




"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

[www.atmicnetworks .com](http://www.atmicnetworks.com)

Warning: Keep connected with our support team
for latest updates

Question: 1

Exhibit:



An engineer works to optimize a website by reducing the page-load time to below 500 ms. The engineer set up a Cisco ThousandEyes page-load test to baseline the current website performance. Which action should be recommended to reduce page-load time?

- A. Optimize the AJAX query calling functions.
- B. Move IMG elements to the bottom of the document body.
- C. Implement lazy loading for objects on the page.
- D. Use a CDN to load fonts faster.

Answer: C

Explanation:

In the context of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), analyzing page-load metrics within Cisco ThousandEyes requires identifying the primary bottlenecks that contribute to the Total Page Load Time. The provided screenshot displays a "Page Breakdown" of 7 resources totaling 953 kB. A critical observation of the pie chart reveals that Images (the teal-colored segment) constitute the vast majority of the page's payload and resource count.

When the goal is to reduce the page-load time from 1023 ms to below 500 ms, the engineer must target the heaviest components. Lazy loading is a design pattern that defers the initialization of non-critical resources at page load time. Instead of loading all images

simultaneously when the user first navigates to the URL, lazy loading ensures that images are only downloaded as they are about to enter the viewport. This significantly reduces the initial DOM load time and the total Page Load Time because the browser does not have to wait for large image files to be fully retrieved before declaring the page "loaded."

Alternative options are less effective in this specific scenario based on the data:

AJAX (XHR/Fetch): The chart shows that XHR and Fetch resources represent a negligible sliver of the total weight; optimizing them would yield minimal gains.


Moving IMG elements: While moving scripts to the bottom can help with rendering, moving image elements to the bottom of the body does not stop the browser from initiating the download requests immediately, thus failing to significantly reduce the total load time.

CDN for Fonts: The "Font" category is also a small fraction of the total 953 kB. While a CDN is a best practice for latency, it does not address the primary "weight" issue caused by the images.

Therefore, implementing lazy loading (Option C) is the most impactful recommendation. It directly addresses the largest resource consumer (Images) identified in the ThousandEyes Page Breakdown, allowing the engineer to reach the sub-500 ms performance target.

Question: 2

Refer to Exhibit:



A network engineer is deploying a Cisco ThousandEyes agent to monitor the network for a SaaS application without affecting the performance of the employee endpoints. Which ThousandEyes agent must be deployed to obtain the network metrics from branch A?

- A. Endpoint Agent
- B. Application Agent
- C. Enterprise Agent

D. Cloud Agent

Answer: C

Explanation:

In the framework of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), selecting the appropriate ThousandEyes agent type is critical to balancing visibility requirements with infrastructure constraints. For Branch A, the primary objective is to gain network-layer metrics (such as latency, packet loss, and jitter) and path visualization for a SaaS application while strictly avoiding any performance impact on employee endpoints.

The Enterprise Agent (Option C) is the correct choice because it is designed for "inside-out" monitoring from within the corporate network environment. These agents are lightweight software probes that can be deployed on existing network infrastructure, such as Cisco Catalyst 9300/9400 switches or Catalyst 8000 Edge Platforms, using Docker containers or virtual machines. By hosting the agent on the branch router or a dedicated local server, the engineer can execute synthetic tests to the SaaS provider's destination. This approach provides the necessary network vantage point from Branch A without requiring any software installation or resource consumption on the individual employee workstations (endpoints).

Other agent types do not satisfy the specific constraints of this scenario:

Endpoint Agents are installed directly on user devices (Windows/macOS) to provide "last-mile" visibility. However, they use the endpoint's CPU and memory, which contradicts the requirement to not affect endpoint performance.

Cloud Agents are maintained by Cisco in global ISP data centers. While they provide "outside-in" visibility, they cannot capture internal branch network characteristics or the specific path from Branch A's internal local area network.

Application Agent is a non-standard term and does not exist as a standalone agent type within the ThousandEyes architecture.

Therefore, deploying an Enterprise Agent within the branch infrastructure ensures that the network engineer obtains high-fidelity network metrics while keeping employee devices entirely unburdened.

Introduction to ThousandEyes

This video provides an essential overview of how ThousandEyes agents function within a CCNP-level enterprise network assurance strategy.

Question: 3

Refer to the exhibit.

```
curl -1 -XPOST https://apl.thauaandeyes.cora/v7/stream \
-H "Content-Type: application/json" \
-H "Authorization: Bearer SSEARER TOKEN" -d '{
  "type": "I",
  "tagHatch": [
```

"key": "TestKey", "value": "TestValue", "objectType": "test"

"atreamEndpointUrl": "https://Zexample.org", "customHeadera" : {
"test": "value"

Which integration type should be configured between ThousandEyes and Grafana?

A. opentelemetry

B. custom-webhook

C. push-api

D. poll-api

Answer: A

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, the evolution of network monitoring includes moving from periodic polling to real-time data streaming. The exhibit displays a curl command targeting the ThousandEyes API v7 /stream endpoint. When integrating ThousandEyes with high-performance observability platforms like Grafana, the standardized and recommended method for machine-to-machine data exchange is through **OpenTelemetry (OTel)**.

According to the ENNA architecture guidelines, the ThousandEyes Streaming API allows users to push granular test metrics (such as network latency, packet loss, and jitter) to external collectors in an OTel-compatible format. In the provided JSON payload, the "type" field is a mandatory parameter that defines the integration protocol. For Grafana, which natively supports OpenTelemetry Protocol (OTLP) via its OpenTelemetry Collector, the value must be set to "opentelemetry" (Option A). This tells the ThousandEyes streaming engine to encapsulate the data according to the OTel semantic conventions, ensuring that Grafana can correctly interpret and visualize the metrics without additional custom parsing logic.

While other options exist in the ThousandEyes ecosystem, they do not fit the specific API call shown for this use case:

Custom Webhooks (Option B) are typically used for event-driven alerts and notifications (e.g., sending a POST request when a threshold is breached) rather than continuous high-fidelity metric streaming.

Push-api and poll-api (Options C and D) are not valid "type" values within the context of the v7 /stream endpoint, as the streaming service specifically utilizes the OpenTelemetry framework for real-time delivery.

By selecting opentelemetry, the network engineer enables a robust "push-based" integration that provides real-time visibility into application performance and network health, leveraging Grafana's advanced dashboarding capabilities to analyze ThousandEyes telemetry data alongside other enterprise infrastructure metrics.

Introduction to ThousandEyes for OpenTelemetry

This video provides a foundational understanding of how ThousandEyes uses modern streaming frameworks to export critical performance data to external observability platforms.

Question: 4

DRAG DROP

Drag and drop the Cisco Network Assurance platforms from the left onto the corresponding business cases on the right.

Answer Area

Catalyst Center	Alert and analyze real-time business impact from applications on business metrics to pinpoint root causes of application problems
AppDynamics	Cloud or on-prem management system leveraging AI to connect, secure, and automate your network operations
ThousandEyes	Provide end-to-end visibility from every user to any application, over any network
Meraki	Cloud-based platform to manage and monitor your distributed network, and IoT technologies with end-to-end visibility

Answer:

Explanation:

AppDynamics

Catalyst Center

ThousandEyes

Meraki

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) architecture, each platform is positioned to address a distinct domain of visibility and management within the modern IT visibility framework.

AppDynamics (matched to Case 1) is specifically engineered for Full-Stack Observability and Application Performance Monitoring (APM). Its unique focus is linking technical application performance—such as code-level execution and database queries—directly to business outcomes and real-time business metrics. By doing so, it allows organizations to pinpoint how application latency impacts revenue or

user satisfaction, making it the primary choice for business impact analysis.

Catalyst Center, formerly DNA Center (matched to Case 2), serves as the foundational controller for campus and enterprise network infrastructure. It leverages AI-driven insights to automate network operations, enforce security policies, and provide at-a-glance health monitoring for local devices and clients. It is the definitive management system for connecting and securing "inside" the corporate network perimeter.


ThousandEyes (matched to Case 3) provides "Internet Intelligence" and is the key solution for monitoring the service delivery chain across environments the enterprise does not own. By utilizing a global network of Cloud, Enterprise, and Endpoint agents, it provides end-to-end visibility from any user location to any application (SaaS or Cloud) over any network (Internet, ISP, or WAN).

Meraki (matched to Case 4) is a cloud-first platform designed for managing highly distributed network environments and IoT deployments. It provides end-to-end visibility through a simplified dashboard that integrates wireless health, switching, and security across thousands of sites, making it ideal for lean IT teams managing distributed retail or branch locations.

Each platform utilizes different data collection methods—active synthetic testing for ThousandEyes and passive monitoring/AI telemetry for Catalyst Center—to ensure comprehensive network assurance across the entire enterprise ecosystem.

Question: 5

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?

- 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: 6

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?

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: 7

Refer to the exhibit.

login Page URL	http* /Aog* fMcroaoAonbn* coeVaacSdOcb-564M44M711-21 etkaaoeWMmO	Overr.de
logout Page URL	hn^e nogm^ooeolMne corna^SOrt 5ftM^tM Wil rafc^OI^	Overdo
Wertry Provider timer	MIper/iai window# net *9cS80ct> 5604 ^54 7-6 211-211^9608660	Override
Seme* Provider Heuer	hQpt ^app tncuwndovet oom	Overrifi*
	W idP ooNlguteU- rwedi lo p*x>j natch Set «^u* Sometvtw* Get a slew um^d 'Audiencia Rvitncoon'	
venfeabon ce^caiet	Certificate	EapvfatJon
	Microsoft Acura Federated SSO Ceroncato	Mar 1# 2023 W 2J31 UTC
		Oicard Change* Run Snglo S»gnO Test

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?

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 a 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: 8

An architect needs to analyze network path metrics from their internal network, specifically from the access layer to a cloud-hosted web server.¹ Which ThousandEyes agent is most appropriate for this task?

A. Synthetic Agent

B. Enterprise Agent

C. Cloud Agent

D. Endpoint Agent

Answer: B

Explanation:

In the framework of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), selecting the correct agent type depends heavily on the vantage point required for the specific observation. For this scenario, the architect must collect metrics from

the internal network access layer—the point closest to where the users or devices reside within the corporate perimeter—towards a cloud-hosted destination.

The Enterprise Agent (Option B) is the most appropriate choice because it is specifically designed to be deployed on infrastructure owned and managed by the organization. These agents are "inside-out" vantage points that can be installed directly on Cisco Catalyst 9300 or 9400 Series switches at the access layer using Docker containers. By deploying an Enterprise Agent at the access layer, the architect gains visibility into the entire network path, starting from the internal LAN, traversing the edge/WAN, and reaching into the cloud-hosted web server. This allows for the identification of issues such as local congestion, ISP peering problems, or cloud provider latency.

Other options do not meet the criteria:

Synthetic Agent (Option A): This is a distractor term. All ThousandEyes agents (Cloud, Enterprise, and Endpoint) are synthetic agents because they all perform active synthetic testing.

Cloud Agent (Option C): These are pre-deployed by Cisco in global ISP data centers and provide an "outside-in" view.¹⁴ While useful for monitoring public-facing availability, they cannot provide visibility into the internal network or the access layer of the organization.

Endpoint Agent (Option D): While these are installed on end-user machines and provide a "usercentric" view, they are generally not used for infrastructure-level path analysis from the access layer switches themselves.

Thus, the Enterprise Agent is the definitive choice for monitoring from the access layer to the cloud.

Question: 9

A network engineer is investigating widespread reports of poor performance for a data center-hosted web application. Which ThousandEyes agent type would be most effective for quickly identifying the root cause?

A. Synthetic Agent

B. Enterprise Agent

C. Endpoint Agent

D. Cloud Agent

Answer: D

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) guidelines, troubleshooting widespread performance issues for a public or data center-hosted application requires an "outside-in" perspective. When reports are widespread,¹⁸ the goal is to determine if the issue is global, regional, or specific to certain ISP paths leading to the data center.

The Cloud Agent (Option D) is the most effective tool for this task because these agents are maintained by Cisco ThousandEyes in over 240+ locations worldwide within Tier 1, 2, and 3 ISPs and cloud provider regions.¹⁹ Because they are pre-deployed and immediately available, a network engineer can instantly run tests from multiple global locations toward the data center-hosted application without having to install any software or manage any infrastructure. This allows the engineer to quickly compare performance metrics (latency, loss, and page load times) across different geographies. If Cloud Agents in London report no issues while those in New York report high packet loss, the engineer can immediately pinpoint the root cause as a regional ISP or peering issue rather than a failure within the data center itself.

Enterprise Agent (Option B): While these could be used if they were already installed in various branch offices, they require ownership of the infrastructure and deployment time. They are better suited for "inside-out" monitoring.

Endpoint Agent (Option C): These are useful for troubleshooting individual user experience but are not the "quickest" way to baseline global performance against a data center application during a widespread event.

Synthetic Agent (Option A): As noted previously, this is a generic term describing the underlying technology used by all ThousandEyes agent types.

Therefore, Cloud Agents provide the necessary breadth and immediate availability to perform rapid root cause analysis for widespread application performance issues.

Question: 10

An architect needs to measure end-user experience for internal web applications and SaaS products.²⁰ Which ThousandEyes agent should be deployed for this purpose?

