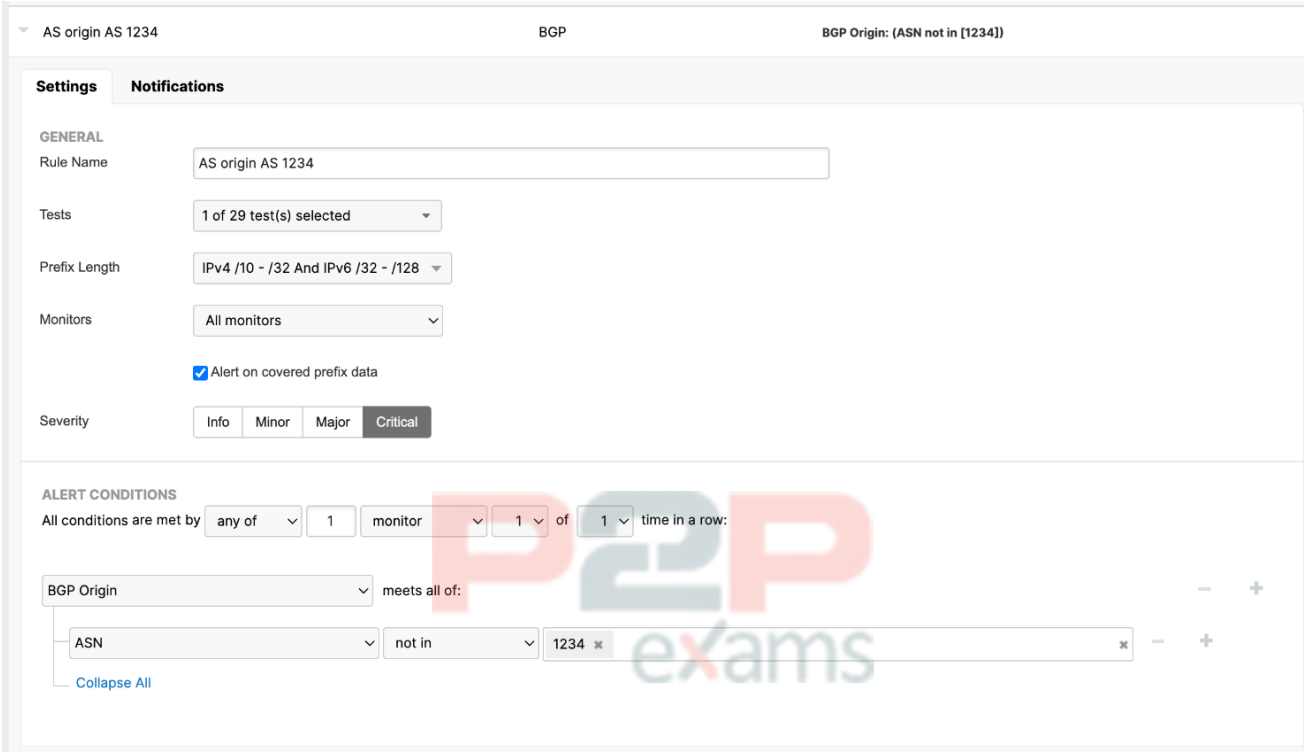

Options:

A- Response Time
B- BGP Reachability
C- Error Type
D- Interface Throughput

Answer:

A, C

Explanation:

In the ThousandEyes platform, alert rules are the primary mechanism for notifying operations teams when performance thresholds are breached. When configuring alerts for Endpoint Agent HTTP Server tests, the available metrics are limited to those that the Endpoint Agent captures during its synthetic browser-based or network-layer probes.

The two correct metrics for this test type are Response Time (Option A) and Error Type (Option C).

Response Time: This metric measures the time-to-first-byte. It is essential for detecting degradations in web application performance as seen from the user's specific machine.

Error Type: This allows administrators to trigger alerts based on specific HTTP response codes or protocol errors (e.g., DNS failures or SSL handshake issues).

