DUMPSBOSS.

Designing and Implementing Enterprise Network Assurance(ENWLSD)

Cisco 300-445

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cisco Enterprise Network Assurance Design	9
Topic 2, Cisco Enterprise Network Assurance Implementation	36
Topic 3, Cisco Enterprise Network Assurance Operations	42
Total	87

QUESTION NO: 1

An Enterprise Agent at a branch reports packet loss to a public SaaS service. The branch LAN and WAN edge show no local errors, and the operations team suspects an upstream provider issue. Which two actions provide the strongest evidence for an ISP escalation? (Choose two)

A. Compare the affected test results with results from another Enterprise Agent or Cloud Agent

B. Review path visualization and loss measurements for the point at which degradation begins

C. Restart the branch WAN edge before collecting additional test evidence

D. Raise the packet-loss threshold to stop alert notifications

E. Review CPU utilization on the SaaS provider's application server

ANSWER: A B

Explanation:

Comparing results from another Enterprise Agent or Cloud Agent helps establish whether the event is limited to the branch path or is more broadly associated with the destination service. Reviewing the affected test's path visualization and loss measurements can show where the degradation begins relative to the enterprise edge and provider handoff. These findings, combined with timestamps and affected destinations, provide actionable evidence for the ISP.

Restarting the WAN edge removes useful evidence and is not justified when it has no local errors. Raising alert thresholds only suppresses the symptom. Reviewing application-server CPU may be useful if the evidence points to the SaaS provider, but it does not establish an upstream network impairment on the branch-to-service path.

QUESTION NO: 2

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:

Page Load is the correct test type because it uses a real browser to request a web page, render its resources, and measure the end-user experience of loading that page. This test category provides browser-specific timing and rendering telemetry, including page-completion and Document Object Model loading measurements. Those measurements can be summarized for an executive audience as high-level application experience and availability indicators, while the IT Operations view can expose detailed operational data for investigation. The detailed view can include the timing of individual page resources and browser-derived performance milestones, enabling engineers to identify whether slow user experience is associated with particular web objects, domains, or loading phases. Cisco ThousandEyes documents that Page Load tests use a browser agent and provide page-load performance data, including a waterfall view of the resources involved in rendering a page. This makes Page Load appropriate when the dashboards are intended to represent the experience of a user accessing a web application rather than merely the reachability of a host or a protocol transaction. See Cisco ThousandEyes Page Load Tests and Cisco ThousandEyes Browser Synthetics.

QUESTION NO: 3

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:

SNMP data collected from the routers is the most appropriate source for establishing a CPU-utilization baseline because it provides direct, device-level telemetry from the core routers themselves. Cisco routers expose CPU utilization through SNMP management information bases, including Cisco process CPU objects such as five-second, one-minute, and five-minute utilization values. A monitoring platform can poll these values at consistent intervals, retain historical measurements, and calculate normal operating ranges, trends, and thresholds for each router. This baseline enables the administrator to distinguish expected utilization patterns, such as scheduled traffic peaks, from abnormal sustained CPU consumption that could affect control-plane and forwarding performance. SNMP also supplies related operational counters and interface data that can be correlated with CPU trends when investigating capacity or performance conditions. Cisco documents the CPU MIB objects used to retrieve these measurements in the CISCO-PROCESS-MIB, including cpmCPUTotal5sec, cpmCPUTotal1min, and cpmCPUTotal5min. See the Cisco IOS XE SNMP CPU and memory monitoring documentation and the CISCO-PROCESS-MIB reference.

QUESTION NO: 4

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:

“Increase the number of agents that must exceed the CPU threshold to trigger the alert” reduces sensitivity by requiring the condition to affect a broader portion of the monitored endpoint population before an alert is generated. This helps distinguish a meaningful, widespread CPU-utilization condition from a localized event on one or a few endpoint devices. In an Endpoint Agent deployment, isolated high CPU can occur during normal activities such as operating-system maintenance, application launches, or local security scans; requiring more affected agents makes the alert criterion more representative of an environment-wide issue.