A. Synthetic Agent

B. Enterprise Agent

C. Cloud Agent

D. Endpoint Agent

Answer: D

Explanation:

In the context of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), measuring the "lived experience" of an end-user requires data collection from the actual device being used to access the services. Unlike server-side or infrastructure-side monitoring, user experience (UX) monitoring must account for local variables like Wi-Fi signal quality, CPU/memory usage, and browser-level performance.

The Endpoint Agent (Option D) is the correct choice for this architecture. It is a lightweight software service installed directly on Windows or macOS workstations, as well as RoomOS devices. The Endpoint Agent provides a dual-monitoring approach: Real User Monitoring (RUM) and Scheduled Synthetic Tests.²⁴ RUM captures actual browser sessions to SaaS (e.g., Salesforce, Microsoft 365) or internal apps, providing a "Experience Score" and a detailed waterfall view of page load components.²⁵ Simultaneously, the agent can run background synthetic network tests to measure latency and path visualization from the user's specific location, whether they are in a branch office, at home on a VPN, or in a coffee shop.

Comparing other agents:

Enterprise Agents (Option B) can simulate a user at a branch office, but they cannot provide insight into the specific health of an individual's laptop or their unique Wi-Fi environment.

Cloud Agents (Option C) are entirely outside the user's network and cannot measure the performance of internal web applications or the "last mile" connectivity of the employee.

Synthetic Agent (Option A) remains a distractor term.

By deploying Endpoint Agents, the architect ensures they have granular, contextual data that correlates application performance directly with the user's device and local network environment.

Question: 11

A network engineer wants to measure their SD-WAN performance metrics. Which agent deployment method is most suitable for this scenario?

- A. Install an agent on the overlay network
- B. Install an agent on the DMZ
- C. Install an agent on their LAN
- D. Install an agent on the underlay network

Answer: D

Explanation:

In the context of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), understanding the visibility gap in SD-WAN environments is essential. While SD-WAN controllers provide native visibility into the overlay network (the logical IPsec tunnels and fabric health), they often lack granular insight into the physical transport or underlay network provided by ISPs or MPLS circuits.

According to the ENNA architecture guidelines, the most suitable method for measuring true SD-WAN performance is to install an agent on the underlay network (Option D). By deploying ThousandEyes Enterprise Agents directly on the transport-facing interfaces (Transport VPN0 in Cisco SD-WAN terminology), engineers can perform hop-by-hop path visualization and measure metrics like packet loss, latency, and jitter across the actual provider path. This is critical because performance degradation in the overlay is almost always a symptom of an issue in the underlay, such as BGP routing instabilities or physical link congestion at an ISP peering point.

Deploying agents on the overlay (Option A) only measures the performance of the tunnel itself, which may hide specific hop-level failures occurring in the public internet. Installing an agent on the LAN (Option C) or DMZ (Option B) adds local network noise to the metrics, making it harder to isolate if a problem exists within the corporate office or the service provider network. By focusing on the underlay, the engineer ensures they have the "internet intelligence" required to hold service providers accountable to SLAs and quickly resolve connectivity issues that impact the SD-WAN fabric.

Question: 12

A network engineer needs to monitor the performance of a business-critical web application accessed by remote employees connecting through a Cisco AnyConnect VPN. Which two agent deployment methods are most suitable for this scenario? (Choose two)

- A. Deploy ThousandEyes Cloud Agents in the same geographical regions as the remote employees.
- B. Integrate ThousandEyes with Cisco AppDynamics to monitor application performance from the server-side.
- C. Deploy ThousandEyes Enterprise Agents on the VPN concentrator where the AnyConnect clients

terminate.

- D. Utilize the ThousandEyes Endpoint Agent and deploy it on a subset of remote employee machines running Cisco AnyConnect.
- E. Configure ThousandEyes tests from Enterprise Agents located in the data center where the web application is hosted.

Answer: A, D

Explanation:

For the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) exam, monitoring remote workforces requires a strategy that captures both the user's local environment and the regional internet health. In a scenario involving Cisco AnyConnect VPN, the "last mile" connectivity of the employee is often the most significant variable in application performance.

Utilizing the ThousandEyes Endpoint Agent (Option D) is the most effective way to monitor this environment. Because the agent resides

directly on the remote employee's machine, it can monitor the performance of the web application both "inside" and "outside" the VPN tunnel. It provides visibility into the local Wi-Fi signal strength, the health of the AnyConnect client, and the latency experienced as traffic traverses the VPN headend. This allows engineers to differentiate between a slow home internet connection and an issue with the VPN concentrator.

Deploying ThousandEyes Cloud Agents (Option A) serves as a critical baseline. By running tests from Cloud Agents in the same regions as the remote employees, the engineer can determine if the "internet" in that region is healthy. If a Cloud Agent in London shows a perfect response time while an Endpoint Agent in London shows high latency, the engineer can immediately isolate the problem to the user's specific setup or the VPN path, rather than a regional ISP outage.

Other options are less suitable for monitoring the remote employee's experience:

AppDynamics (Option B) provides server-side code visibility but cannot see the user's home Wi-Fi or local network path.

Enterprise Agents on the VPN concentrator (Option C) can monitor the path from the data center to the app, but they cannot see the path from the user to the concentrator.

Enterprise Agents in the data center (Option E) provide an "inside-out" view of the app's health but miss the entire remote access experience.

Question: 13

Which of the following is an example of active monitoring in network performance management?

- A. Analyzing SNMP data to observe interface utilization on a router
- B. Capturing packets on a network segment to identify the top talkers
- C. Sending a continuous ping from one office to another to measure latency
- D. Collecting NetFlow records to analyze traffic patterns over time

Answer: C

Explanation:

Within the framework of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), network monitoring is categorized into two primary methodologies: active and passive monitoring. 1 Active monitoring (Option C) is characterized by the generation of synthetic or "probes" traffic specifically designed to measure network performance. 2 These probes simulate real-world user activity, such as HTTP requests, DNS queries, or ICMP pings, to baseline performance metrics like latency, jitter, and packet loss.

The core benefit of the active approach is its independence from actual user traffic. By sending a continuous ping or synthetic HTTP probe, an engineer can verify path availability and performance even during off-peak hours when no real users are on the network. In the context of Cisco ThousandEyes—a central platform in the ENNA certification—this is the primary mode of operation for Cloud, Enterprise, and Endpoint agents. For instance, a ThousandEyes network test proactively sends packets to a target IP or URL to visualize

the hop-by-hop underlay and overlay paths.

Conversely, options A, B, and D represent passive monitoring techniques. Passive monitoring involves observing and analyzing traffic that is already traversing the network.³ Methods such as SNMP (Option A) provide device-level health data like CPU load and interface utilization, while packet captures (Option B) and NetFlow (Option D) analyze the characteristics of existing user flows to determine top talkers or traffic patterns. While passive monitoring is excellent for volume and utilization analysis, it lacks the proactive capability to test a path's performance before a user encounters a failure. Therefore, sending a synthetic probe like a continuous ping is the definitive example of active monitoring.

Question: 14

What is a primary advantage of passive monitoring over active monitoring?

- A. Passive monitoring can measure the network's performance under synthetic conditions.
- B. Passive monitoring can provide real-time data on network performance without adding traffic to the network.⁴
- C. Passive monitoring allows for the generation of test traffic to simulate user behavior.
- D. Passive monitoring can directly measure the performance of specific network services or protocols.

Answer: B

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) architecture, a critical design consideration is the impact of the monitoring solution on the production environment. The primary advantage of passive monitoring (Option B) is its non-intrusive nature; it provides insights into network performance and traffic composition without injecting additional "synthetic" overhead into the data plane.

Passive techniques—such as Cisco Meraki Insight (MI), NetFlow, and SNMP—rely on the telemetry generated by existing user traffic or the device's own control plane. For example, Meraki Insight analyzes HTTP/S flows as they naturally pass through a Meraki MX appliance to derive application performance scores, rather than sending separate probes.⁵ This ensures that the monitoring tool itself does not consume bandwidth or contribute to network congestion, which is particularly vital in bandwidth-constrained branch environments or on high-utilization links.

In contrast, active monitoring (Options A and C) requires the deliberate generation of synthetic traffic, which can potentially skew results if the volume is too high or if the network is already at capacity. While active monitoring is essential for proactive troubleshooting (measuring performance before users complain), passive monitoring is the preferred method for long-term historical analysis of real user experience and infrastructure utilization because it captures what is actually happening on the wire. Option D is a shared capability; both types can measure specific services, but only passive monitoring does so while remaining transparent to the network load. Therefore, the lack of added traffic is the definitive advantage of the passive approach.

Question: 15

A network administrator observes a recurring pattern in their Cisco SD-WAN: during peak business hours, users at a specific branch office experience poor voice call quality, characterized by choppy audio and delays.⁶ The administrator suspects that network congestion is contributing to this issue and wants to leverage ThousandEyes WAN Insights to improve the situation proactively. Which capability of WAN Insights is most relevant to this scenario?

- A. Monitoring the public internet paths to the voice call service provider and identifying outages or performance bottlenecks.⁷
- B. Analyzing historical SD-WAN performance data during peak hours and recommending alternative paths that prioritize voice traffic based on its SLA.
- C. Generating synthetic voice traffic to proactively test network paths and identify potential congestion points during peak hours.⁸
- D. A set of network management tools that leverage SNMP and flow protocols into a single dashboard.
- E. Providing real-time alerts on security threats targeting voice traffic within the SD-WAN.⁹

Answer: B

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, ThousandEyes WAN Insights is a predictive analytics solution designed to transform network management from a reactive to a proactive model.¹⁰ In the scenario provided, the voice quality issues are recurring and peak-hour dependent, indicating a need for optimization rather than just standard real-time alerting.¹¹

The most relevant capability for this scenario is analyzing historical SD-WAN performance data to generate path recommendations (Option B). WAN Insights integrates with Cisco Catalyst SD-WAN Manager (vManage) and vAnalytics to ingest massive volumes of telemetry.¹² It uses advanced statistical models to analyze the performance of all active circuits (MPLS, Internet, etc.) over time. If the system identifies that a different available path consistently delivers a higher Quality of Experience (QoE) or better adherence to the voice SLA during those peak windows, it generates a recommendation to adjust the Application-Aware Routing (AAR) policy.¹³¹⁴

This predictive approach allows the administrator to fine-tune network policies in advance, ensuring that sensitive traffic like voice is automatically rerouted to the most stable path before the congestion impacts the users. Option A describes standard ThousandEyes synthetic testing, while Option C is incorrect because WAN Insights specifically uses existing SD-WAN telemetry rather than generating its own synthetic voice probes.¹⁵ Option D and E describe general NMS or security features not specific to WAN Insights' predictive mission. Thus, the proactive recommendation of alternative paths based on historical SLA adherence is the core function of WAN Insights for improving voice quality.¹⁷¹⁸¹⁹

Question: 16

ThousandEyes WAN Insights integrates with Cisco SD-WAN to provide visibility into network performance and generate path

recommendations. Which two data sources from the SD-WAN

environment are essential for WAN Insights to function? (Choose two)

- A. Configuration data from vManage, such as device templates and centralized policy settings.
- B. Historical network performance metrics collected by vAnalytics, including loss, latency, and jitter for SD-WAN tunnels.
- C. Data collected from ThousandEyes endpoint agents installed on end-user devices within the SD-WAN.
- D. Application traffic flow data from the SD-WAN data plane, categorized by vManage application lists.
- E. Syslog messages from SD-WAN edge routers, indicating device events and errors related to tunnel establishment.

Answer: B, D

Explanation:

The architecture for Designing and Implementing Enterprise Network Assurance (300-445 ENNA) specifies that ThousandEyes WAN Insights relies on deep integration with the Cisco SD-WAN management stack. To generate its predictive path recommendations, the platform must ingest specific telemetry data that reflects both the network's behavior and the applications traversing it.

The first essential data source is historical network performance metrics collected by vAnalytics (Option B). Before WAN Insights can be activated, vAnalytics must be enabled to collect and enrich raw network telemetry from the edge routers.³⁴ This data includes granular metrics for every SD-WAN tunnel, such as packet loss, latency, and jitter.³⁵ WAN Insights analyzes these historical trends to forecast future path quality and determine which transport circuits are most likely to meet application SLAs over a long-term period.

The second essential data source is application traffic flow data (Option D). WAN Insights must understand which applications are currently active in the fabric to prioritize recommendations for "business-critical" services like Office 365, Webex, or custom internal apps.³⁸ This information is ingested as flow records from the SD-WAN data plane and categorized based on the Application Lists defined in Cisco Catalyst SD-WAN Manager (vManage).

Options A and E are configuration or logging data that, while useful for general management, are not the raw telemetry inputs used by the WAN Insights predictive engine. Option C is incorrect because WAN Insights explicitly uses infrastructure telemetry rather than ThousandEyes agent-based synthetic data for its SD-WAN fabric calculations. By combining vAnalytics performance metrics and application flow data, WAN Insights can provide the "Predictive Path Recommendations" that are a hallmark of modern network assurance.

Question: 17

What type of agent is typically installed on Cisco SD-WAN devices as part of the ThousandEyes integration with vManage?

