


Cisco.300-445.v2026-05-04.q32

Exam Code:	300-445
Exam Name:	Designing and Implementing Enterprise Network Assurance
Certification Provider:	Cisco
Free Question Number:	32
Version:	v2026-05-04
# of views:	127
# of Questions views:	320
https://www.freepdfdumps.com/Cisco.300-445.v2026-05-04.q32.html	

NEW QUESTION: 1



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: ([SHOW ANSWER](#))

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 browser-based 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.

NEW QUESTION: 2

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 (LEAVE A REPLY)

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 than¹⁵ generating its own synthetic voice¹⁶ probes. 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

NEW QUESTION: 3

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

Answer: A,C,D (LEAVE A REPLY)

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.

NEW QUESTION: 4

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 ([LEAVE A REPLY](#))

Detailed 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.

NEW QUESTION: 5

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 ([LEAVE A REPLY](#))

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.

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.

NEW QUESTION: 6

Refer to the exhibit.

The screenshot shows a configuration page for a network agent to server test. The fields are as follows:

- Target: 10.10.10.10
- Protocol: TCP
- Port: 80
- Probing Mode: Prefer SACK (selected), Force SACK, Force SYN
- Path Trace Mode: ☐ In Session
- Interval: 2 minutes
- Agents: Select agent(s)
- Alerts: ☒ Enable
- Alert Rules: 1 of 3 alert rules selected
- Edit Alert Rules button

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 7

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: (SHOW ANSWER)

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.

NEW QUESTION: 8

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 (LEAVE A REPLY)

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.

NEW QUESTION: 9

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 (LEAVE A REPLY)

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.

NEW QUESTION: 10

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: (SHOW ANSWER)

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. This data includes granular metrics for every SD-WAN tunnel, such as packet loss, latency, and jitter. 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. 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.

NEW QUESTION: 11

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 (LEAVE A REPLY)

Detailed 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.

NEW QUESTION: 12

Refer to the exhibit.

The screenshot displays the Cisco ThousandEyes 'Monitor Application' interface. On the left, a sidebar shows 'Endpoint Agents' and 'Synthetic Tests' with a list of tests. The main area is titled 'Monitor Application' and shows a grid of popular applications for monitoring. The applications listed are: Custom Application, Google Suite, Microsoft 365, Microsoft Login, Microsoft Teams, Salesforce, ServiceNow, Slack, Webex, and Zoom. Each application card includes a logo, name, description, and '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: (SHOW ANSWER)

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.

NEW QUESTION: 13

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.** 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`
- B.** No action required; the agent will pick up the configuration automatically
- C.** Restart the container using `app-hosting stop appid agentname` followed by `app-hosting start appid agentname`
- D.** Reinstall the agent using the `app-hosting install` command with the correct proxy settings

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 14

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 15

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 (LEAVE A REPLY)

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 time-sensitive, manually entering it into a static test configuration (Option B) is impossible for an automated, recurring test. While some complex workarounds (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.

NEW QUESTION: 16

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 (LEAVE A REPLY)

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.

Valid 300-445 Dumps shared by Actual4test.com for Helping Passing 300-445 Exam! Actual4test.com now offer the **newest 300-445 exam dumps**, the Actual4test.com 300-445 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-445 dumps with Test Engine here:

https://www.actual4test.com/300-445_examcollection.html (70 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 17

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

(Choose two)

AS origin AS 1234 BGP BGP Origin: (ASN not in [1234])

Settings Notifications

GENERAL

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**

ALERT CONDITIONS

All conditions are met by any of 1 monitor 1 of 1 time in a row:

BGP Origin meets all of:

ASN not in 1234

Collapse All

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

Answer: A,C (LEAVE A REPLY)

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.

NEW QUESTION: 18

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: (SHOW ANSWER)

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 time-consuming 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.

NEW QUESTION: 19

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: (SHOW ANSWER)

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.

NEW QUESTION: 20

Refer to the exhibits.







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 ([LEAVE A REPLY](#))

NEW QUESTION: 21

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 22

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: (SHOW ANSWER)

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 thenumberof users but doesn't solve the underlying issue of poor channel health for those already connected.

NEW QUESTION: 23

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 (LEAVE A REPLY)

In theDesigning 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 totune 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.

NEW QUESTION: 24

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 (LEAVE A REPLY)

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.

NEW QUESTION: 25

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 (LEAVE A REPLY)

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.

NEW QUESTION: 26

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 27

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 (LEAVE A REPLY)

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.

NEW QUESTION: 28

Refer to the exhibit.



```
curl -i -XPOST https://api.thousandeyes.com/v7/stream \
-H "Content-Type: application/json" \
-H "Authorization: Bearer $BEARER_TOKEN" -d '{
  "type": "[redacted]",
  "tagMatch": [
    {
      "key": "TestKey",
      "value": "TestValue",
      "objectType": "test"
    }
  ],
  "streamEndpointUrl": "https://example.org",
  "customHeaders": {
    "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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 29

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 ([LEAVE A REPLY](#))

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.

NEW QUESTION: 30

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: ([SHOW ANSWER](#))

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.

NEW 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: (SHOW ANSWER)

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.

Valid 300-445 Dumps shared by Actual4test.com for Helping Passing 300-445 Exam!
Actual4test.com now offer the **newest 300-445 exam dumps**, the Actual4test.com 300-445 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-445 dumps with Test Engine here:

https://www.actual4test.com/300-445_examcollection.html (70 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)

NEW QUESTION: 32

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 (LEAVE A REPLY)

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 time-series 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.

Valid 300-445 Dumps shared by Actual4test.com for Helping Passing 300-445 Exam! Actual4test.com now offer the **newest 300-445 exam dumps**, the Actual4test.com 300-445 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 300-445 dumps with Test Engine here:

https://www.actual4test.com/300-445_examcollection.html (70 Q&As Dumps, **30%OFF**

Special Discount: Freepdfdumps)