“Adjust the alert rule to require more rounds of data to exceed the threshold” also reduces noise because the threshold breach must persist across multiple measurement intervals. This adds temporal confirmation to the alert condition and filters short-lived CPU spikes that are not operationally significant. Together, population-based qualification and multi-round persistence make alerting more actionable while retaining visibility into sustained or broadly distributed endpoint performance problems. Cisco ThousandEyes alert rules support configuring conditions and alert behavior to focus notifications on events that meet the intended operational impact criteria. See the ThousandEyes Alert Rules documentation and the ThousandEyes Endpoint Agents documentation.

QUESTION NO: 5

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

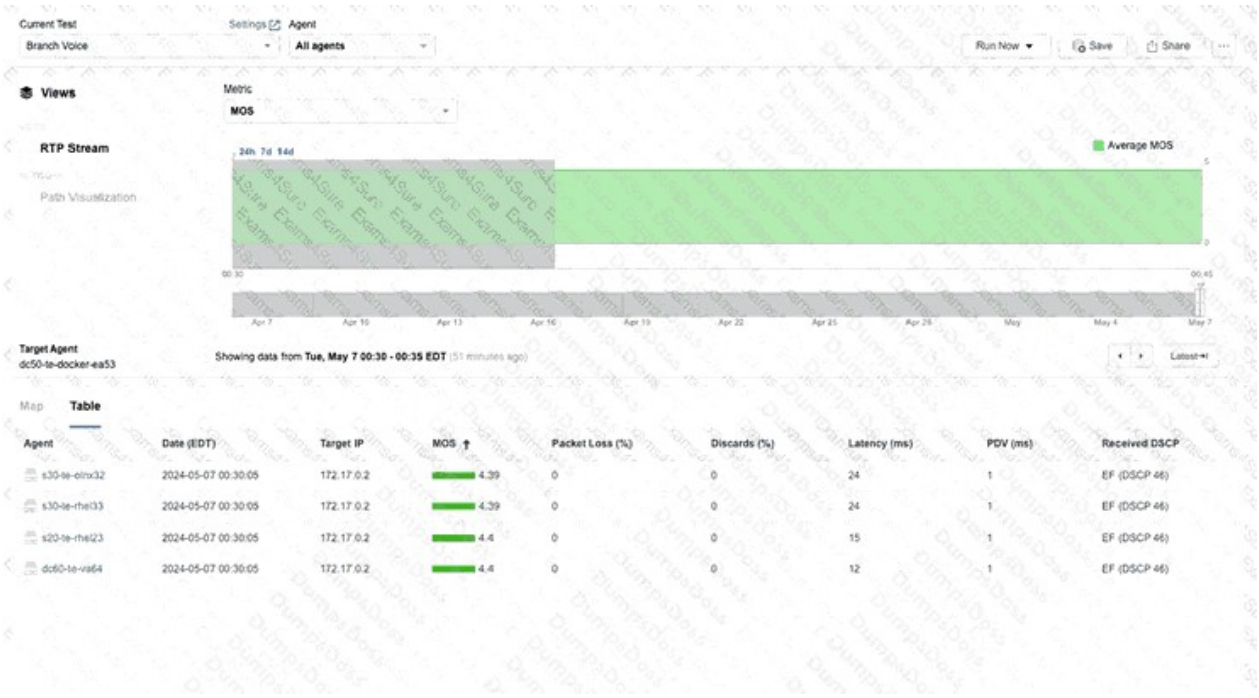
Current Test
Branch Voice
Settings
Agent
All agents
Run Now
Save
Share
Views
Metric
MOS
RTP Stream
Path Visualization
24h 7d 14d
Average MOS
Target Agent
dc50-te-docker-ea53
Showing data from Tue, May 7 00:30 - 00:35 EDT (51 minutes ago)
Map
Table
Agent
Date (EDT)
Target IP
MOS
Packet Loss (%)
Discards (%)
Latency (ms)
PDV (ms)
Received DSCP
s30-te-olnx32
2024-05-07 00:30:05
172.17.0.2
4.39
0
0
24
1
EF (DSCP 46)
s30-te-rhel33
2024-05-07 00:30:05
172.17.0.2
4.39
0
0
24
1
EF (DSCP 46)
s20-te-rhel23
2024-05-07 00:30:05
172.17.0.2
4.4
0
0
15
1
EF (DSCP 46)
dc60-te-vx64
2024-05-07 00:30:05
172.17.0.2
4.4
0
0
12
1
EF (DSCP 46)