- A. Cloud Agent
- B. Browser Agent
- C. Enterprise Agent¹

D. Endpoint Agent²

Answer: C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, the integration of ThousandEyes with the Cisco SD-WAN fabric is a cornerstone of modern wide-area network visibility. The primary agent type used in this native integration is the Enterprise Agent (Option C).

According to the ENNA implementation guidelines, ThousandEyes Enterprise Agents run as containerized applications directly on the Cisco IOS XE software stack (version 17.6.1 and later) of supported hardware, such as Cisco Catalyst 8000 Edge Platforms and ISR 4000 Series routers.³ This deployment is orchestrated centrally through Cisco Catalyst SD-WAN Manager (formerly vManage), which handles the lifecycle management, including downloading the Docker-based agent image from the Cisco software repository and pushing the initial configuration to the branch edge devices.⁴

Once deployed, these Enterprise Agents function as internal vantage points that can measure performance across both the underlay transport and the SD-WAN overlay. By placing the agent on the router, engineers can execute synthetic tests to monitor mission-critical SaaS applications or data center services directly from the branch's perspective, avoiding the need for extra hardware probes at each site. This "agent-on-router" architecture significantly reduces both CapEx and OpEx while providing high-fidelity network intelligence into ISP peering, regional outages, and fabric health.

Cloud Agents (Option A) are managed by Cisco in global data centers and cannot be "installed" on customer edge routers.

Endpoint Agents (Option D) are designed for user workstations or laptops and do not provide the infrastructure-level visibility required for SD-WAN transport monitoring.

Browser Agent (Option B) is not a specific standalone agent type in this context; rather, browser-level testing is a capability often performed by Endpoint or Enterprise Agents.

Therefore, the Enterprise Agent is the correct verified agent type for Cisco SD-WAN vManage integration.

Question: 18

What is the primary purpose of integrating ThousandEyes with Meraki?

- A. To deploy Endpoint Agents for VPN connectivity monitoring
- B. To monitor external applications and services from SD-WAN sites
- C. To enhance cloud security and compliance
- D. To manage user access policies and permissions

Answer: B

Explanation:

The Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework highlights the integration between ThousandEyes and Cisco Meraki as a solution for "cross-domain assurance".⁵ The primary purpose of this integration is to monitor external applications and services from SD-WAN sites (Option B).

In a distributed Meraki environment, IT teams often struggle with visibility into the "Internet as a WAN," where performance issues may occur outside the local network perimeter. By embedding ThousandEyes Enterprise Agents natively within Meraki MX appliances, organizations can bridge the gap between internal LAN metrics and external service health.⁶ This integration allows for proactive monitoring of SaaS platforms (like Microsoft 365, Salesforce, and Webex) and other public-facing dependencies using synthetic probes. It complements the native Meraki Insight (MI), which provides passive monitoring of real user traffic, by adding active path visualization and hop-by-hop analysis across the Internet.

Key advantages of this integration include:

One-Click Activation: Enabling the ThousandEyes agent directly from the Meraki Dashboard without additional hardware.⁷

Pre-configured Templates: Using built-in test templates for common SaaS applications to accelerate troubleshooting.⁸

Isolation of Fault Domains: Quickly determining if a user's lag is caused by a local Wi-Fi issue (via Meraki wireless metrics) or an ISP routing problem (via ThousandEyes path data).⁹

While ThousandEyes does provide visibility for VPN and security, Options A, C, and D are not the primary focus of the specific Meraki-ThousandEyes integration architecture, which is centered on extending application performance assurance to distributed branch locations.

Question: 19

What type of data does ThousandEyes use to diagnose when integrated with Cisco Secure Client?

- A. Data related to network hardware configurations
- B. Data related to user activity and behavior
- C. Network performance data from the user's device
- D. Data related to secure web gateway performance

Answer: C

Explanation:

Under the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) guidelines, the integration between ThousandEyes and Cisco Secure Client (formerly AnyConnect) is designed to provide visibility into the hybrid workforce experience. The integration primarily leverages network performance data from the user's device (Option C) to diagnose connectivity and application issues.

This data is collected by the ThousandEyes Endpoint Agent, which is deployed as a module within the Cisco Secure Client.¹⁰ The agent captures real-time telemetry from the user's laptop or workstation, including Wi-Fi signal strength, CPU and memory utilization, and local network gateway latency. Specifically, for remote users, it monitors the health of the VPN tunnel and the performance of applications as seen from the end-user's vantage point.

When a user reports a "slow application," this integrated telemetry allows IT teams to determine if the root cause is the user's home Wi-Fi, a saturated VPN concentrator, or an issue within the ISP underlay. The agent performs Automated Session Testing (AST), which maps the network path as soon as a user joins a critical meeting or accesses a SaaS tool, providing a granular view of every hop between the device and the service destination. Unlike hardware configuration data (Option A) or behavioral analytics (Option B), the focus here is strictly on the network performance metrics that impact digital experience.

Therefore, the collection of device-centric network performance data is the core function of the Cisco Secure Client and ThousandEyes integration.

Question: 20

What advantage does the integration of ThousandEyes with Cisco technologies offer for troubleshooting?

- A. It eliminates the need for manual data entry.
- B. It provides real-time virtual assistance to end-users.
- C. It automates network configuration changes based on user feedback.
- D. It allows for quick identification and resolution of performance issues.¹¹

Answer: D

Explanation:

The Designing and Implementing Enterprise Network Assurance (300-445 ENNA) certification emphasizes that the primary value proposition of the Cisco "Assurance Stack" is the reduction of Mean Time to Identification (MTTI) and Mean Time to Resolution (MTTR).

The integration of ThousandEyes across Cisco's portfolio—including Catalyst, Meraki, and SD-WAN—allows for quick identification and resolution of performance issues (Option D).¹²

By embedding ThousandEyes agents into the existing network infrastructure, Cisco enables "end-to-end visibility" that spans domains the enterprise traditionally does not control, such as the public internet and SaaS environments.¹³ During troubleshooting, this cross-platform visibility allows network engineers to immediately correlate internal network health (from Catalyst Center or Meraki Dashboard) with external path visualization (from ThousandEyes). For example, when a user experiences poor video quality in a Webex meeting, the integration allows the engineer to "crosslaunch" from the Webex Control Hub directly into a ThousandEyes path view.¹⁴ This pinpoint accuracy avoids the "blame game" between network, application, and ISP teams by providing a "single source of truth" regarding where the packet loss or latency is occurring.

While some automation exists (Option C), the core benefit isn't automatic configuration changes based on feedback, but rather providing the actionable insights required for manual or policy-based remediation. Similarly, while it streamlines data gathering, its ultimate purpose is the speed of resolution in complex, hybrid environments.

Question: 21

The network team has deployed Webex RoomOS Endpoint Agents and integrated Webex Control

Hub with ThousandEyes. The VoIP team wants to know which metrics they can collect from the Webex Control Hub view. Where does the VoIP team find the network data?

- A. Devices
- B. Network Path
- C. Users
- D. Settings

Answer: B

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, the integration between ThousandEyes and Webex Control Hub provides a streamlined troubleshooting experience for collaboration services. For the VoIP team to access the specific ThousandEyes network telemetry—such as latency, loss, and jitter—they must navigate to the Network Path (Option B) section within the Troubleshooting tab of the Control Hub.¹⁵

The Network Path visualization is a direct result of the ThousandEyes Endpoint Agent data being pulled into the Webex interface.¹⁶ When a user or RoomOS device experiences poor audio or video quality during a meeting, the Control Hub's troubleshooting view displays a "Network Path" line under the participant's details.¹⁷ By clicking on this line, the VoIP team can see a hop-by-hop breakdown of the entire route from the collaboration device to the Webex media node. This view highlights specific hops where performance is "Poor" (red), "Fair" (yellow), or "Good" (green) based on predefined thresholds for latency (>400ms) or loss (>5%).

While "Devices" (Option A) is where the agents are activated, and "Users" (Option C) allows for selecting a specific participant, the actual telemetry metrics and the visualization of the network route are strictly located in the Network Path view. This integration eliminates the need for the VoIP team to leave the Webex environment for initial triage, as they can identify if a problem is local to the branch office or deep within a service provider's network directly from the "Network Path" dashboard.

Question: 22

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?

- 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: 23

What is an important consideration when choosing a time period for collecting data to establish a baseline for interface utilization on a critical network link?

- A. Selecting the time period with the lowest network traffic volume.
- B. Ensuring the time period aligns with the organization's financial year.
- C. Capturing both peak and off-peak traffic patterns for a representative view.
- D. Limiting the time period to minimize the amount of data that needs to be analyzed.

Answer: C

Explanation:

Establishing an accurate baseline is a cornerstone of the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) methodology. A baseline must represent the "typical" behavior of the network across its standard operating cycle to be useful for anomaly detection and capacity planning. When monitoring interface utilization on a critical link, it is vital to capture both peak and off-peak traffic patterns (Option C).

Most enterprise networks exhibit cyclical traffic patterns, often referred to as "diurnal" cycles, where utilization spikes during business hours (peak) and drops significantly at night or on weekends (off-peak). If an engineer only collects data during the lowest volume period (Option A), the baseline will be unrealistically low, leading to "false positive" alerts when normal business traffic returns.

Conversely, only capturing peak data might hide background synchronization tasks or backup windows that occur during off-hours.

By selecting a representative time period—typically one full week or a month—the engineer ensures the baseline accounts for variations such as Monday morning "logon storms," mid-week steady states, and weekend maintenance. This comprehensive view allows the assurance platform (such as Cisco Catalyst Center or ThousandEyes) to calculate a standard deviation and establish dynamic thresholds. If utilization deviates from these historically representative norms, the system can trigger an alert based on a true anomaly rather than a predictable daily surge. Aligning with a financial year (Option B) is irrelevant to technical performance, and limiting the data (Option D) compromises the baseline's statistical validity.

Question: 24

Your organization wants to be notified of an event as soon as it is triggered by an alert threshold. This notification should be sent to your ITSM and generate an incident so it can be responded to appropriately. What kind of integration should you use?

- A. ServiceNow Integration
- B. DNA Center Integration
- C. Custom Webhooks
- D. Alerts API

Answer: A

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, the objective of modern network assurance is to bridge the gap between "visibility" and "action". When an organization requires an automated workflow to handle network performance anomalies, the most efficient architecture is the native ServiceNow Integration (Option A). This integration is categorized as a "Custom-Built" or native integration within the ThousandEyes platform, designed specifically to facilitate the delivery of direct notifications into a ServiceNow account.

According to the ENNA implementation standards, the ThousandEyes ServiceNow integration utilizes the ServiceNow Incident Management module. When a predefined alert rule (such as a 5% packet loss threshold on a critical SaaS path) is violated, ThousandEyes triggers an event and immediately pushes the alert data to ServiceNow via an OAuth-authenticated connection. Within ServiceNow, this data is used to automatically generate an Incident, complete with relevant metadata such as the test name, agent

location, and the specific metrics that triggered the violation. This automation eliminates the manual overhead of "copy-pasting" alert details from a monitoring dashboard into a ticketing system, thereby significantly reducing the Mean Time to Identification (MTTI). While Custom Webhooks (Option C) can achieve a similar result by sending JSON payloads to a REST API, they require additional development effort to parse the data on the receiver side. The native ServiceNow integration provides a pre-configured template that maps ThousandEyes alert fields directly to ServiceNow incident fields, offering a "one-click" setup experience that is preferred for enterprise-grade deployments. Options B and D are irrelevant for the specific goal of ITSM incident generation. Therefore, for direct ITSM notification and incident creation, the native ServiceNow Integration is the verified recommendation.

Question: 25

You have been tasked with creating a dashboard in your organization's Observability platform. This dashboard should have data that is streamed in real-time and used to populate data for tables, graphs, charts, and other formats. What kind of integration should you use?

- A. API Endpoints
- B. OpenTelemetry
- C. DNA Center Integration
- D. Alert Thresholds

Answer: B

Explanation:

Within the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework, the transition from "polling" to "streaming" is a major architectural shift. To populate real-time dashboards in external observability platforms like Grafana, Splunk, or AppDynamics, the architect should utilize the OpenTelemetry (OTel) integration (Option B).

ThousandEyes for OpenTelemetry is a push-based API built on the standardized OpenTelemetry Protocol (OTLP). Unlike traditional REST API polling (Option A), which retrieves data at fixed intervals and can be subject to rate limiting and latency, the OTel integration allows ThousandEyes to stream granular network metrics as they are collected. These metrics—including latency, loss, jitter, and HTTP response times—are exported in a standardized format that is natively understood by modern observability backends. This allows the platform to populate complex visualizations such as timeseries graphs, heatmaps, and multi-metric tables in near real-time, providing a "single pane of glass" view that correlates network performance with application and infrastructure telemetry.

A key advantage of the OTel approach is data portability and correlation. By applying metadata tags to ThousandEyes tests, the data can be filtered and categorized within the external dashboard to match the organization's business logic (e.g., grouping by region or application tier). This enables SREs and NetOps teams to quickly identify if a performance dip in an application dashboard correlates with a spike in internet latency measured by ThousandEyes. Options C and D do not provide the streaming data pipeline required for

real-time external dashboard population. Thus, OpenTelemetry is the definitive choice for high-fidelity, real-time observability integration.

Question: 26

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.