Other options are incorrect because they apply to different agent types or monitoring layers:

BGP Reachability (Option B): BGP metrics are not applicable to Endpoint Agent tests. These metrics require the BGP route visualization layer, which is typically associated with Cloud or Enterprise Agents monitoring public internet paths.

Interface Throughput (Option D): This is an infrastructure-level metric rather than a synthetic application metric. While ThousandEyes can monitor device health via SNMP, it is not a metric used in a synthetic HTTP Server test.

By utilizing Response Time and Error Type, an architect can ensure that the operations team is alerted as soon as users experience slow loading or actual failures when accessing critical internal or SaaS-based web applications.

Question 5

Question Type: MultipleChoice

ThousandEyes offers several native integrations for receiving instant event notifications triggered by alerts. Which of the following integrations are available directly within the ThousandEyes platform? Select all that apply.

Options:

A- ServiceNow
B- PagerDuty
C- MS Teams
D- Splunk
E- AWS
F- AppDynamics
G- Webex
H- Slack

Answer:

A, B, C, D, F, G, H

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) architecture, ThousandEyes provides a suite of native, 'out-of-the-box' integrations specifically designed for alert notifications. These integrations are configured via the Manage > Integrations or Alert Rules > Notifications tabs and allow the platform to push instant event data to a variety of collaboration, IT operations, and observability tools.

The verified list of native notification integrations includes:

ServiceNow (A): Facilitates direct incident management and automated ticketing workflows.

PagerDuty (B): Allows for automated incident escalation and on-call routing based on ThousandEyes alerts.

MS Teams (C) & Slack (H): Enable real-time chatops by pushing alert details and permalinks directly into specified channels for team collaboration.

Splunk (D): Utilizes the Cisco ThousandEyes App for Splunk to ingest alert data for historical analysis and correlation within a SIEM or logging platform.

AppDynamics (F): Sends alert notifications directly into an AppDynamics instance, allowing application owners to correlate network issues with APM metrics.12

Webex (G): Integrates with the Webex Control Hub and specific Webex spaces to provide unified visibility for collaboration performance.34

AWS (Option E) is not a native alert notification destination 5in this context; instead, ThousandEyes integrates with AWS for 'Cloud Insights' (ingesting VPC Flow Logs) and 'Test Recommendations' rather than acting as a receiver for event notifications like a chat or ITSM tool. By utilizing these native integrations, enterprise teams ensure that the right stakeholders are notified through their preferred communication channels the moment a performance threshold is breached, drastically improving response times across the organization.

Question 6

Question Type: MultipleChoice

Which type of test are we using for these dashboards (Executive and IT Operations)?

Options:

A- HTTP server
B- Page Load
C- Agent to server
D- FTP

Answer:

B

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) data analysis guidelines, identifying the underlying test type is the first step in interpreting a dashboard. By observing the metrics displayed across both the Executive and IT Operations dashboards, we can definitively identify the test type as Page Load (Option B).

A Page Load test is a Web-layer test that uses a browser engine (Chromium) to fully render a page and collect advanced metrics beyond simple availability. Evidence for this test type in the dashboards includes:

Page Completion Time: Displayed on the Executive dashboard map, this metric is specific to how much of the page successfully rendered.

DOM Load Time: Found in the IT Operations dashboard, the Document Object Model (DOM) time is a browser-centric metric that measures when the page structure has finished loading.

Waterfall Charts: While not a specific answer option, these are the foundation of Page Load tests, allowing for the timing of individual web components.

Simple HTTP Server tests (Option A) only measure availability and response codes, missing the rendering metrics seen here. Agent to server (Option C) is a network-layer test that provides path visualization but no browser-level data. FTP (Option D) is a protocol-specific test not used for web application monitoring. Therefore, the presence of DOM and Page Completion metrics confirms the Page Load test type.

Question 7

Question Type: MultipleChoice

An engineer deployed a Cisco ThousandEyes Enterprise Agent on a Meraki MX to monitor a critical SaaS application. Which kind of monitoring has the engineer set up?

