



# Cisco Network Assurance Prep

300-445 ENNA

Self-Study Prep Guide

by Cisco ThousandEyes Customer Engineering



# 1.1 Agent Types

 Cognitive Level: Analyze

Determine agent types, such as synthetic user agent, scripting agent, and local collection agent to meet network assurance and security requirements

## Overview

This task assesses your ability to select and deploy the appropriate agent for specific scenarios and business requirements. With agents available on multiple platforms, it is essential to understand supported platforms, deployment methods, and how to align agent deployments with business needs. Hands-on experience and familiarity with industry use cases for each agent type are highly recommended.

While the exam blueprint mentions specific agent types like *synthetic user agent*, *scripting agent*, and *local collection agent*, ThousandEyes uses a different terminology. ThousandEyes focuses on providing network visibility from the agent's point of view and the network path to the destination. To do that, ThousandEyes has 3 types of agents: Cloud Agents, Enterprise Agents, and Endpoint Agents. The available tests also cover application-level metrics, allowing engineers to correlate application issues with network metrics.

## Synthetic User Agent

A synthetic user agent generates real traffic to assess the status of a target, which differs from traditional passive monitoring techniques such as SNMP and flow protocols. All ThousandEyes agents are considered synthetic agents since they actively generate traffic to measure performance.

## Scripting Agent and Local Collection Agent

The term "Scripting Agent" doesn't exist within the ThousandEyes glossary. Instead, it refers to the general concept of how agents collect information. In ThousandEyes, data collection is performed by tests which will be covered in detail in Section 2.3. The closest term used is Transaction Scripts, which is a type of test that allow us to emulate the user workflow using an application to collect web metrics. For example, if you wanted to gather metrics of the user login to an application, you would use this test type.

"Local Collection Agent" is a generic term that depends on the context of the observation. The meaning of a local agent would vary based on the specific use case.

These terms allude to the collection of metrics to provide network visibility. ThousandEyes Agents collect the data by executing tests. Cloud and Enterprise Agents support Network, DNS, Web (Page Load, Transaction, and API tests can only run when the BrowserBot component is enabled while installing the Enterprise Agent) and Voice test types. Endpoint Agents support scheduled, dynamic, and real user test types. Additionally, Endpoint Agents collect local network information where the Agent is connected.

Test Types for Cloud, Enterprise, and Endpoint Agents are covered in greater detail in Domain 2.

# ThousandEyes Agent Types

ThousandEyes agents serve as points of observation on the network where they are deployed. They require test assignment to collect data and are also known as "vantage points." These lightweight Linux-based software agents enable layered monitoring tests that provide insights into network performance but also can help provide metrics from the application layer that can be correlated to understand the impact of network issues on the applications and services delivered to end users or isolate the cause to an application/service problem.




Figure 1.1-1: ThousandEyes Agent Types

## Enterprise Agents

Enterprise Agents are installed on customer's infrastructure to provide insights from inside networks, the network path to applications and services, including the ISP, the Internet, ISP of the application or service provider all the way to the network of the target being monitored.

There is a wide range of platforms that support the Enterprise Agent. For the latest supported platforms, versions, and requirements, visit the product documentation on [Installing Enterprise Agents](#). It is important for the candidate to get experience installing the Enterprise Agent on as many offerings as possible. Some of the easiest options available, considering the candidate might already have access to these platforms, are deploying the Enterprise Agent as a:

- Linux Package <sup>1</sup>
- Docker container
- Virtual appliance (VMWare, VMWare Fusion, Microsoft HyperV)
- Physical appliance (Raspberry Pi)

<sup>1</sup> Linux Package OS versions are continuously updated. For the most up to date information visit [Enterprise Agent Support Lifecycle](#).

## Cloud Agents

Cloud Agents are deployed and maintained by Cisco ThousandEyes in different geographical locations, mainly ISPs and Cloud Providers. These agents can be used without any prior configuration steps, except for assigning tests to them. See [Cloud Agents](#) for more information.