- 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 in 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: 27

A network engineer deploys a ThousandEyes Docker agent on a switch using app-hosting. The agent needs to communicate through a proxy server, but this configuration was missed during the initial deployment. The engineer adds the proxy settings to the app-hosting configuration. What is the next step to ensure the agent uses the proxy and appears online in the ThousandEyes portal?

- A. Restart the container using `app-hosting stop appid agentname` followed by `app-hosting start appid agentname`
- B. Reinstall the agent using the `app-hosting install` command with the correct proxy settings
- C. Execute the full agent lifecycle: `app-hosting stop appid agentname`, `app-hosting deactivate appid agentname`, `app-hosting activate appid agentname`, `app-hosting start appid agentname`
- D. No action required; the agent will pick up the configuration automatically

Answer: C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) implementation guide, the Cisco IOS XE Application Hosting lifecycle is a critical concept for troubleshooting and configuration management. When an application environment variable—such as proxy settings—is modified in the app-hosting configuration, a simple restart of the container is insufficient.

The correct procedure involves moving the application back to the Configured state before bringing it back to Running. This requires a four-step lifecycle execution (Option C):

Stop: Terminates the current running instance of the container.

Deactivate: Releases the hardware resources (CPU, Memory, Virtual Interfaces) and, crucially, detaches the previous runtime configuration.

Activate: Re-allocates the resources and injects the new configuration parameters (including the updated proxy settings) into the container's environment.

Start: Initiates the container with the newly applied configuration.

Without the deactivate/activate sequence, the container will continue to use the environment variables present at the time of its initial activation, causing the agent to remain offline as it fails to reach the ThousandEyes portal without the proxy. Reinstalling (Option B) is an unnecessary and timeconsuming step that involves re-downloading or re-copying the image, while No action (Option D) will result in no change to the agent's connectivity status.

Question: 28

What are the different ways to deploy a ThousandEyes Agent in a Switch? (Choose all that apply)

- A. Application Hosting
- B. Catalyst Center (formerly DNA Center)
- C. Catalyst SD-WAN Manager (formerly vManage)
- D. From the ThousandEyes Portal in the "Enterprise & Cloud Agent" section
- E. All of the above

Answer: E

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) architecture, the deployment of ThousandEyes Enterprise Agents on switching infrastructure is designed to be highly flexible, catering to different management styles and network scales. Modern Cisco switches, specifically the Catalyst 9300 and 9400 Series, support the execution of containerized applications directly on the switch's hardware through the Cisco IOS XE Application Hosting framework.¹

Application Hosting (Option A) refers to the manual method using the Cisco IOS XE Command Line Interface (CLI). An engineer can use app-hosting commands to install, configure, and manage the Docker-based Enterprise Agent.² This provides granular control over resource allocation and networking for the container. Catalyst Center (Option B) provides a GUI-driven, automated workflow to deploy agents at scale across an entire campus fabric.³ It simplifies the lifecycle management by handling the installation and activation across multiple switches simultaneously. For switches acting as edge devices in an SD-WAN fabric, Catalyst SD-WAN Manager (Option C) orchestrates the deployment through centralized policies and templates, allowing the agent to be pushed as part of a device configuration. Finally, while the ThousandEyes Portal (Option D) is the management plane for tests, it is also where the Docker image and Account Group Token are obtained for any of the aforementioned deployment methods.

Each method provides the same end result: an Enterprise Agent running "on-box" to provide deep visibility into the network path starting directly from the access or core layer. Therefore, All of the above (Option E) is the correct answer as it encompasses the manual, automated, and orchestrated methods available for switch-based deployment as per the 300-445 ENNA official guidelines.

Question: 29

What Meraki platform supports ThousandEyes?

- A. Meraki MX (Security Appliances)
- B. Meraki MR Series (Wireless Access Points)
- C. Meraki MS Series (Switches)
- D. Meraki MV (Smart Cameras)
- E. Meraki MG (Cellular Gateways)
- F. All of the above

Answer: A

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, the native integration of ThousandEyes within the Meraki portfolio is specifically focused on the Meraki MX security and SD-WAN appliances. This integration allows users to install ThousandEyes Enterprise Agents directly on supported Meraki MX hardware, providing critical visibility into the performance of services that distributed users rely on.⁴

The Meraki MX (Option A) supports the ThousandEyes Enterprise Agent as a containerized service that can be activated with a single click within the Meraki Dashboard.⁵ This integration is designed to provide better monitoring and testing capabilities for customers interested in improving the quality of their experience and adding the appropriate SD-WAN policies to optimize network performance.

By running the agent natively on the MX, administrators can monitor the performance of external applications and services directly from the SD-WAN sites, bridging the gap between local branch health and global internet performance.⁶

While Cisco continues to expand its assurance integrations across the portfolio, the MR (Wireless), MS (Switching), MV (Camera), and MG (Cellular) series do not currently support the native, "on-box" execution of the ThousandEyes Enterprise Agent. For these other platforms, visibility is typically achieved through passive monitoring via Meraki Insight (MI) or by deploying ThousandEyes agents on connected client devices (Endpoint Agents) or nearby infrastructure. Therefore, the Meraki MX is the only verified platform supporting native ThousandEyes agent deployment in this context.

Question: 30

A network engineer deploys a ThousandEyes Docker agent on a switch using app-hosting.⁷ The agent needs to communicate through a proxy server, but this configuration was missed during the initial deployment. The engineer adds the proxy settings to the app-hosting configuration. What is the next step to ensure the agent uses the proxy and appears online in the ThousandEyes portal?

- A. Restart the container using app-hosting stop appid agentname followed by app-hosting start appid agentname

- B. Reinstall the agent using the app-hosting install command with the correct proxy settings
- C. Execute the full agent lifecycle: app-hosting stop appid agentname, app-hosting deactivate appid agentname, app-hosting activate appid agentname, app-hosting start appid agentname
- D. No action required; the agent will pick up the configuration automatically

Answer: C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) implementation guide, the Cisco IOS XE Application Hosting lifecycle is a critical concept for troubleshooting and configuration management. When an application environment variable—such as proxy settings or account tokens—is modified in the app-hosting configuration, a simple restart of the container is insufficient to apply those changes.

The correct procedure involves moving the application back through its lifecycle states to ensure the new environment variables are injected. This requires executing the full agent lifecycle (Option C): Stop, Deactivate, Activate, and Start.

Stop: Terminates the current running process of the container.

Deactivate: Releases the allocated hardware resources (CPU, Memory) and, most importantly, clears the runtime environment.

Activate: Re-allocates the resources and injects the updated configuration parameters into the container's environment variables.

Start: Initiates the container with the newly applied configuration settings.

After the application is up and running, the ThousandEyes-agent process attempts to connect to the controller in the cloud environment. Without the deactivate/activate sequence, the container continues to use the environment variables present at the time of its initial activation, causing the agent to remain offline as it fails to reach the ThousandEyes portal without the proxy. Reinstalling (Option B) would work but is an unnecessarily lengthy process, while No action (Option D) will result in no change to the agent's connectivity status.

Question: 31

Which deployment option should a network administrator use to deploy the ThousandEyes Endpoint Agent to all users on their internal domain using a Microsoft Domain Controller?

- A. Microsoft Intune
- B. Group Policy Objects
- C. JAMF

D. Power Shell

Answer: B

Explanation:

For the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) exam, understanding enterprise-scale deployment of the Endpoint Agent (EPA) is vital for ensuring comprehensive user-centric visibility. When an organization utilizes a Microsoft Domain Controller to manage its internal domain, the standard mechanism for automated software distribution is Group Policy Objects (GPOs) (Option B).

Deploying the ThousandEyes Endpoint Agent via GPO allows a network administrator to push the .msi installer package to all Windows-based workstations joined to the domain. This method is highly effective for large-scale environments because it ensures that the agent is installed automatically upon machine startup or user login, without requiring manual intervention from the end-user or the IT staff at each individual machine. The GPO can be configured to perform a "silent installation," which runs in the background, ensuring a seamless experience for the employees while providing the IT team with the necessary performance metrics.

While Microsoft Intune (Option A) is a modern cloud-based endpoint management solution, the question specifically references a Microsoft Domain Controller, making GPO the most direct and traditional choice for that specific infrastructure. JAMF (Option C) is specifically for Apple (macOS/iOS) device management and would not be used for a Windows domain deployment. PowerShell (Option D) can be used for scripting, but it lacks the centralized policy enforcement and automatic compliance checking provided by GPOs in an Active Directory environment. Therefore, Group Policy Objects is the verified deployment method for this scenario.

Question: 32

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

- 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 "GPOUpdate" 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: 33

Which strategy is most effective for a scalable, secure, and minimally disruptive deployment of ThousandEyes Endpoint Agents to Windows users?

- A. Manually install the Endpoint Agent on each device
- B. Use a centralized software deployment tool (e.g., GPOs, Intune) that supports silent installation to deploy the Endpoint Agent
- C. Email employees a download link for the Endpoint Agent and request they install it on their devices
- D. Provide a web portal where employees can log in and download the Endpoint Agent

Answer: B

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) best practices, the goal of a successful endpoint assurance rollout is to maximize coverage while minimizing the burden on both IT staff and end-users. The most effective strategy for achieving this is to use a centralized software deployment tool (e.g., GPOs, Intune) that supports silent installation (Option

B).

Centralized deployment tools allow administrators to push the ThousandEyes Endpoint Agent to thousands of machines simultaneously from a single management console.¹² By using silent installation parameters (such as /quiet or /qn for MSI packages), the agent is installed in the background without requiring any user interaction or even showing a setup wizard.¹³ This "minimally disruptive" approach ensures that business operations are not interrupted. Furthermore, centralized tools provide a "secure" deployment method because the IT team retains control over the version of the agent being installed and can ensure that the Account Group Token is correctly embedded, which is necessary for the agent to register with the ThousandEyes dashboard.

Contrasting this with other methods:

Manual Installation (Option A): This is not scalable and is highly labor-intensive, making it unsuitable for large enterprises.

Emailing links (Option C) or Web Portals (Option D): These methods rely on the end-user to perform the installation correctly and to have the necessary local administrative privileges. This often leads to low adoption rates, inconsistent configurations, and security risks associated with users running installers manually.

By leveraging existing enterprise infrastructure like Active Directory GPOs or Microsoft Intune, organizations can ensure a 100% compliant deployment that immediately begins providing the "user-centric" visibility required for modern network assurance.

Question: 34

Refer to the exhibit.

Target

10.10.10.10

Protocol

TCP

v

Port

80

Probing Mode

1 Prefer SACK

Force SACK

Force SYN

Path Trace Mode

☐ In Session

Interval

2 minutes

Agents

Select agent(s)

Alerts

☒ Enable

1 of 3 alert rules selected

Edit Alert Rules

Which setting should be enabled for this network Agent to Server test to avoid test traffic being detected by firewalls as malicious?

- A. Path Trace Mode: In Session
- B. Protocol: TCP
- C. Port: 5000
- D. Probing Mode: Force SYN

Answer: A

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) architecture, a critical challenge in active synthetic monitoring is ensuring that test probes accurately reflect user traffic without being dropped by intermediate security appliances. Standard network tests often utilize separate connections for measuring performance metrics (latency/loss) and for path discovery (hop-by-hop visualization). This behavior can be flagged by stateful firewalls or Intrusion Prevention Systems (IPS) as suspicious or malicious scanning activity.

To mitigate this, the engineer should enable Path Trace Mode: In Session (Option A). When this mode

is active, ThousandEyes performs path discovery using the exact same TCP session established for the performance measurement. By embedding path discovery probes within an active, established session, the traffic appears to firewalls as part of a legitimate, ongoing communication stream rather than an independent series of probes with varying TTL values that might trigger "anti-spoofing" or "scanning" alerts.

Reviewing the alternative options:

Protocol: TCP (Option B): While using TCP is generally more firewall-friendly than ICMP, the exhibit shows TCP is already selected. The issue is not the protocol itself, but how the path discovery probes are handled relative to the session.

Port: 5000 (Option C): Changing the port to a non-standard value like 5000 often makes traffic more likely to be scrutinized or blocked by default firewall policies compared to standard web ports like 80.

Probing Mode: Force SYN (Option D): Forcing SYN packets is a technique used to bypass certain types of load balancers but does not address the fundamental issue of path discovery probes being seen as a separate, malicious scan by stateful inspection engines.

Therefore, enabling In Session path trace mode is the most effective way to ensure consistent visibility through security-hardened environments.

Question: 35

Refer to the exhibit.

Basic Configuration		Advanced Settings	
Domain	intemal-domain.com	IN	
Interval	10 minutes	▼	
Agents	1 of 8 selected		
DNS Servers	ns47.domaincontrol.com. x 1 more		Q Lookup Servers
Alerts	Q Enable		
	2 of 5 alert rules selected		▼ Edit Alert Rules

A network admin has been tasked with monitoring the IPv6 record and name server resolution times with different agents. Select the two actions that the engineer must take to meet the requirements.

- A. Create a DNS Server test monitoring the A record
- B. Create a DNS Server test monitoring the AAAA record
- C. Create a DNS Trace test monitoring the ANY record
- D. Create a DNS Server test monitoring the NS record
- E. Create a DNS Trace test monitoring the NS record

Answer: B, D

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, DNS testing is a fundamental component of ensuring application availability and performance. The ThousandEyes platform provides two primary types of DNS tests: DNS Server and DNS Trace.