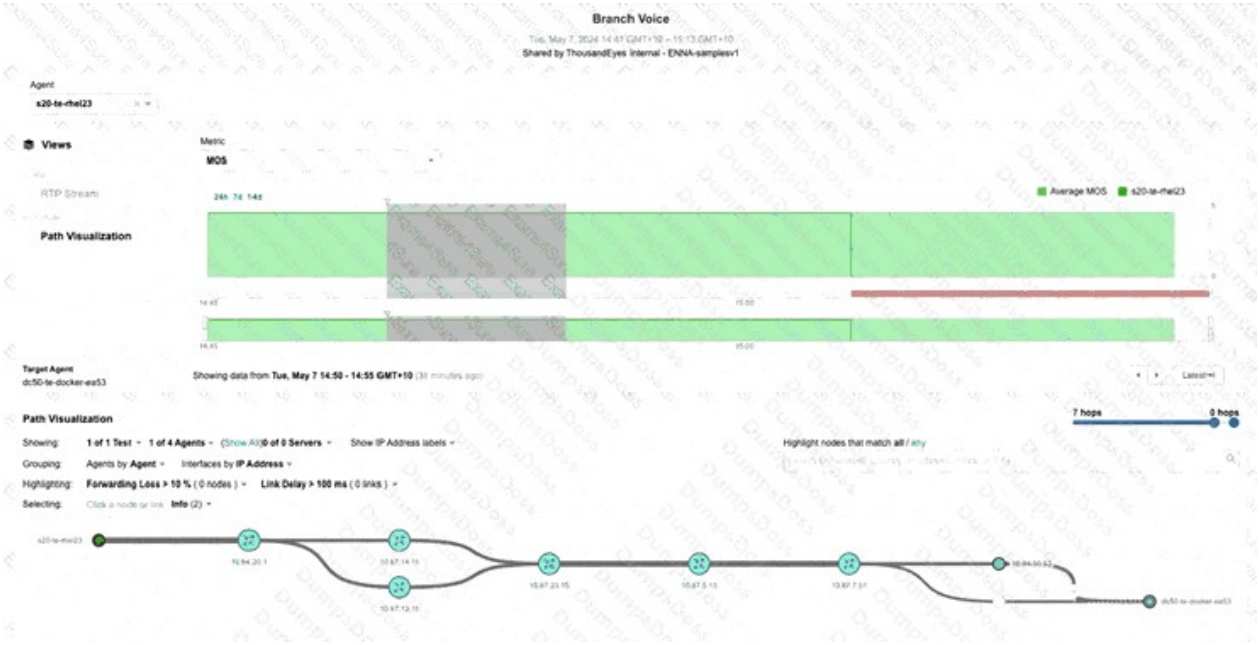
Branch Voice
Shared by ThousandEyes Internal - ENNA-samplesv1
Agent
s20-te-rhel23
Views
Metric
MOS
RTP Stream
Path Visualization
24h 7d 14d
Average MOS
s20-te-rhel23
Target Agent
dc50-te-docker-ea53
Showing data from Tue, May 7 14:50 - 14:55 GMT+10
Path Visualization
7 hops
0 hops
Showing:
1 of 1 Test
1 of 4 Agents
0 of 0 Servers
Show IP Address labels
Grouping:
Agents by Agent
Interfaces by IP Address
Highlighting:
Forwarding Loss > 10 % (0 nodes)
Link Delay > 100 ms (0 links)
Selecting:
Info (2)