## Endpoint Agents

Endpoint Agents are installed on end-user operating systems such as Windows and Mac. They are also supported by RoomOS-11 devices except for DX, SX, MX. The Endpoint Agent can also be deployed as a module of the Cisco Secure Client bundle. The resources section contains links to installation guides both manual and at scale. To briefly summarize the options, please look at the table below.

| Type                | Manual Installation                 | At-Scale Deployment                                                                               |
|---------------------|-------------------------------------|---------------------------------------------------------------------------------------------------|
| Windows             | <a href="#">Single agent</a>        | <a href="#">Intune Guidance for Software Deployment Teams Group Policies (Browser Extensions)</a> |
| Mac OS              | <a href="#">Single agent</a>        | <a href="#">Munki</a>                                                                             |
| Cisco Secure Client | <a href="#">Cisco Secure Client</a> | <a href="#">Silent/Mass Installation</a>                                                          |
| Cisco Webex RoomOS  | <a href="#">Webex Control Hub</a>   | <a href="#">Webex Control Hub</a>                                                                 |



# Hands-on Activities

We recommend the following activities:

## Activity 1: Enterprise Agent Installation

1. Sign up for a ThousandEyes trial here:



2. Deploy an Enterprise Agent. If you would like to deploy using the Enterprise Agent as a Linux Package, use the Cisco U Tutorial:



If you prefer, you can also use the [Getting Started Guide for Enterprise Agents](#). This tutorial covers the installation steps as well as a brief verification of the agent installation and how it should show up on the ThousandEyes GUI.

3. Extra: Create an HTTP server test to `https://thousandeyes.com` with default values and assign it to the Enterprise Agent you created in step 2.

**i** ThousandEyes agents require internet access to communicate and register to our platform. If you require a proxy, please check the [documentation](#) on how to configure it.

## Activity 2: Endpoint Agent Installation

1. Sign up for a ThousandEyes trial if you have not already done it in activity 1.
2. Install the Endpoint Agent on your computer if you have a supported operating system (Windows or MacOS).
3. Extra: Create an HTTP server test to `https://thousandeyes.com` with default values and assign it to your Endpoint Agent.

You will cover task 1.1 and start with a brief introduction to Domain 2 by following activities 1 and 2.

## Resources

- [Getting Started with Cloud and Enterprise Agents](#)
- [Getting Started with Endpoint Agents](#)
- [Deploy a ThousandEyes Agent in AWS Tutorial](#)
- [Cloud Agents](#)
- [Enterprise Agents](#)
- [Endpoint Agents](#)

# Sample Questions

## 1.1 Question 1

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

- ☒ A) Synthetic Agent
- ☒ B) Enterprise Agent
- ☒ C) Cloud Agent
- ☒ D) Endpoint Agent

## 1.1 Question 2

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

## 1.1 Question 3

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

- ☒ A) Synthetic Agent
- ☒ B) Enterprise Agent
- ☒ C) Cloud Agent
- ☒ D) Endpoint Agent

# 1.2 Agent Location

 Cognitive Level: Analyze

Determine agent location to meet network assurance and security requirements

## Overview

Determining the agent location involves selecting the optimal observation point to collect actionable metrics for problem-solving. When dealing with issues in a service or application that users and customers rely on, consider the following question: From which locations or devices do I need metrics to help pinpoint the root cause of a network problem? In other words, when an engineer needs to troubleshoot a service impacted by the network, they must identify the exact source of the issue to develop a solution. From the ThousandEyes perspective, understanding where to deploy agents and which type of agent to use in each situation is crucial.

Traditional troubleshooting methods include packet captures, SNMP, flow protocols (e.g., NetFlow, sFlow), syslog, etc., all of which are classified as passive monitoring. This guide focuses on synthetic monitoring, which generates traffic from the observation point or agent to the target to retrieve metrics and report a real-time view of the user experience, instead of merely observing the system.

Agent location is related to the protocols that can be used, how they are enabled, the supported platforms for deploying agents, and the specific use cases.

When determining agent location, it is essential to consider the end user of the service, application, or network being assured. The focus should be on delivering an excellent experience to end users by ensuring that the network meets requirements and providing evidence to quickly resolve problems affecting the user.

## Key Concepts

Agent location considerations:

- **Network topology:** What is the architecture and workflow from my users to applications and services? In other words, what do I need to monitor, what are the requirements, what is the best agent and test type that needs to be deployed and where? (i.e., Network engineer needs visibility from their enterprise network to a specific web server. Hence, it would make sense to deploy an Enterprise Agent in the Data Center and configure an HTTP server test to the URL of that web server which will provide BGP, network and web metrics).
- **Security requirements:** Ensuring proper data protection and access control. Keep in mind that the Enterprise Agent requires direct Internet Access in order for the Agent to register to the ThousandEyes platform, obtain test configuration before it can run tests. Proxy and firewall on the path need to be considered.
- **Infrastructure:** Evaluate the best agent to meet your needs. If you have a device (i.e., specific Catalyst 9300/9400, Nexus 9K, Meraki MX, etc.), you can enable the Enterprise Agent on them. You can also deploy it on your virtualization platform (i.e., HyperV, VMWare), as a Linux package or a Docker container on a server (i.e., Ubuntu server), or if you have RoomOS devices, you can enable the ThousandEyes Endpoint agent. If you don't have infrastructure in a specific location, you might be able to find a Cloud Agent in the geographical area.