To satisfy the specific requirements of monitoring IPv6 records and name server resolution times, the engineer must utilize the DNS Server test. This test type is designed to measure the performance of specific DNS servers by querying them for particular resource records.

Monitoring IPv6 Records: In the DNS system, IPv6 address records are known as AAAA records.

Therefore, to monitor the availability and resolution of IPv6 addresses for a domain, the engineer must create a DNS Server test and set the query type to AAAA (Option B).

Monitoring Name Server Resolution Times: To measure how long it takes for a DNS server to return the authoritative name servers for a domain, the engineer must query for NS (Name Server) records. By creating a DNS Server test targeting the NS record (Option D), the platform can provide granular metrics on the response time and availability of those specific records from the perspective of the selected agents.

DNS Trace tests (Options C and E) are less suitable for measuring specific resolution times of a single server; instead, they are used to visualize the entire delegation path from the root servers down to the authoritative servers to identify where a "break" in the chain occurs. Furthermore, an A record query (Option A) only provides the IPv4 address, failing the requirement to monitor IPv6. By selecting DNS Server tests for AAAA and NS records, the administrator ensures they have the precise performance data needed to baseline and troubleshoot the organization's DNS health.

Question: 36

Add New Test

New Test

Layer

Test Type

Test Name

Test Description

Basic Configuration

URL

Schedule

Interval

Agents

Q Searching...

☐

Q Europe

0 of 1 agent

Ci abfentan-va-1

0 North America

0 of 7 agents

O agt-ciscolive-skills-1550

☐ CertBrk

eastl-agent-1

O tf-agt-ciscolive

WebinarAgent!

O west-2-ag-1

Q west-2-ag-2

Select agent(s)

Show

Enabled

S3

S3

S3

S3

Views Enabled for This Test

Web

Page Load

HTTP Server

Overview

Path Visualization

BGP Route Visualization

Built-In Labels

(3) Cloud

JJ Enterprise

™ Enterprise Cluster

© IPv4 Compatible

© IPv6 Compatible

\$ Proxied

Single Homed

Dual Chromium

An engineer is trying to configure a Page Load test and is trying to assign the east1-agent-1 to run it, but cannot see that agent listed. What is the reason that the designated agent does not appear?

- A. The agent is not running
- B. The agent is disabled
- C. The agent is still registering
- D. The agent does not support Page load tests

Answer: D

Explanation:

In the framework of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), selecting the appropriate agent for specific test types is a critical configuration step. The ThousandEyes platform utilizes different agent binaries and capabilities depending on the underlying hardware and software environment. Specifically, Page Load tests and Transaction tests require the agent to have a browser engine (Chromium) installed and functional to render web content and measure DOM-related metrics.

According to the ENNA implementation guidelines, not all Enterprise Agents support these browserbased tests. The exhibit displays a "New Test" configuration screen where the engineer is searching for "east1-agent-1". While other agents like "west-2-ag-1" are visible, the designated agent is missing because it does not support Page load tests (Option D). This common occurrence happens when an Enterprise Agent is deployed on resource-constrained hardware—such as certain Cisco Catalyst 9000 switches or Cisco ISR 4000 series routers—where the Chromium browser package is either not supported or has not been enabled due to memory and CPU limitations.

When an agent lacks the required capabilities for a specific test layer (e.g., Web layer tests), it is automatically filtered out from the selection list in the ThousandEyes portal to prevent invalid test configurations. To identify compatible agents, an engineer can look for the "Browser" or "Dual Chromium" labels in the agent settings. If the "east1-agent-1" were merely offline (Option A), disabled (Option B), or still registering (Option C), it would typically still appear in the list but with a status indicator showing its unavailability, rather than being entirely hidden from the test assignment menu. Therefore, the absence of the agent in a web-layer test configuration explicitly indicates a lack of functional support for the required browser-based metrics.

Question: 37

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?

- 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: 38

To monitor communication and measure network performance from branch offices in San Francisco and Texas to the data center in North Virginia, which combination of test type and target is the most appropriate?

- A. Agent-to-server test type and Cloud Agent

- B. Cloud Agent and HTTP Server
- C. Enterprise Agent and Agent-to-agent test type
- D. HTTP Server and DNS Server
- E. Agent-to-server test type and DNS Server

Answer: C

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) architecture, monitoring site-to-site connectivity between managed locations requires vantage points at both ends of the connection. In this scenario, the organization needs to measure performance between branch offices (San Francisco and Texas) and a central data center (North Virginia).

The most appropriate combination is using Enterprise Agents with the Agent-to-agent test type (Option C). Enterprise Agents are software probes deployed on infrastructure owned or controlled by the organization, such as switches or routers in the branch offices and servers in the data center. Unlike Agent-to-server tests, which only provide one-way metrics and approximate path data, an Agent-to-agent test provides bi-directional network metrics. This includes detailed throughput, packet loss, latency, and jitter measurements in both directions, which is critical for identifying asymmetric routing or link saturation that might affect application performance.

Alternative options are less suitable for this specific internal monitoring requirement:

Cloud Agents (Options A and B): These are located in global ISP data centers and cannot be placed inside the organization's private data center or branch LAN to measure internal path characteristics.

HTTP/DNS Server tests (Options B, D, and E): These target specific services but do not provide the granular, bi-directional network-layer performance data (like throughput or path visualization) that an Agent-to-agent test delivers.

Agent-to-server (Option E): While useful for reaching a target that cannot host an agent, it loses the benefit of the destination agent's ability to measure the "return" path independently.

By deploying Enterprise Agents at each site and configuring Agent-to-agent tests, the engineer gains the highest level of visibility into the health of the private WAN or SD-WAN fabric connecting the branches to the data center.

Question: 39

Refer to the exhibit.

Monitor Application Step 1 of 3 - Select an application

Q. Filter by application name

Popular Applications

10 tests

Test Name

- ☐ Slack - Tenant
- ☐ Slack - App
- ☐ Slack - API
- ☐ Webex - Dynamic
- ☐ Webex - HTTP Server
- ☐ Salesforce - Lightning
- ☐ Vonage - Scheduled - HTTP Server
- ☐ Teams Transport Relay ICMP
- ☐ Teams Transport Relay TCP
- ☐ https://app.thousandeyes.com

Custom Application
Created by ThousandEyes

Google Suite
A template for monitoring your Google Suite
Created by ThousandEyes

Microsoft 365
A template for monitoring your Microsoft 365 suite
Created by ThousandEyes
Certified

Microsoft Teams
A template for monitoring Microsoft Teams
Created by ThousandEyes

ServiceNow
A template for monitoring your ServiceNow instance
Created by ThousandEyes

Webex
A template for monitoring Webex
Created by ThousandEyes

Salesforce
A template for monitoring Salesforce
Created by ThousandEyes

Slack
A template for monitoring your Slack instance
Created by ThousandEyes

Zoom
A template for monitoring Zoom
Created by ThousandEyes

An engineer is tasked with configuring a new test to monitor a web application from the employee's point of view. What two actions should be taken to fulfill the requirement?

- A. Create a new custom application monitor
- B. Create a new google suite monitor
- C. Add a new scheduled test to the monitor
- D. Add a new dynamic test to the monitor
- E. Add a new test template

Answer: A, C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, monitoring the digital experience of internal employees requires leveraging ThousandEyes Endpoint Agents to

simulate or record traffic from the user's specific vantage point. When an engineer needs to monitor a specific, potentially internal or non-standard web application, they must utilize the Custom Application workflow within the Endpoint Experience settings.

According to the ENNA implementation standards, the first necessary action is to Create a new custom application monitor (Option A). In the ThousandEyes portal under Endpoint Experience > Test Settings > Synthetic Tests, the "Monitor Application" button provides access to pre-defined templates for popular services like Microsoft 365 or Webex. However, for a unique enterprise application, the engineer must select "Custom Application" to define the application's identity, including its name and the relevant domain or URL.

The second required action is to Add a new scheduled test to the monitor (Option C). Endpoint Agents perform Scheduled Tests at regular, predefined intervals—such as every 5 or 10 minutes—to proactively check the application's availability, response time, and network path health without requiring user interaction. By adding a scheduled HTTP Server or Network test to the custom monitor, the engineer ensures a consistent baseline of the application's performance as seen from the employee's machine.

Reviewing the incorrect options:

Google Suite monitor (Option B): This is a specific template for Google Workspace and is not suitable for a custom application.

Dynamic Tests (Option D): These are specifically designed for collaboration tools like Zoom or Microsoft Teams to capture ad-hoc sessions; they are generally not used for baselining a standard custom web application.

Test Template (Option E): While all monitors are based on templates, "adding a new template" is not a configuration action but rather a result of the monitoring setup.

Question: 40

You want to create an endpoint label that automatically includes all Endpoint Agents connected to your corporate network. If your agents are named using the format agentname-network, what filter would you use in the hostname field to achieve this?

- A. -corporate
- B. agentname-
- C. agent*corporate
- D. There is no wildcard configuration available

Answer: A

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework, managing a large-scale deployment of Endpoint Agents requires efficient categorization through Agent Labels. Agent labels are used to group agents for targeted test assignment and data filtering based on specific attributes like location, OS, or naming conventions.

ThousandEyes supports the use of wildcards (specifically the * symbol) within label filters to allow for dynamic and scalable grouping. In this scenario, the organization has adopted a naming convention of agentname-network. To capture all agents connected to the "corporate" network regardless of the unique agentname prefix, the engineer should use the filter *-corporate (Option A) in the hostname or agent name field. The asterisk functions as a placeholder for any string of characters, ensuring that any agent ending with "-corporate" is automatically added to the label as soon as it registers with the ThousandEyes platform.

Other options are incorrect for the following reasons:

agentname-* (Option B): This would only match agents that literally start with the string "agentname-", which is a placeholder and not a representative filter for a group of uniquely named agents.

agent*corporate (Option C): This filter is overly broad and lacks the specific hyphen delimiter used in the organization's naming convention, potentially leading to the inclusion of unintended agents.

No wildcard (Option D): Wildcards are a documented and essential feature of the ThousandEyes label system to facilitate automated management in enterprise environments.

By implementing wildcard-based labels, the network administrator ensures that new agents are seamlessly integrated into the assurance strategy without the need for manual manual intervention.

Question: 41

What type of endpoint agent test will gather browser activity?

- A. Scheduled tests
- B. Dynamic tests
- C. Real user tests
- D. Network Access tests

Answer: C

Explanation:

The Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum identifies Real User Tests (RUM) as the primary method for capturing actual browser-based interactions. While synthetic tests provide proactive baselines, Real user tests (Option C) provide visibility into the "lived experience" of an employee as they navigate corporate websites and SaaS platforms.

Real user tests utilize a lightweight browser extension (available for Chrome, Edge, and Safari) that is part of the Endpoint Agent deployment. When a user visits a domain defined in a "Monitored Domain Set," the extension automatically begins recording performance metrics. This includes the time spent on each page, DNS resolution time, and a full browser waterfall of every object (images, scripts, CSS) loaded during the session. Crucially, it calculates an "Experience Score" for each session, allowing IT teams to identify if a user's frustration is caused by slow page rendering or underlying network issues.

In contrast:

Scheduled tests (Option A) and Dynamic tests (Option B) are forms of Synthetic Monitoring. They simulate traffic using a headless browser or simple network probes to check availability but do not record the actual clicks or browsing activity of the real human user.

Network Access tests (Option D) is not a standard test category in the ThousandEyes endpoint ecosystem; rather, network connectivity is a layer monitored within the other test types.

Therefore, Real user tests are the definitive tool for organizations needing to analyze actual employee browser behavior and the resulting application performance.

Question: 42

You want to monitor Microsoft Teams using ThousandEyes endpoint agents. Which tests are available for this type of application monitoring?

- A. Scheduled tests
- B. Dynamic tests
- C. Scheduled, dynamic and real user tests
- D. Scheduled and dynamic tests

Answer: D

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) best practices for collaboration monitoring, Microsoft Teams (as well as Webex and Zoom) is monitored using a combination of Scheduled and Dynamic tests (Option D).

Scheduled Tests: These are used to proactively monitor the reachability and performance of the Microsoft Teams infrastructure. The Teams template typically includes scheduled HTTP Server tests to `teams.microsoft.com` and `login.microsoftonline.com`, along with network-layer tests to the Teams Transport Relay (e.g., `worldaz.tr.teams.microsoft.com`). These provide a continuous baseline of connectivity regardless of whether a user is currently in a call.

Dynamic Tests: These are a unique feature for collaboration monitoring. When an Endpoint Agent detects that the Teams application has initiated a real-time media session (audio or video call), it automatically triggers a Dynamic Test to the specific IP and port being used for that call. This allows IT teams to see the exact network path—including hop-by-hop latency and loss—of the actual call traffic while the session is active.

Real User Tests (Option C) are generally not used for Teams monitoring in the same way they are for web-based SaaS apps. Because most users utilize the Teams desktop client rather than the browser, the RUM browser extension cannot capture the rich telemetry of the desktop application's internal signaling and media streams. Therefore, the specialized "Automated Session Testing" provided by Dynamic Tests, combined with the proactive health checks of Scheduled Tests, constitutes the complete assurance strategy for Microsoft Teams.