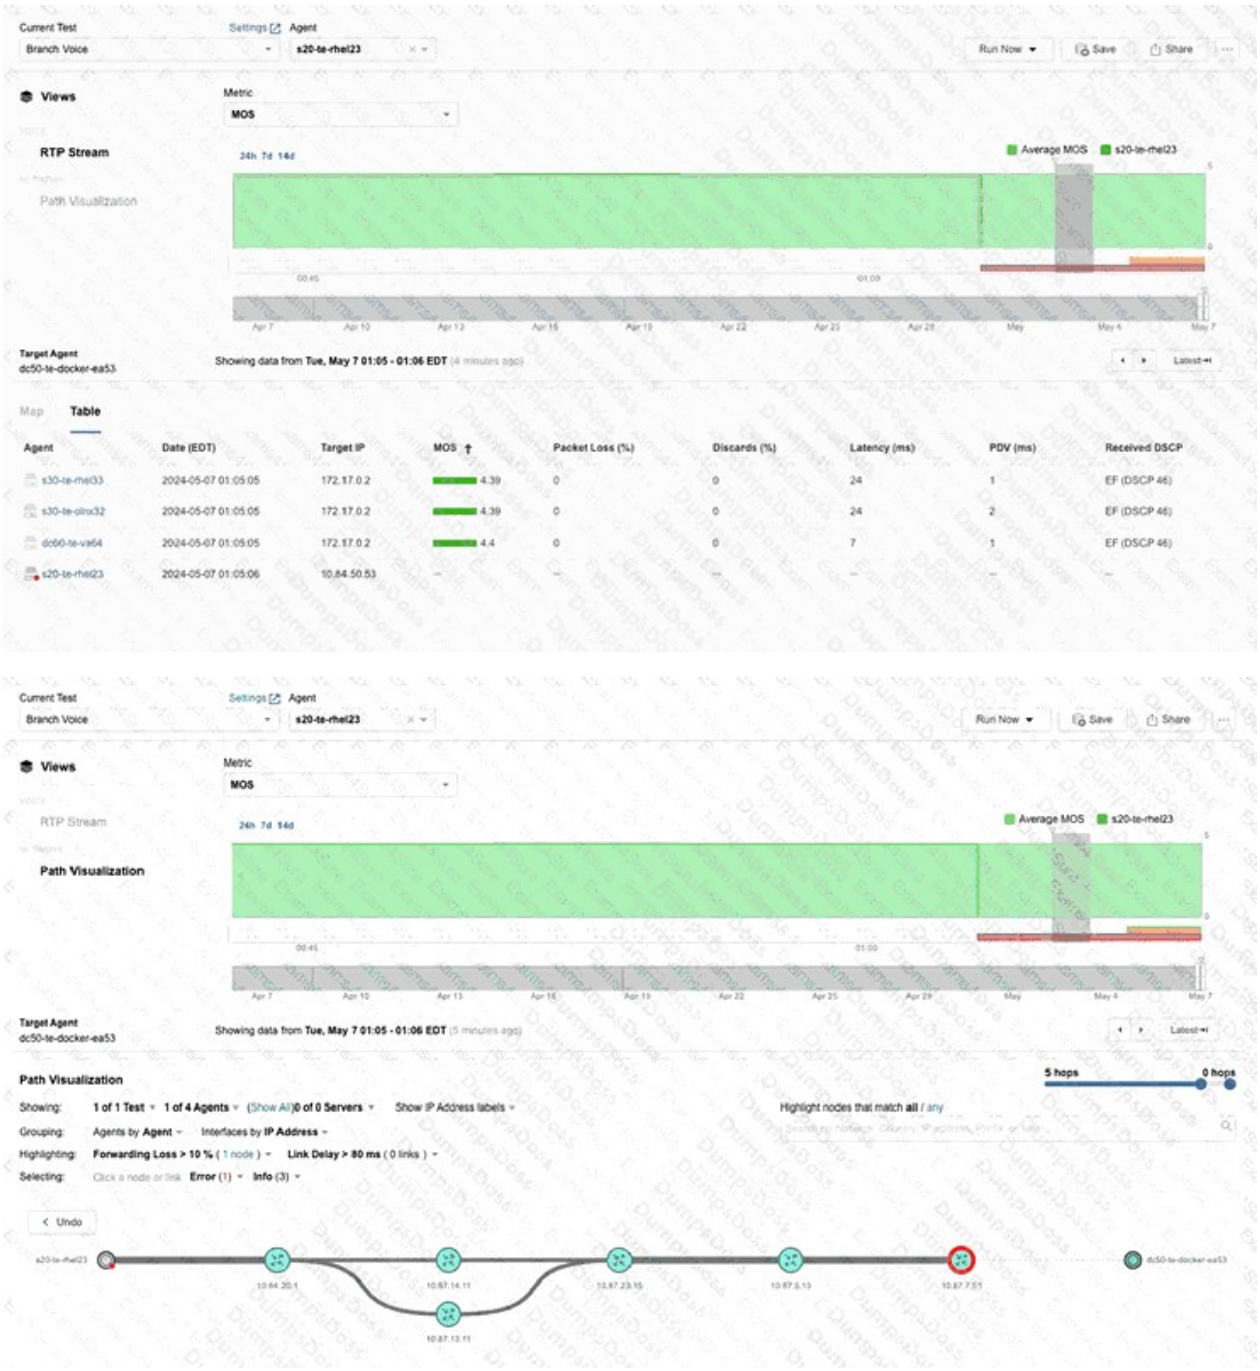

A. Involve the Voice team as the RTP test does not return any relevant results for the agent located at site 20 (identified as s20 in the exhibit)

B. Verify the routing changes on device 10.87.7.51

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

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

ANSWER: C

Explanation:

Verify the docker host 10.84.50.53 and ensure the agent container is running. The exhibits indicate that the remote-site agent associated with site 20 is not producing usable RTP-test data. Before interpreting voice-path metrics or escalating the symptom as a voice-service issue, the collection point itself must be operational. An Enterprise Agent deployed in Docker depends on both the host's reachability and the health of its running container. If the Docker host is unavailable, if the container has stopped or is unhealthy, or if it cannot establish its required outbound connectivity, scheduled tests cannot execute and no valid RTP measurements can be returned.

The appropriate next step is therefore to access 10.84.50.53, confirm that the Docker service is active, inspect the Enterprise Agent container status and logs, and restart or remediate the container as needed. Once the agent has reconnected and appears online in the monitoring platform, rerun the RTP test to obtain valid path, loss, latency, and jitter observations. Cisco ThousandEyes documents Docker-based Enterprise Agent deployment and the need for the agent to maintain connectivity to the ThousandEyes platform: Docker Enterprise Agent installation. Container health and lifecycle verification are also foundational Docker operational practices: Docker container operations.

QUESTION NO: 6

Refer to the exhibit.