In this section, we'll analyze where to deploy ThousandEyes Agents based on two main factors:

1. Agent by Vantage Point (Outside-in monitoring, Inside-Out monitoring and Last-mile monitoring)
2. Agent Location by Use Cases

## Agent by Vantage Point

### OUTSIDE-IN MONITORING (CLOUD AGENTS)

Cloud Agents are globally distributed vantage points, managed and maintained by the ThousandEyes Operations team, deployed in tier 2 and tier 3 Internet Service Providers, Internet exchange points, and cloud providers such as AWS, Google, Azure, and Alibaba. These vantage points are capable of running all network, DNS, web, transaction, and voice layer tests available within the ThousandEyes platform, and are available for use by all ThousandEyes customers on a unit consumption basis.

You can use Cloud Agents to provide an "outside-in" or comparison view in places where campuses are located or where users and customers will be accessing sites. One of the advantages of using Cloud Agents is that you don't have to deploy any servers and can get service and network health visibility right away; then later you can add Enterprise Agents based on your requirements.

Cloud Agents are strategically deployed close to end-users in key locations worldwide, providing valuable insights into the user experience from diverse geographical perspectives. They are hosted in various environments, including Broadband ISPs, cloud providers, Webex data centers, and mobile edge compute facilities.

For more information about the location of ThousandEyes Cloud Agents, see the [Cloud Agent World Map](#).

### INSIDE-IN MONITORING (ENTERPRISE AGENTS)

Enterprise Agents are vantage points deployed locally by customers within their own infrastructure to monitor their data centers, cloud VPCs/VNETs, branch offices, and other internal or Internet-based network assets. These agents are ideal for monitoring the experience from internal networks, sites, or branch offices to applications and services hosted externally, such as those in cloud providers, SaaS applications, or API services.

Enterprise Agents are also useful for monitoring network paths and performance metrics between branch offices and data centers, between sites, or even between different cloud providers. To deploy Enterprise Agents, you typically need to have ownership of or access to the infrastructure where the agents will be installed.

Enterprise Agents can be installed as a package on supported Linux distributions or as a pre-packaged virtual appliance that can be deployed on various hypervisor or hardware platforms. They should be placed as close to the users as possible to accurately represent the user experience. In most cases, this means deploying a single Enterprise Agent at each office location, although multiple agents may be necessary for large campuses or to monitor different networks with distinct traffic policies.

For more information, see the [Enterprise Agents](#) documentation.

### LAST-MILE MONITORING (ENDPOINT AGENTS)

Endpoint Agents are lightweight software components installed on employee PCs to monitor applications through a browser plug-in. These agents provide visibility into last-mile performance issues related to Wi-Fi, bandwidth capacity, ISP routing, VPN gateways, and SaaS availability, as they follow the user's location.

Large-scale deployments can be managed using tools like the Cisco Secure Client.



Note that the ThousandEyes controller ( <https://c1.eb.thousandeyes.com> ) automatically handles all Endpoint Agent software updates when the agents check in every 10 minutes or when switching networks.

For more information on managing and troubleshooting Endpoint Agents, refer to the [ThousandEyes Endpoint Agent](#) documentation.

## Use Cases

Use cases represent different scenarios where you can deploy agents to solve specific problems, which will guide your decision on where to deploy the agent.

### Monitoring Hybrid Work

Monitoring hybrid work environments involves considering various domains and tools relevant to the end-user perspective. It is crucial to consider all the components that interact to enable hybrid workers to work from anywhere and how IT can support them when issues arise. Key components to consider when designing a network assurance platform for hybrid work include:

- **Collaboration Tools:** Monitor tools such as Slack, Microsoft Teams, Webex, and Zoom to ensure a smooth user experience. Consider the agent deployment location, required protocols, and types of tests needed to collect relevant metrics.

For this use case, the Endpoint Agent might be the best fit for most scenarios because these agents have dynamic tests that can detect a Webex, Zoom or MS Teams ongoing call that the user is participating in and collect network metrics from the user's PC to the collaboration servers. RoomOS devices also get the Endpoint Agent and can collect metrics for calls.