Question: 43

An engineer needs to create a test to execute a user's workflow where the user has to log in to OneDrive and download a file. The test has to implement a retry mechanism. The engineer has limited scripting experience. What are the actions that the engineer needs to take?

- A. Create the script from the Office365 > OneDrive - Download File template
- B. Install the ThousandEyes Recorder IDE and record the user flow
- C. Check the transaction-scripting-examples repository for sample scripts
- D. All of the above

Answer: D

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, transaction monitoring is essential for validating complex, multi-step user workflows. When an engineer with limited scripting experience needs to monitor a OneDrive file download process, the ThousandEyes ecosystem provides several tools to simplify the creation of robust Transaction tests.

The correct approach is to leverage all available resources, making Option D the definitive answer. First, the ThousandEyes Recorder IDE (Option B) is a critical tool for non-scripters. It allows an engineer to perform the actual workflow—navigating to OneDrive, logging in, and initiating the download—in a browser environment on their local machine while the tool records every click and keyboard entry. The Recorder automatically generates the corresponding JavaScript code using the ThousandEyes transaction library. Second, the

platform provides pre-built script templates (Option A), such as those for Office 365 and OneDrive, which include baseline logic and best practices for these specific services. Third, the official transaction-scripting-examples repository on GitHub (Option C) is a maintained source of code snippets. This is particularly useful for implementing advanced logic, such as a retry mechanism, which ensures that the test does not report a "failure" due to a transient network hiccup but instead attempts the action again before triggering an alert.

As shown in the provided exhibit, a typical script uses import statements from @thousandeyes and selenium-webdriver to control the browser. By combining the Recorder for the basic flow, a Template for service-specific nuances, and Sample Scripts for logic enhancements like retries, the engineer can deploy a highly reliable assurance test without deep coding expertise. Therefore, all three actions are highly recommended and valid within the ENNA implementation framework.

Question: 44

You're responsible for monitoring the performance of a company e-commerce website. You're considering using ThousandEyes Synthetic Web Tests. Which of the following functionalities of ThousandEyes Synthetic Web Tests would be MOST beneficial for monitoring the e-commerce checkout process?

- A. HTTP server monitoring
- B. Transaction monitoring
- C. DNS monitoring
- D. Routing visibility

Answer: B

Explanation:

According to the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework, application health is not just about reachability, but about the functionality of business-critical paths. For an e-commerce checkout process, which involves multiple interactive steps such as "Add to Cart," "Enter Shipping Info," and "Finalize Payment," Transaction monitoring (Option B) is the most beneficial functionality.

Transaction tests provide a high-level view of application performance by executing a script that simulates a real user's interaction with the site's various components. Unlike HTTP server monitoring (Option A), which merely verifies that a single URL returns a specific status code (like 200 OK), Transaction monitoring validates that the underlying application logic and database transactions are working as expected across the entire session. If a user can reach the homepage but the "Checkout" button fails due to a slow backend API call or a JavaScript error, only a Transaction test will identify the failure and provide a granular waterfall chart to show exactly where the process broke down.

While DNS monitoring (Option C) and Routing visibility (Option D) are crucial for identifying foundational

connectivity issues, they are insufficient for monitoring complex transactional workflows. A DNS failure would prevent access to the site entirely, but a "broken" checkout process often occurs while the network and DNS appear perfectly healthy. Transaction monitoring bridges this gap by providing "layer 7" visibility into the end-to-end user experience, making it the definitive choice for e-commerce performance assurance.

Question: 45

An engineer needs to create a test that requires authentication configuration to monitor an API. The test must send a POST request with client credentials parameters to get a token. The token then needs to be sent out on a GET request to be authorized to get the resource.

Scheme

None

Basic

NTLM

Kerberos

OAuth

HTTP Version

Prefer HTTP/2

HTTP/1.1 only

Request Method

GET

POST

POST Body

Follow Redirects



Enable

User Agent

Default

Custom Headers

Root Request

Enter HTTP headers...



What must be done to meet the requirements? (Choose two)

- A. Configure the HTTP server test to use Basic authentication for client credentials
- B. Configure the HTTP server test to use NTLM authentication for client credentials
- C. Configure the HTTP server test to use OAuth authentication for client credentials
- D. Parameters are not supported by HTTP server OAuth authentication; use a Transaction script instead
- E. Parameters are not supported by HTTP server OAuth authentication; use an API test instead

Answer: D, E

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, monitoring modern

APIs often requires handling complex authentication flows, such as OAuth 2.0 with specific client credential parameters. The provided exhibit illustrates the HTTP Authentication options available within a standard HTTP Server test: None, Basic, NTLM, Kerberos, and OAuth.

According to the ENNA implementation standards, while the HTTP Server test type supports OAuth (Option C), its native implementation is limited. Specifically, it is designed to use a pre-existing token or a simple token refresh flow; it does not support the injection of custom parameters in the initial POST request required to obtain a token in many enterprise client-credential scenarios. Basic and NTLM (Options A and B) are legacy protocols that rely on simple username/password headers and cannot facilitate the multi-step token exchange process described.

To fulfill the requirement of a two-step flow (POST for token, followed by GET for resource), engineers must use more flexible test types:

Transaction scripts (Option D): These allow the engineer to write custom JavaScript using the ThousandEyes transaction library to programmatically handle the POST request, parse the resulting JSON token, and then pass that token into the subsequent GET request's header.

API tests (Option E): These are purpose-built for API monitoring and natively support the definition of variables and multi-step requests where the output of one call (the token) serves as the input for the next.

By utilizing these advanced test types, the engineer can successfully navigate complex authentication requirements that the standard HTTP Server test cannot accommodate.

Question: 46

You are tasked with creating a ThousandEyes transaction test to monitor the login process of a web application that uses SAML-based SSO with MFA. The MFA step involves a one-time password (OTP) generated by a mobile app. How can you configure the ThousandEyes test to successfully navigate this login process?

- A. Configure the test to automatically enter the OTP from the mobile app.
- B. Manually enter the OTP in the test configuration each time it changes.
- C. Use a ThousandEyes webhook to retrieve the OTP from a third-party service.
- D. Exclude the MFA step from the transaction test and focus only on the SAML login.

Answer: D

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework, a common architectural hurdle is monitoring applications protected by Multi-Factor Authentication (MFA). ThousandEyes transaction tests utilize automated browser sessions to simulate user behavior, but they face inherent limitations when interacting with "out-of-band" security mechanisms.

Specifically, ThousandEyes agents cannot natively interact with external hardware tokens, biometric prompts, or mobile-app-based OTP generators (Option A). Since an OTP is dynamic and timesensitive, manually entering it into a static test configuration (Option B) is impossible for an automated, recurring test. While some complex workhooks (Option C) might theoretically interface with a virtual MFA service, this introduces significant security risks and architectural complexity that is generally discouraged in a standard assurance design.

The verified and most practical approach is to exclude the MFA step from the transaction test and focus only on the SAML login (Option D). For monitoring purposes, IT teams often create a "synthetic user" account within the Identity Provider (IdP) that is specifically exempted from MFA policies when originating from known ThousandEyes Enterprise Agent IP addresses. This allows the transaction script to validate the availability and performance of the SAML-based SSO redirect, the credential challenge, and the final application landing page. This strategy ensures that the network and application health can be baselined without the test being blocked by a security gate it was never intended to pass.

Question: 47

You are investigating intermittent failures in a ThousandEyes transaction test targeting a web application that uses Basic Authentication. The failures occur randomly across different agents and times of day. What steps would you take to troubleshoot and resolve the issue? (Select all that apply)

- A. Disable Basic Authentication in the test configuration to isolate the problem.
- B. Verify the correctness of credentials by manually logging into the application from different locations.
- C. Analyze the ThousandEyes waterfall charts and HTTP response codes to identify potential bottlenecks or errors.
- D. Contact the web application vendor to report the issue and inquire about possible server-side problems.

Answer: B, C, D

Explanation:

Troubleshooting intermittent performance issues is a core component of the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum. When a transaction test using Basic Authentication fails randomly, the engineer must employ a multi-layered diagnostic approach to determine if the fault lies with the credentials, the network path, or the application server.

The first step is to verify the correctness of credentials (Option B). Because Basic Authentication encodes credentials in the header, any change in user permissions or account lockouts will trigger immediate failures. Performing manual logins from various geographical locations helps rule out location-based access control lists (ACLs) or regional IdP sync issues.


The most data-rich diagnostic step is to analyze the ThousandEyes waterfall charts and HTTP response codes (Option

C). The waterfall view allows the engineer to see if the failure happens during the initial 401 Unauthorized challenge or after the credentials are sent. If the charts show high "Wait Time" or 5xx errors, the issue is likely server-side latency or instability. If the "Connect Time" is high, the problem may be network-layer congestion.

Finally, if the telemetry indicates that the network path is healthy but the server is intermittently returning errors or timing out, the engineer should contact the web application vendor (Option D). Providing the vendor with the specific ThousandEyes "Share Link" allows them to see the exact same packet-level and browser-level data, proving that the issue is not with the client's network but with the server-side infrastructure. Disabling authentication (Option A) is not a valid troubleshooting step for a test designed specifically to monitor an authenticated workflow.

Question: 48

Users at a remote corporate site (identified as s30 in the exhibit) are experiencing issues with a critical Enterprise Application hosted in the Data Center. The site connects to the central campus through an MPLS network.



The following exhibits show the network status before and after the issue began. Based on the information presented, what is the most likely cause of the problem and what actions would you take next as a Network Operations Engineer?

A. Escalate to the transmission media team and have the optic fiber between 10.84.30.1 and

10.87.16.53 checked.


- B. Review the bandwidth utilization at this site.
- C. Reach out to the team that owns the Enterprise Application and have the server reviewed.
- D. Check the routing tables on the MPLS network devices for any recent changes.



Answer: A


Explanation:

Question: 49

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








- 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)
- Verify the routing changes on device 10.87.7.51
- Verify the docker host 10.84.50.53 and ensure the agent container is running.
- Analyze the jitter and latency trends on the affected voice paths to identify potential network congestion.

Answer: C

Explanation:

Question: 50

Refer to the exhibits.



VIEWS

CTF • Run Now ▼ Add a filter *

1^1 Save and Share

Dynamic Tests

Scheduled Tests

HTTP Server

Network

Real User Tests

Local Networks

Path Visualization Map Table

Showing: VPN Overlay - Add 8 filter -

Grouping: Agents by Network Interfaces by Network

Highlighting: Forwarding Loss > 10 % (0 nodes) • Link Delay > 100 ms (0 links) •

Selecting: Click a node or link

Showing data from Sun, May 5 23:26 ■ 23:27 GMT+2

! ▶ Latest-!

24h 7d 14d Metric Loss

■ Average Loss | Number of Agents



Highlight nodes that match all / any Search on Network, Country, IP address.

Prefix, or W

Node Labels by None

4

Node SA (213.238.64.0/18)

1

4

NETIASPOT-5GHz-p60D

4

Netla SA (AS-12741)

1

Unknown

1

172.67.68.44 (172.67.68.44)

0 hops

VIEWS

CTF • Run Now ▼ Add a filter *

Save and Share

Dynamic Tests

Scheduled Tests

HTTP Server

Network

Real User Tests

Local Networks

Showing: VPN Overlay -

Path Visualization

Add a filter *

Grouping: Agents by Network ▼ Interfaces by Network ▼

Highlighting: Forwarding Loss > 10 % (0 nodes) • Link Delay > 100 ms (0 links) •

Selecting: Click a node or link

24h 7d 14d Metric Loss

■ Average Loss ■ Number of Agents



Showing data from Sun, May 5 23:27-23:28 GMT+2

Highlight nodes that match all / any

Search on Network, Country, IP address

Node Labels by None

4

No Traces

The endpoint has the following IP credentials:

192.168.100.9/24, DNS: 8.8.8.8, 8.8.4.4, GW: 192.168.100.1

Based on the views presented in the exhibits, what led to the error occurring on Sun, May 5 23:27 GMT +2?

- A. The test target stopped responding.
- B. The FQDN of the test target is non-existent.
- C. The DNS servers assigned to the endpoint are unreachable.
- D. The DNS settings on the endpoint are incorrect.