Options:

A- active monitoring
B- passive monitoring
C- agentless monitoring
D- server monitoring

Answer:

A

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, understanding the distinction between different monitoring methodologies is fundamental to architecting an effective assurance strategy. When an engineer deploys a ThousandEyes Enterprise Agent on a Meraki MX appliance, they are implementing active monitoring.

Active monitoring, as defined in standard network assurance frameworks like RFC 7799, involves the generation of synthetic traffic or 'probes' that are sent across the network to a specific destination. These probes, which can utilize protocols such as ICMP, TCP, or HTTP/S, simulate real user transactions to measure performance metrics including latency, packet loss, jitter, and path visualization. The Enterprise Agent acts as a dedicated vantage point, executing these tests at scheduled intervals to provide a proactive baseline of network and application health. This allows the engineer to identify performance degradation or outages even when no real users are actively using the application, ensuring that issues are detected before they impact the business.

It is important to contrast this with passive monitoring (Option B). In the Meraki ecosystem, Meraki Insight (MI) natively performs passive monitoring by observing and analyzing actual user traffic flows (HTTP/S data) as they traverse the MX appliance without injecting additional traffic. While passive monitoring is excellent for understanding real-world user experience and server response times, it relies on existing traffic and cannot provide hop-by-hop path visualization across the Internet in the same way active synthetic probing does.

By integrating the ThousandEyes Enterprise Agent---which runs as a containerized service within the MX architecture---the engineer gains the benefits of active monitoring directly from the branch edge. This eliminates the need for separate hardware and provides deep, 'outside-in' and 'inside-out' visibility into SaaS application performance. Therefore, the deployment of a ThousandEyes agent explicitly enables active monitoring (Option A) to supplement the native passive capabilities of the Meraki platform.

Question 8

Question Type: MultipleChoice

An administrator has set up GPO properly, but realized ThousandEyes EPA was not deployed on one of the office PCs.9 What is the appropriate first step?

Options:

A- After GPO deployment, an administrator account must log in to deploy the EPA
B- Check that the PC belongs to the needed domain
C- Reboot the PC, this will restart GPO on the server
D- Reboot the Server, this will restart GPO on the PC

Answer:

B

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, troubleshooting deployment issues is as critical as the initial configuration. When utilizing Group Policy Objects (GPOs) for the distribution of the ThousandEyes Endpoint Agent (EPA), certain prerequisites must be met for the policy to take effect on a target machine.10

The most fundamental requirement is that the target PC must be a member of the domain or the specific Organizational Unit (OU) where the GPO is linked. Therefore, the appropriate first step is to check that the PC belongs to the needed domain (Option B). GPOs are scoped based on Active Directory membership; if a computer is in a workgroup or a different domain/OU that is not targeted by the policy, it will never receive the instruction to install the software, regardless of how 'properly' the GPO is configured on the server.11

The other options are technically incorrect or represent a misunderstanding of GPO mechanics:

Administrator Login (Option A): EPA installation via GPO is typically configured to run under the System account context during startup, meaning it is not dependent on a specific user (administrator or otherwise) logging in to trigger the deployment.

Rebooting to restart GPO on the server (Option C): Rebooting a client PC triggers a 'GPUpdate' request from the client to the server, but it does not 'restart' the GPO service on the server side.

Rebooting the Server (Option D): This is an invasive and unnecessary step that will not force a policy update on a specific client PC if that PC is not correctly joined to the domain or is experiencing a local networking issue.

By verifying domain membership first, the administrator ensures the basic trust and communication path required for GPO delivery is functional before moving to more complex troubleshooting steps like checking event logs or network connectivity.

Question 9

Question Type: MultipleChoice

A network administrator wants to establish a baseline for CPU utilization on their core routers. Which data source would be MOST appropriate for this purpose?