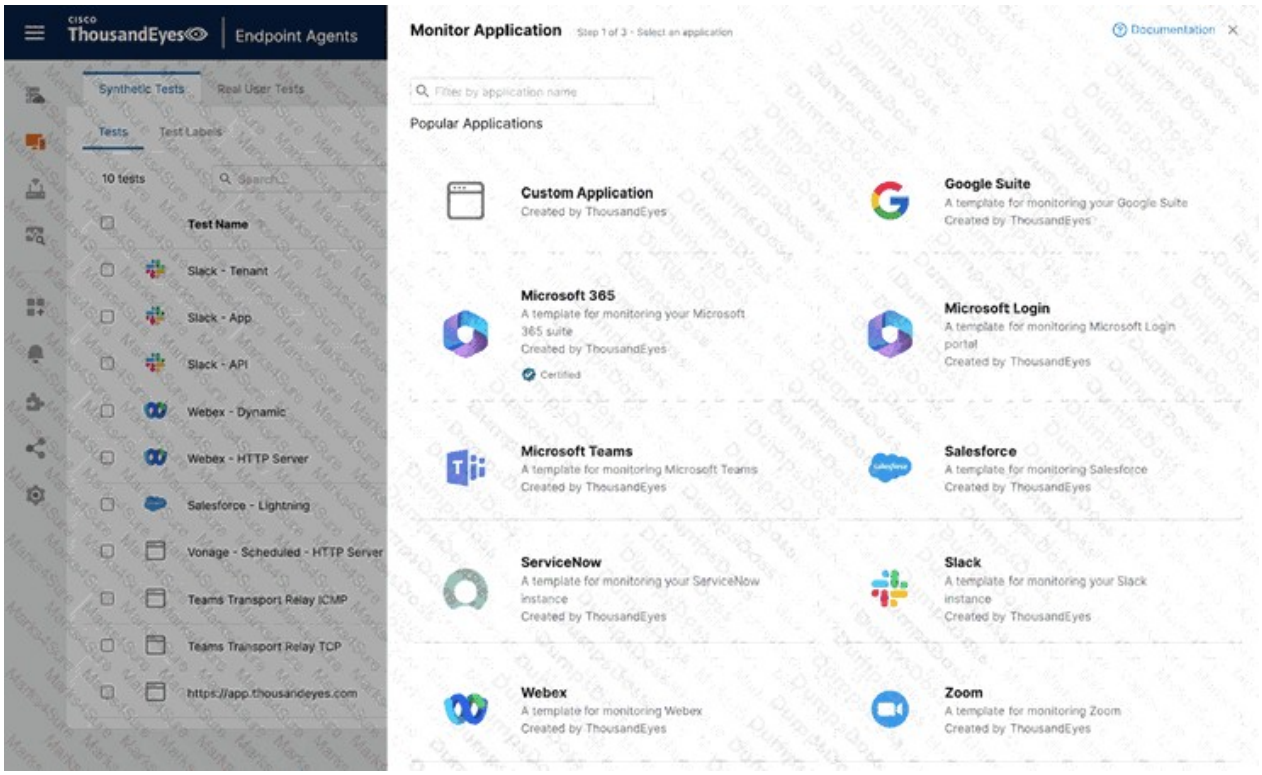

An engineer is tasked with configuring a new test to monitor a web application from the employee's point of view. What two actions should be taken to fulfill the requirement?

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:

Create a new custom application monitor and **Add a new scheduled test to the monitor** provide monitoring from the employees' endpoint perspective. A custom application monitor is appropriate when the target is an organization-specific web application rather than a predefined SaaS application. It defines the application that Endpoint Agents should evaluate, including the relevant web destination and the employee-agent scope.

A scheduled test supplies the proactive, repeating measurement needed to assess that application consistently. Endpoint Agents run scheduled tests from enrolled employee devices at the configured interval, collecting web-application performance and availability observations from the actual endpoint network context. This creates repeatable visibility into user experience, including the client-side perspective that cannot be obtained solely from cloud-agent tests. Together, the

custom monitor identifies the application to observe and the scheduled test establishes the recurring synthetic measurements required for ongoing assurance.

Cisco ThousandEyes documentation describes Endpoint Experience as visibility from user endpoints and explains the use of scheduled tests for proactive endpoint measurements. See the ThousandEyes Endpoint Experience documentation and the Endpoint Agent tests documentation.

QUESTION NO: 7

Thousand23Eyes 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 e25ssential for WAN Insights to function? (Choose two)

A. Configuration data from vManage, su27ch as device templates and centralized policy settings.28

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:

Historical network performance metrics collected by vAnalytics, including loss, latency, and jitter for SD-WAN tunnels, are essential because WAN Insights uses the retained performance history of SD-WAN paths to identify trends and evaluate the expected quality of available transports. The tunnel-level telemetry provides the quantitative basis for predicting whether a given path is likely to satisfy an application's performance requirements. vAnalytics must therefore be enabled and collecting data before WAN Insights can produce meaningful path recommendations.

Application traffic flow data from the SD-WAN data plane, categorized by vManage application lists, is also essential because recommendations must be relevant to the traffic being steered. Application awareness associates observed traffic with the application definitions and intent used by the SD-WAN policy framework. Combining application context with historical loss, latency, and jitter enables WAN Insights to prioritize a path recommendation according to the performance needs of the affected application, rather than treating all traffic identically. Cisco describes WAN Insights as using Cisco SD-WAN telemetry and vAnalytics information to provide predictive insights and recommendations. See the ThousandEyes WAN Insights documentation and Cisco's Cisco SD-WAN analytics overview.

QUESTION NO: 8

An organization needs to monitor a REST API used by a mobile application. The assurance requirement is to validate DNS resolution, TCP and TLS connection behavior, HTTP response codes, server wait time, and the returned response from several regions. Browser rendering and multistep user interaction are not required. Which ThousandEyes test type is the most appropriate?

A. HTTP Server test

B. Page Load test

C. Transaction test

D. Agent-to-agent test

ANSWER: A

Explanation:

An HTTP Server test is the most appropriate choice because it is designed to measure an HTTP or HTTPS service and provides the application- and network-layer timing needed for an API availability and performance requirement. It can validate response codes, connection behavior, server wait time, and response content without introducing browser-rendering or workflow-script complexity.

A Page Load test is browser based and is intended to measure the loading of a rendered web page and its page resources. A transaction test is intended for multistep browser interactions, such as login and checkout workflows. An Agent-to-agent test measures network performance between Enterprise Agents rather than validating an HTTP API response.