Answer: C

Explanation:

Question: 51

The Endpoint stopped appearing online after it was moved to another network. The customer reviewed the endpoint logs but did not identify anything suspicious. The customer also confirmed that the endpoint was online on the old network, and the new network is fully operational. Other endpoints that were moved to the new network are also online. Since the new network is small, the admin is using static IP assignment. What is the best way to bring the endpoint online?

```
2024-05-07 13:20:12.452 DEBUG [3068.4 3 5 2] agent.OefaultPersisTentconnector^OefaultPercistentConnection.cpp/1GG ■ RECONNECT 14 currently scheduled -or running, ignoring schedule attempt for CONNECT
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService^OefaultAgentSettingsService.cpp:496 Checkin failed with error code: 1
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService^DefaultAgentSettingsService.cpp:5 58 ■ Calling connect on persistent Connection
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService^OefaultAgentSettingsService.cpp:576 - Will consider polling For scheduled tests in 30s
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService^DefaultAgentSettingsService.cpp:432 - Scheduling checkin in lam
2024-05-07 13:28:14.468 DEBUG [3068.4 35 2] agent.DefaultPersistentConnection^OefaultPersistentConnection.cpp:198 - Running pre-schedule checks for CONNECT, isRetry=false, prev=CONNECT, current=RECONNECT, next=NONE
2024-05-07 13:20:14.468 DEBUG [3068.4352] agent.OefaultPersistentConnection^DefaultPersistentConnection.cpp:206 ■ RECONNECT is currently scheduled or running, ignoring schedule attempt for CONNECT
2024-05-07 13:20:24.753 DEBUG [3068.8240] system.EtwTraceConsumer^DefaultEtwTraceConsumer,app:108 Flushed ETW trace buffers to file, path="C:\Vrogr\mData\ThousandEyes\Endpoint Agent\etwTraces\ThousandEyes Endpoint Agent.etl",
2024-05-07 13:20:24.755 size=320kB
2024-05-07 13:20:24.812 DEBUG [3068.8240] net.SocketMonitoringgWinSocketMonitoringService.cpp:335 - Processed ETW event batch; published=0, stale=213, deduped=0, duplicatesDueToCaal=celnterv=0, skipped=0!
2024-05-07 13:20:24.812 DEBUG [3068.7712] scheduler.task.StatscwwittergUsageStatsCommitter.cpp:43 - Starting UsageStatsCommitter task
2024-05-07 13:20:24.813 DEBUG [3068.7712] tests.scheduledtestsmetric^OefaultScheduledTestsMetricsService.cpp:54 - End metric tracking period at: 2024-May-07 13:20:24.812000
2024-05-07 13:20:24.813 DEBUG [3068.7712] scheduler.task.statscwwittergUsageStatsCommitter.cpp:95 - Updated usage stats object with 69 entries.
2024-05-07 13:20:28.498 DEBUG [3068.2148] scheduler.lmplegOofaultTaskScheduler,app:173 - Scheduled task: to:agent:UcageStat&Committer\0x1, delay=60000 ns
2024-05-07 13:20:28.498 DEBUG [3068.2148] net.NinHttpClient^NinHttpClient.cpp:2115 - Request to: w://cl.eb.thousandeyes.com/rela/connect timed out: 12802: The operation timed out
2024-05-07 13:20:28.498 DEBUG [3068.8240] net.WinHttpClient^NinHttpClient.cpp:2115 Request to: https://cl.eb.thousandeyes.com/VstaU15J5On timed out: 12862: The operation timed out
2024-05-07 13:20:28.498 DEBUG [3068.4352] net.DefaultWebsocketStompConnection^OefaultWebsocketStompConnection.cpp:135 - Websocket connection wait: 42108ms
2024-05-07 13:20:28.498 ERROR [3068.4352] net.DefaultWebsocketStompConnection^OefaultWebsocketStompConnection.cpp:148 asynOpen (Websocket): Input/output error [getfentczs]
2024-05-07 13:20:28.498 ERROR [3068.4352] agent.OefaultPersistentConnection^OefaultPersistentConnection.cpp:317 ■ asynOpen: Input/output error [generic-5]
2024-05-07 13:20:28.498 DEBUG [3068.4352] agent.DefaultPersistentConnection^OefaultPersistentConnection.cpp:485 Running post-run checks for RECONNECT, next=NONE, retry=trere
2024-05-07 13:20:28.498 DEBUG [3068.4352] agent.DefaultPersistentConnection^DefaultPersistentConnection.cpp:198 ■ Running pre-schedule checks for RECONNECT, isRetry=true, prev=RECONNECT, current=NONE, next=NONE
2024-05-07 13:20:28.498 DEBUG [3068.4 352] agent.OefaultPersi=entCvncnector^Oe-faultPer51StentConnectlcm. cop: 386 - Set schedule delay=158&ns, jit ter *32 3ms, nextDelay=21Mms
2024-05-07 13:20:28.771 INFO [3068.4352] agent.DefaultPersistentConnection^DefaultPersistentConnection.cpp:226 - Scheduling RECONNECT after 1823ms
2024-05-07 13:20:28.771 INFO [3068.8072] agent.services.CheckinService^CheckinSery/Oe,app:422 ■ Checking in with ThousandEyes backend. Last successful Check-in: 2024-May-07 13:16:04.698000. Last attempt! 2024-May-07 13:18:24.548000
2024-05-07 13:20:30.327 DEBUG [3068.8072] api.usp^UserSpaceProxyClient.cpp:50 ■ Relaying HTTP request https://cl.eb.thousandeyes.com/eyebrow/enterprise/agent/checkin to user space proxy
2024-05-07 13:20:30.327 DEBUG [3068.4352] agent.DefaultPersistentConnection^DefaultPersistentConnection.cpp:244 ■ Running RECONNECT now
2024-05-07 13:20:30.327 DEBUG [3068.4352] net.DefaultWebsocketStompConnection^OefaultWebsocketStompConnection.cpp:308 STOMP connection closed successfully
2024-05-07 13:20:30.327 DEBUG [3068.4352] net.DefaultWebsocketStompConnection^DefaultWebsocketStompConnection.cpp:327 ■ Websocket connection closed successfully
```

- A. It may be an issue with the lack of space in the new network. The endpoint should be moved back to the old network.
- B. The endpoint agent should be reinstalled to come online. This always helps.
- C. The endpoint will automatically come online in 10-15 minutes, no action is needed.
- D. Endpoint IP settings must be checked along with connectivity to c1.eb.thousandeyes.com.

Answer: D

Explanation:

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

troubleshooting endpoint connectivity requires a systematic analysis of local device logs and network environment changes. According to the log message provided in Exhibit there was a specific timeout error when the agent attempted to reach the ThousandEyes backend. The logs explicitly state: Request to:

wss://c1.eb.thousandeyes.com/relay/connect timed out: 12002: The operation timed out.

Since the endpoint was relocated to a new network utilizing static IP assignment, and other endpoints on the same network are functioning correctly, the issue is highly localized to this specific device's configuration. In a static IP environment, human error during the manual entry of the IP address, subnet mask, default gateway, or DNS server is a common root cause of connectivity failure. If the gateway is incorrect, the agent cannot reach the internet; if the DNS is incorrect, it cannot resolve c1.eb.thousandeyes.com. Therefore, the best way to bring the endpoint online is to verify the accuracy of the IP settings and test direct connectivity to the ThousandEyes relay URL.

The alternative options do not address the technical evidence found in the logs:

Option A: Moving the device back is a regressive step that does not resolve the need for the device to function in its new location.

Option B: Reinstalling software is a "brute force" method that rarely fixes underlying network-layer misconfigurations like an incorrect static IP.

Option C: A timeout error (12002) indicates a persistent connectivity blockage that will not resolve itself without administrative intervention.

Checking the static IP credentials ensures the device has a valid path to the internet to establish its secure websocket (WSS) connection to the ThousandEyes platform.

Question: 52

Review the exhibits. Based on the evidence, which action is most likely to solve the issue?

Map fable Waterfall Gamma ' Dependent Applications

Pago https://vritpoin.0/basB-ajtr, Liseupasa

Screenshot

Waterfall

Overview

Search by URL or Domain... Component Type

Component	Response Code	Domain	Provider	Size(kB)	Duration (ms)	Waterfall T
Opass	401 [Headers]	httpbin...	Amazon.com, I...	0	590	

[Request Headers](#) [Response Headers](#)
[Request Headers](#) [Response Headers](#)

:authority: httpbin.org method: GET
:path: /basic-auth/user/pass
:scheme: https
accept-encoding: gzip, deflate, br sec-fetch-dest: document sec-fetch-mode: navigate sec-
fetch-site: none sec-fetch-user: ?1 upgrade-insecure-requests: 1
user-agent: Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like
Gecko) Chrome/118.0.5993.70 Safari/537.36 x-thousandeyes-agent: yes

access-control-allow-credentials: true access-
control-allow-origin: * content-Length: 0
date: Wed, 24 Apr 2024 19:28:12 GMT
server: gunicorn/19.9.0
wm-authenticate: Basic realm="Fake Realm"


- A. Modify the firewall rules to allow connections to the target domain
- B. Modify the authentication credentials
- C. Change the HTTP request method to PATCH
- D. Modify the target URL to an available API endpoint

Answer: B

Explanation:

Question: 53

Review the exhibits. Based on the evidence, what seems to be the underlying issue?



Request Headers Response Headers

authority: httpbin.org :method: GET :path: /mypage :scheme: https accept-encoding: gzip, deflate, br
 sec-fetch-dest: document sec-fetch-mode: navigate sec-fetch-site: none sec-fetch-user: ?1 upgrade-
 insecure-requests: 1
 user-agent: Mozilla/5.0 (XU; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko)
 Chrome/118.0.5993.70 Safari/537.36 x-thousandeyes-agent: yes

Request Headers Response Headers

access-control-allow-credentials: true access-
 control-allow-origin: * content-length: 0
 content-type: text/html; charset=utf-8 date:
 Wed, 24 Apr 2024 19:48:03 GMT location:
 /mypage
 server: gunicorn/19.9.0

- A. There is a network connectivity problem preventing us from reaching the target URL
- B. One of the DOM elements cannot be found in the server
- C. The request timed out waiting for the server to respond
- D. There is a misconfiguration in the application server


Answer: D

Explanation:

Question: 54

Carefully review the exhibits. Which detail indicates the network issue might be caused by a BGP Hijack?





- A. Availability Drop
- B. AS 16509 change to AS 10297
- C. HTTP Server response delay
- D. Packet Loss

Answer: B

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, identifying routing anomalies such as BGP Hijacking is a critical aspect of external network assurance. A BGP Hijack occurs when an Autonomous System (AS) illegitimately announces a prefix it does not own, causing traffic intended for the rightful owner to be misdirected.

By contrasting Exhibit 3.4-1 (Before the Outage) and Exhibit 3.4-2 (During the Outage), a definitive change in the routing path is observed. In the baseline view (Exhibit 3.4-1), traffic from the Los Angeles agent reaches the destination network 54.239.104.0/23, which is correctly associated with Amazon.com, Inc. (AS 16509).

However, during the outage (Exhibit 3.4-2), the path visualization shows the traffic being redirected into a different network: eNET Inc. (AS 10297). At this new location, the traffic encounters 100% Forwarding Loss at the node ten7-8.core-2.xlhost.com.

concrete indicator of a BGP Hijack. While symptoms such as Availability Drop (Option A), Response Delay (Option C), and Packet Loss (Option D) are clearly visible in the ThousandEyes telemetry, these symptoms are generic and could be caused by various other issues like physical link failures or congestion. The shift in the AS path—specifically to an AS that does not legitimately host the target's IP prefix—provides the forensic evidence needed to identify the root cause as a routing takeover. This visualization allows network administrators to bypass internal troubleshooting and immediately focus on external remediation with upstream providers.

Question: 55

Considering the observed network behavior and the information in the exhibits, which action would be the most appropriate next step for the network administrator to take?

- A. Contact the internal network team to investigate potential misconfigurations on the local routers
- B. Reach out to the Internet Service Provider (ISP) to report the suspected BGP hijacking incident
- C. Implement traffic filtering rules on the firewall to block traffic originating from AS 10297
- D. Restart the DNS server to refresh its cache and potentially resolve the observed issue

Answer: B

Explanation:

The Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework emphasizes that the goal of internet intelligence is to enable rapid and accurate escalation to the party responsible for a service degradation. Based on the evidence of a BGP Hijack identified in the previous question, the issue is occurring entirely within the public internet ecosystem.

The most appropriate next step is to reach out to the Internet Service Provider (ISP) to report the suspected BGP hijacking incident (Option B). Since the traffic is being misdirected by an external Autonomous System (AS 10297) before it reaches the intended destination (AS 16509), the fix must occur at the routing policy level of the major transit providers. The network administrator should provide the ISP with the ThousandEyes "Share Link" or screenshots showing the path change and the unauthorized AS announcement, as this data serves as proof to accelerate the ISP's mitigation efforts, such as implementing prefix filters or contacting the offending network.

Other options are ineffective for this specific scenario:

Option A: The path visualization shows that traffic is successfully leaving the local network and reaching the public internet; the problem is many hops away from the internal routers.

Option C: Blocking traffic from AS 10297 does not solve the problem of your traffic being attracted to it. The hijack affects how the rest of the world (including your ISP) sees the route to your destination.

Option D: DNS is not the issue; the agent successfully resolved the hostname to the correct IP, but the BGP layer misdirected the packets at the routing level.

By identifying the issue as an external routing event, the administrator avoids wasting internal resources and directly triggers the necessary external remediation.

Question: 56

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

The screenshot shows the configuration page for an alert rule titled "AS origin AS 1234" under the "BGP" category. The "Settings" tab is active, showing the following configuration:

- Rule Name:** AS origin AS 1234
- Tests:** 1 of 29 test(s) selected
- Prefix Length:** IPv4 /10 - /32 And IPv6 /32 - /128
- Monitors:** All monitors
- Alert on covered prefix data:** ☒
- Severity:** Info, Minor, Major, Critical (Critical is selected)

The "ALERT CONDITIONS" section shows the rule is configured to trigger when "All conditions are met by any of 1 monitor 1 of 1 time in a row". The condition is defined as:

- BGP Origin:** meets all of:
 - ASN:** not in 1234

There is a "Collapse All" link at the bottom of the conditions list.