Options:

A- DNS resolution time from ThousandEyes tests
B- HTTP Server response times from ThousandEyes tests

C- SNMP data collected from the routers
D- Network path visualization from ThousandEyes tests

Answer:

C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework, baselining is the process of establishing a 'normal' performance profile for network infrastructure to enable the detection of anomalies. When the metric of interest is the internal health of a physical device, such as CPU utilization on a core router, SNMP (Simple Network Management Protocol) is the industry-standard data source.

SNMP provides direct visibility into the device's control plane and hardware performance. By polling specific Object Identifiers (OIDs) from the router's Management Information Base (MIB), a monitoring system like Cisco Catalyst Center or a third-party NMS can collect granular data on CPU cycles, memory allocation, and temperature. This 'inside-out' telemetry is essential for baselining because it reflects the actual resource consumption of the router during various traffic loads.

Conversely, ThousandEyes tests (Options A, B, and D) provide 'outside-in' synthetic data. While DNS resolution time (Option A), HTTP response times (Option B), and Path Visualization (Option D) are excellent for measuring end-to-end service delivery and network transit health, they do not report on the router's internal hardware state. For instance, a router could have 99% CPU utilization (indicating a potential crash), yet a ThousandEyes path test might still show a 'green' path if the data plane (ASICs) is still forwarding packets efficiently. Therefore, to establish a reliable baseline for hardware-specific metrics like CPU, SNMP data (Option C) is the only appropriate source among the choices.

Question 10

Question Type: MultipleChoice

Refer to the exhibit.

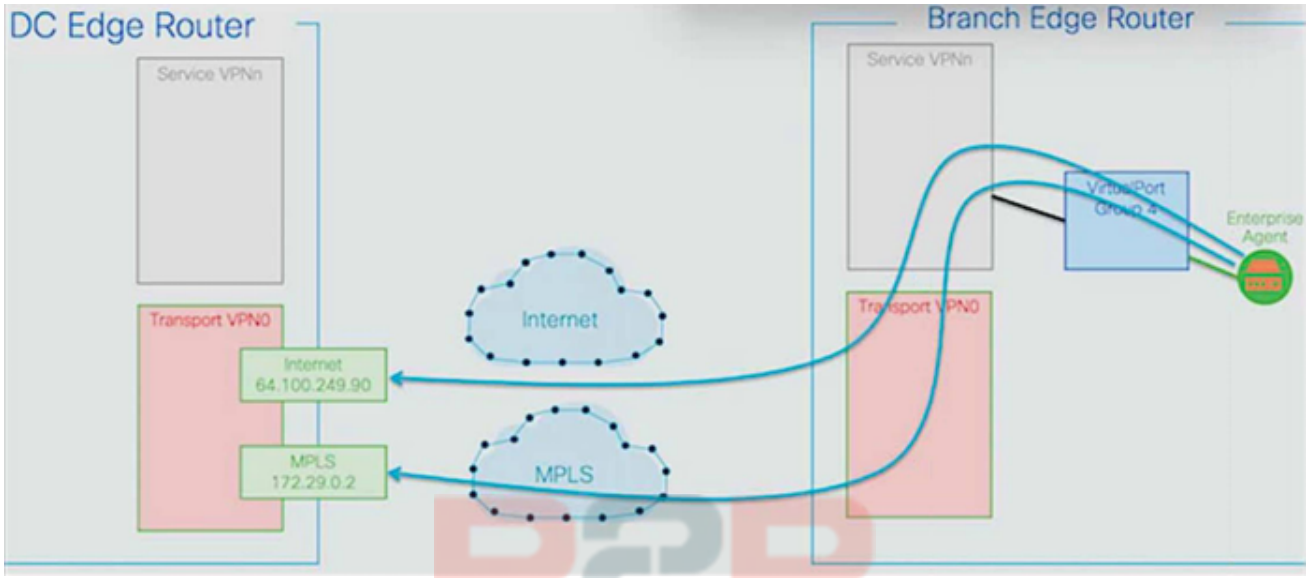

An engineer must use Cisco ThousandEyes testing to monitor their Cisco Catalyst SD-WAN fabric. Which SD-WAN component is being monitored by ThousandEyes?