You can still cover this use case with Cloud & Enterprise Agents in a different way. ThousandEyes has agents installed on Webex data centers so you can run agent-to-agent tests to get the metrics and path towards those endpoints. You can also configure HTTP server tests and have them assigned to Cloud & Enterprise Agents to monitor the web response and network from the collaboration servers. It is recommended to use templates while creating these tests.

- **SaaS and API:** With the migration of workloads to the cloud, monitoring SaaS platforms (e.g., Salesforce, Microsoft 365, Google Suite, ServiceNow) and APIs consumed by application workflows is essential to avoid business disruption and ensure a stable user experience.

Depending on what perspective you want to collect metrics for, you can:

- Deploy Endpoint Agents: If you want to monitor the user experience to specific URLs for the SaaS applications used by them, for example, a specific Salesforce instance.
- Deploy Enterprise Agents when you want to monitor the experience from a branch office, more specifically, from your internal networks to specific SaaS applications, services or Cloud Provider workloads. The test type depends directly on the application/service that needs to be monitored.
- Use Cloud Agents when you want to monitor the experience of your customers. For example, if you are hosting an API or custom application on your Data Center, Cloud Agents will offer several different geographical locations.
- **Wi-Fi Monitoring:** Ensure reliable Wi-Fi connectivity for hybrid workers.

Deploy Endpoint Agents. Endpoint agents are the only type of Agents that collect wireless metrics from the operating

system as well as from tests running to the local gateway.

- **Secure Access:** Utilize tools like Cisco Secure Access Experience Insights to monitor and secure user access.

Deploy Endpoint Agents. The Endpoint Agent is part of the Cisco Secure Client bundle.

## Monitoring SASE (Secure Access Service Edge)

Monitoring SASE involves overseeing a cloud-based network architecture that combines network security functions with WAN capabilities. Key components include:

- SD-WAN (Software-Defined Wide Area Network)
- Cisco Umbrella
- Meraki SD-WAN
- Catalyst SD-WAN

## Monitoring Core Network

Monitoring the core network infrastructure involves overseeing:

- SD-WAN & WAN
- VPN Headend
- Access Networks
- Wi-Fi
- Meraki SD-WAN
- Catalyst SD-WAN
- Catalyst Switching

## Monitoring Hybrid Cloud

Hybrid cloud monitoring involves overseeing both on-premises data center infrastructure and cloud services, including:

- DC Switching
- DC Routing
- Cloud Services
- Cloud Network
- Cisco Nexus
- Multicloud

## Monitoring FSO (Full Stack Observability)

Monitoring FSO involves overseeing applications and their underlying infrastructure using tools such as:

- Open Telemetry
- Cisco AppDynamics

## Resources

- [Beginner's Guide To Enterprise Network Monitoring](#)
- [Endpoint Agent use cases deployment and data privacy models](#)
- [A Simple, Secure Way to Connect Your Branches](#)
- [Cisco Security Application](#)
- [ThousandEyes on Meraki](#)
- [Workforce Digital Experience](#)

## Sample Questions

### 1.2 Question 1

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

### 1.2 Question 2

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.

# 1.3 Active and Passive Monitoring



Cognitive Level: Remember

Describe active and passive monitoring (RFC 7276 and RFC 7799)

## Overview

This section covers the principles of active and passive monitoring techniques as defined in [RFC 7276](#) and [RFC 7799](#).

## Key Concepts

### Active Monitoring

Active monitoring involves injecting test traffic into the network to measure performance metrics.

RFC 7276 defines several key terms related to active network monitoring:

- **Active Measurement** - A form of measurement that relies on packets or sequences of packets that are transmitted across a network to permit a measurement to be performed.
- **Active Metric** - A metric calculated from an active measurement performed across the path between two points, using probe packets.
- **Active Measurement System** - A system that performs active measurements.
- **Probe Packet** - A packet transmitted across a network to permit an active measurement to be performed.
- **Synthetic Traffic** - Traffic generated by an active measurement system and transmitted into a network to perform measurements.

### ACTIVE MONITORING CHARACTERISTICS

Active monitoring relies on injecting dedicated measurement packet streams into the network solely for measurement purposes. This approach generates additional test traffic on the network. Active monitoring allows for the measurement of end-to-end or partial path performance and provides the capability to test specific protocols or services by generating appropriate test packets. It offers greater control over the sampling time and frequency of measurements. Some examples of active monitoring protocols include ping, traceroute, OWAMP, and TWAMP.