- 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: 57

The alert shown in the exhibit is designed to detect which of the following network security issues?

- A. Route poisoning
- B. DNS poisoning
- C. BGP hijacking
- D. DNS hijacking

Answer: C

Explanation:

As shown in Exhibit 4.1-1 (image_6e19fb.png), the alert rule is configured for the BGP layer with the name "AS origin AS 1234". The critical logic is defined in the Alert Conditions section, which specifies: BGP Origin meets all of: ASN not in 1234.

This specific configuration is the standard method used in ThousandEyes to detect BGP hijacking (Option C). BGP hijacking occurs when an unauthorized Autonomous System (AS) announces a network prefix that it does not own, effectively stealing the traffic meant for the legitimate owner. By setting an alert that triggers when a monitored prefix is seen originating from an Autonomous System Number (ASN) that is not the expected authorized ASN (in this case, AS 1234), ThousandEyes provides immediate notification of a routing takeover.

Identifying this via BGP Origin alerts is superior to simple connectivity checks because it catches the issue at the routing control plane.

DNS Issues (Options B and D): These involve the translation of names to IPs and would be detected via DNS-specific tests rather than BGP Origin checks.

Route Poisoning (Option A): This is a generic term often referring to the manipulation of distancevector protocol metrics to make a route unreachable; while related to routing, BGP Origin checks are the specific tool for detecting external prefix hijacks.

By monitoring the ASN origin, network engineers can quickly distinguish between a local ISP outage and a global security event where their traffic is being misdirected to an illegitimate network.

Question: 58

Refer to the exhibit.

Add New Alert Rule

Alert Type **Web** **HTTP Server**

Rule Name **ENNA-HTTP-Alert**

Settings Notifications

Test
s
Agent
s
Severity
y

0 of 29 test(s)
selected
All
agents

Info	Minor	Major	Critical
------	-------	-------	----------

ALERT CONDITIONS

All v conditions are met by any 1 agent 1 v time in a row



A network engineer is tasked with configuring an alert that will trigger if the HTTP server responds with a server error. What alert conditions should be configured to meet the specified requirements?

- A. Error type is any
- B. Wait Time is Dynamic (New) with Medium sensitivity
- C. Response Time $\geq$ Static 500ms
- D. Response Code is server error(5XX)12

Answer: D34

Explanation:

56

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) framework, configuring effective alert rules is critical for distinguishing between standard network noise and actionable

application-layer failures. For Web - HTTP Server tests, ThousandEyes allows engineers to monitor both network-level metrics (like Connect time) and application-level indicators (like HTTP response codes).

The requirement is to trigger an alert specifically when the HTTP server responds with a server error. In the HTTP protocol, server errors are categorized as the 5XX series of status codes (e.g., 500 Internal Server Error, 503 Service Unavailable, 504 Gateway Timeout). To meet this requirement, the engineer must configure a location alert condition where the Metric is set to Response Code and the condition value is server error(5XX) (Option D).

Reviewing the other options:

Error type is any (Option A): While this would capture server errors, it would also trigger for 4XX client errors (like 404 Not Found) and network-layer timeouts, making it too broad for a specific "server error" requirement.

Wait Time is Dynamic (Option B): This monitors the time-to-first-byte using statistical baselining. While high wait times often precede 5XX errors, this condition only alerts on latency, not on the actual error code itself.

Response Time (Option C): Similar to wait time, this monitors performance speed rather than the logical success or failure of the server's response.

By specifically selecting Response Code: server error(5XX), the engineer ensures that the operations team is only notified when the application backend is experiencing a functional failure, rather than just a slow response or a client-side misconfiguration.

Question: 59

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?

- 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 $\geq$ %
- D. Scheduled Tests > Endpoint End-to-End (server), Memory load $\geq$ %

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 $\geq$ % (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 ($\geq$ %). 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.

Question: 60

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

- 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: 61

Analyzing the IT operations dashboard, which agent has a better HTTP Connect Time?

- A. San Jose CA (AT&T)
- B. Mexico City Mexico (TelMex)

Answer: B

Explanation:

In the IT Operations Dashboard, performance metrics are displayed per agent to facilitate granular troubleshooting. By reviewing the latest metrics table, the Mexico City Mexico (TelMex) agent displays a Connect time of 0.51 ms, which is significantly lower (and therefore "better") than the corresponding time for the San Jose CA agent. Connect time measures the duration for the TCP three-way handshake; a lower value indicates faster network-layer establishment.

Question: 62

In the IT operations dashboard, what is the alert trigger reason?

- A. Page Load Packet Loss
- B. Network jitter
- C. Network packet loss
- D. Page Load Latency

Answer: C

Explanation:

The IT Operations Dashboard includes a section at the beginning that explicitly displays active alert rules and their status. According to the dashboard configuration, the reason for the current alert trigger is Network packet loss. This indicates that the underlying network path for the application is experiencing packet drops exceeding the defined threshold, even if the application layer remains partially functional.

Question: 63

A network monitoring engineer is tasked with creating a widget that displays the average packet loss from an agent installed as a Linux package. What is the data source and measure that should be selected?

- A. Endpoint Agents and Median
- B. Cloud & Enterprise Agents and Mean
- C. Routing and Standard Deviation
- D. Devices and nth Percentile

Answer: B

Explanation:

In Designing and Implementing Enterprise Network Assurance (300-445 ENNA), dashboard widgets must be mapped to the correct telemetry data source and statistical measure. When an agent is installed as a Linux package on an organization's infrastructure, it is classified as an Enterprise Agent.

The correct configuration for this widget is Cloud & Enterprise Agents and Mean (Option B).

Data Source: Since the Linux-based agent is an Enterprise Agent, it shares the same data source category as Cloud Agents in the ThousandEyes dashboard configuration menu.

Measure: To display the "average" packet loss as requested, the Mean (the arithmetic average) is the appropriate statistical measure.

Other options are unsuitable:

Option A: Endpoint Agents are those installed on Windows/macOS user devices, not typically a standard Linux server package used for infrastructure monitoring. Median would show the middle value, not the average.

Option C: Routing refers to BGP data, which does not report packet loss in the same way as synthetic network tests.

Option D: Devices refers to hardware monitored via SNMP, and an nth Percentile is a specific statistical cutoff (like 95th percentile) rather than a simple average.

Question: 64

In the IT operations dashboard, while comparing the latest metrics, what is the time difference between Page Load time and DOM time?

- A. 120.6 ms
- B. 125.3 ms
- C. 100 ms
- D. 150.4 ms

Answer: B

Explanation:

This analysis requires extracting two specific browser-rendering metrics from the IT Operations Dashboard. By placing the cursor over the latest metrics data points:

Page Load time: 1479.6 ms.

DOM load time: 1354.3 ms.

By subtracting the DOM time from the Page Load time ($1479.6 \text{ ms} - 1354.3 \text{ ms} = 125.3 \text{ ms}$), the engineer identifies the duration spent rendering the page after the initial document structure was established.

Question: 65

A CPU utilization alert for Endpoint Agents is triggering too frequently, creating alert noise. Which of the following steps would help reduce the sensitivity of the alert rule? (Select two)

- A. Increase the number of agents that must exceed the CPU threshold to trigger the alert
- B. Lower the CPU utilization percentage in the alert condition
- C. Adjust the alert rule to require more rounds of data to exceed the threshold
- D. Enable the alert rule on more Endpoint Agents

Answer: A, C

Explanation:

Alert fatigue is a major operational challenge in network assurance. To make a CPU utilization alert for Endpoint Agents less sensitive and reduce "noise," an engineer must implement statistical filters that ignore transient spikes or isolated events.

The two most effective methods are:

Increase the agent count (Option A): By requiring a higher number or percentage of agents to simultaneously exceed the threshold, the system ensures the alert represents a widespread environmental issue rather than a single user running a CPU-intensive background process.

Require more rounds of data (Option C): Instead of alerting on a single measurement (1 of 1 round), the engineer can configure the rule to require the threshold to be breached for multiple consecutive checks (e.g., 2 of 3 rounds). This filters out brief, non-impactful CPU spikes that occur naturally during OS updates or browser startups.

Other options would have the opposite effect:

Option B: Lowering the percentage threshold (e.g., from 90% to 50%) would cause the alert to trigger much more frequently, increasing noise.

Option D: Enabling the alert on more agents increases the pool of potential triggers, which typically leads to more notifications unless the logic in Option A is also applied.

Question: 66

You're analyzing NetFlow data for a network supporting voice and video traffic. The data shows consistent spikes in delay and jitter during peak hours. Which optimization would you recommend?

- A. Implement a complete QoS redesign
- B. Increase bandwidth on all network links
- C. Tune the existing QoS configuration to prioritize voice and video traffic
- D. Replace all network hardware with newer models

Answer: C

Explanation:

In the Designing and Implementing Enterprise Network Assurance (300-445 ENNA) curriculum, capacity planning and optimization are driven by telemetry data such as NetFlow. When NetFlow identifies that delay and jitter—metrics highly impactful to real-time traffic—spike during peak hours, it indicates that high-priority packets are competing for resources with bulk data.

The most appropriate recommendation is to tune the existing QoS configuration (Option C). This involves adjusting the Queuing and Scheduling policies on the routers to ensure that voice and video traffic (typically marked with EF and AF41/AF42 DSCP values) is serviced before other traffic classes during periods of congestion. This solution is targeted, cost-effective, and directly addresses the observed jitter issues without the need for massive capital expenditure.

Reviewing other options:

Option A: A complete QoS redesign is often unnecessary and too invasive for solving peak-hour jitter if a basic QoS framework is already in place.

Option B: Increasing bandwidth on "all" links is a "brute force" approach that is expensive and fails to address the underlying problem of traffic prioritization.

Option D: Hardware replacement is a last resort and would not resolve delay/jitter if the new hardware still lacks a properly tuned QoS policy.

Question: 67

SNMP data indicates that a wireless access point is experiencing high channel utilization and increased retransmissions. What optimization would you recommend to improve voice call quality for users on this access point?

- A. Increase the transmit power of the access point
- B. Change the access point to a different, less congested channel
- C. Disable all non-voice traffic on the wireless network
- D. Implement strict admission control for all wireless clients

Answer: B

Explanation:

In wireless network assurance, high channel utilization and increased retransmissions are clear indicators of RF interference or over-subscription. Retransmissions are particularly damaging to voice call quality because they introduce significant jitter and late-arrival packet loss that the jitter buffer cannot overcome.

The recommended optimization is to change the access point to a different, less congested channel (Option B). This directly reduces the competition for airtime (Channel Utilization) and the likelihood of collisions from neighboring access points (Co-Channel Interference), which in turn lowers the retransmission rate. Modern wireless controllers using Radio Resource Management (RRM) often automate this, but a manual adjustment based on SNMP telemetry is a valid operational fix.

Other options are technically flawed:

Option A: Increasing transmit power often increases interference and channel utilization for surrounding cells, making the problem worse for everyone.


Option C: Disabling all non-voice traffic is an extreme measure that is rarely feasible in a real business environment.

Option D: Admission control limits the number of users but doesn't solve the underlying issue of poor channel health for those already connected.

Question: 68

The following exhibit shows the Capacity Planning results for a router interface connected to an ISP,

which provides a 1Gbps connection. Based on the evidence, which action is most likely to fix the observed behavior?



- A. Request a link increase from the ISP
- B. Reconfigure maximum capacity for the interface
- C. Restrict the Web Sites that can be visited from the site
- D. Reconfigure business hours settings

Answer: B

Explanation:

In the context of Designing and Implementing Enterprise Network Assurance (300-445 ENNA), capacity planning requires accurate baselining of interface bandwidth against its theoretical and provisioned limits. Analyzing Exhibit 4.5 Question 4 (image_79d4fc.jpg) reveals a significant discrepancy between the physical reality of the link and its configuration within the monitoring tool.

The exhibit displays a capacity planning dashboard with a calendar heatmap and traffic graphs. The heatmap for February through May shows a high frequency of "Severe" (red) utilization blocks. Looking at the Egress graph for Tue May 07 2024, the traffic spikes clearly exceed 40.0 Mbps.

Crucially, the dashboard indicates an "Egress Capacity" of 49.5 Mbps and reports that the "Highest consumption" was 48.0 Mbps, representing 97% of the available bandwidth.

However, the question states that the ISP provides a 1 Gbps (1000 Mbps) connection. Since the actual traffic being sent is less than 50 Mbps, the link is nowhere near physical saturation. The "Severe" alerts and high utilization percentages are occurring only because the monitoring software (likely ThousandEyes or a similar NMS) is configured with a Maximum Capacity of only 49.5 Mbps for this interface. This misconfiguration causes the tool to calculate utilization based on a much smaller "pipe" than what actually exists, leading to false-positive alerts.

Therefore, the most likely action to fix this observed behavior is to reconfigure maximum capacity for the interface (Option B) to match the 1 Gbps specification.

Option A is unnecessary because the current link is only being utilized at ~5% of its 1 Gbps potential.

Option C is a restrictive policy change that is not justified given the actual available headroom.

Option D might shift how data is displayed but will not fix the underlying mathematical error in utilization calculations.