Options:

A- underlay
B- IPsec tunnels
C- overlay
D- GRE tunnels

Answer:

A

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, understanding the visibility gap between the SD-WAN overlay and the transport underlay is a core competency. The provided exhibit illustrates a ThousandEyes Enterprise Agent deployed on a Branch Edge Router performing tests across two distinct paths: Internet (reaching a destination at 64.100.249.90) and MPLS (reaching a destination at 172.29.0.2).

According to the ENNA architecture guidelines, ThousandEyes is primarily utilized to provide hop-by-hop visibility into the underlay network. While SD-WAN controllers like vManage provide native monitoring for the overlay---the logical IPsec tunnels (Option B) that form the SD-WAN fabric---they often lack granular visibility into the physical service provider paths (the underlay) that carry those tunnels. The exhibit specifically highlights the agent probing the transport networks (Transport VPN0) directly, bypassing the overlay tunnels to measure the raw performance of the ISP and MPLS circuits.

By monitoring the underlay (Option A), the engineer can identify if high latency or packet loss is caused by a specific hop within the service provider's infrastructure or at a peering point. This

'underlay visibility' is critical for troubleshooting SD-WAN performance issues where the overlay may report a tunnel down, but the root cause lies in a BGP routing change or physical fiber cut in the provider network. ThousandEyes Enterprise Agents, natively integrated into Catalyst 8000 and ISR 4000 platforms, allow for this persistent underlay monitoring without additional hardware.

Overlay (Option C): While ThousandEyes can monitor overlay performance, the exhibit's focus on the raw IP addresses (Internet and MPLS) in the transport VPN indicates an underlay test.

IPsec/GRE Tunnels (Options B & D): These represent the transport mechanisms of the overlay. ThousandEyes probes the path under these tunnels to ensure the transport health is sufficient to support the fabric.

Question 11

Question Type: MultipleChoice

A company is noticing sporadic slowdowns in their web application performance, impacting user experience. They suspect it might be related to high CPU utilization on employee laptops, potentially caused by background processes. Which ThousandEyes alert type and condition combination would be most effective in identifying if endpoint CPU performance is contributing to this issue?

Options:

A- Real User Tests > Network Tests and Path Trace, End-to-End Packet Loss
B- Scheduled Tests > Endpoint Path Trace, Path length > #
C- Real User Tests > Endpoint, CPU utilization %
D- Scheduled Tests > Endpoint End-to-End (server), Memory load %

Answer:

C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, isolating performance issues on end-user machines requires a combination of application-layer experience data and local system telemetry. When an architect suspects that local hardware constraints, such as high CPU usage, are degrading the digital experience, the Endpoint Agent is the required vantage point.

The correct combination is Real User Tests > Endpoint, CPU utilization % (Option C). This alert

configuration is effective for several reasons:

Targeted Visibility: It directly monitors the user's device (Endpoint) where the hardware bottleneck is suspected to reside.

Real-Time Context: By utilizing Real User Tests (RUM), the agent gathers performance data during actual user activity (e.g., browsing the web application). This ensures that the CPU spikes are correlated with active application usage, providing a representative view of the impact on experience.

Threshold Alerting: The condition monitors for CPU utilization exceeding a defined percentage threshold (%). This allows IT teams to be notified only when the CPU load reaches problematic levels that are known to cause application sluggishness or browser 'freezing'.

Alternative options are less effective for this specific troubleshooting goal:

Option A: Monitors network loss, which would not identify a local CPU bottleneck.

Option B: Monitors path length (hops), which is a routing metric irrelevant to local hardware performance.

Option D: Monitors memory load during scheduled tests. While memory is a factor, the specific suspicion in this scenario is CPU utilization, and scheduled tests do not always capture the same load impact as a real user interacting with the browser.