### Passive Monitoring

Passive monitoring relies on observing existing traffic as it passes through the network, without injecting any test packets.

RFC 7799 defines several key terms related to passive network monitoring:

- **Passive Measurement** - A form of measurement that does not depend on packets or sequences of packets injected into the network being measured.
- **Observation Point** - A location in the network where packets can be observed for passive measurement purposes.
- **Observation Domain** - The set of all observation points within a network at which passive measurements are made.
- **Flow** - A sequence of packets that have some set of packet header values in common.
- **Flow Record** - A data record containing information about a specific flow that was constituted and observed at an observation point.
- **Flow Key** - A specific combination of packet header values used to define a flow.

## PASSIVE MONITORING CHARACTERISTICS

Passive monitoring relies on observing existing packet streams as they naturally occur in the network, serving measurement purposes without the need for injected test traffic. This approach monitors real user traffic and behavior, providing insights without disrupting network operations. It allows for the measurement of various metrics such as traffic volume and the mix of applications and protocols in use. However, it requires sufficient levels of real traffic to be effective and cannot test specific protocols on demand. Deployment typically involves tapping into network links or configuring span ports on switches. Passive monitoring can encompass all traffic or use sampling techniques to reduce overhead. Examples of protocols used for passive monitoring include IPFIX, sFlow, and PSAMP.

## Hybrid Methods

Instead of being completely separate approaches, active and passive monitoring techniques can be combined as *hybrid methods*. For example, you could add measurement fields to an existing data stream, or attach measurement traffic onto already existing data streams.

## Comparison

This table summarizes the key differences between active and passive monitoring:

| Feature            | Active Monitoring                      | Passive Monitoring                             |
|--------------------|----------------------------------------|------------------------------------------------|
| Test traffic       | Sends out simulated traffic            | Observes real user traffic                     |
| Network impact     | Adds extra work for the network        | No impact on existing traffic                  |
| Metrics measured   | Delay, data loss, jitter, reachability | Traffic volume, traffic types, network usage   |
| Sampling control   | Controlled testing frequency           | Depends on actual traffic levels               |
| Setup requirements | Needs dedicated testing endpoints      | Needs a network tap or span port               |
| Troubleshooting    | Can test specific problems on demand   | Provides a broad view but needs enough traffic |
| Standards          | RFC 4656 (OWAMP), RFC 5357 (TWAMP)     | RFC 7011 (IPFIX), RFC 3176 (sFlow)             |

To get a complete picture of network performance and stability, it's best to use both active and passive monitoring.

## Resources

- [RFC 7276](#)
- [RFC 7799](#)

## Sample Questions

### 1.3 Question 1

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

## 1.3 Question 2

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.
- ☒ 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.



# 1.4 ThousandEyes WAN Insights

 Cognitive Level: Remember

Describe ThousandEyes WAN Insights

## Overview

ThousandEyes WAN Insights is a predictive feature that recommends optimal network paths for user applications within SD-WAN networks. It accomplishes this by:

- **Data Collection and Analysis:** Collecting raw network traffic data from SD-WAN routers to determine path capacity.
- **Forecasting and Recommendations:** Forecasting network conditions by analyzing this data (provided through Cisco Catalyst SD-WAN Manager integration) and recommending the best paths for applications.
- **Actionable Insights:** Enabling network administrators to review and apply these recommendations from ThousandEyes to Cisco Catalyst SD-WAN Manager.

This information helps network and IT teams proactively avoid experience degradation by using the recommended paths. Additionally, WAN Insights is useful for capacity planning because it provides visibility of all network paths and their utilization for the monitored applications and sites.

### Going into the exam, remember:

- WAN Insights is a predictive feature within ThousandEyes that uses performance data from SD-WAN routers in the fabric to make network path recommendations.
- WAN Insights requires integration with Cisco Catalyst SD-WAN Manager (formerly known as vManage).

## Resources

- [WAN Insights | ThousandEyes Documentation](#)
- [Closed Loop Automation in SD-WAN via ThousandEyes - DEVNET-1608](#)

-  **YouTube:** WAN Insights Tutorial | Optimize Experiences Across Cisco SD-WAN

(Video link: <https://www.youtube.com/watch?v=9pJuX0ZeCfA>)

- [Cisco Blog: ThousandEyes WAN Insights](#)

# Sample Questions

## 1.4 Question 1

What is ThousandEyes WAN Insights, and how does it complement Cisco's SD-WAN network infrastructure? Select all that apply.

- ☒ A) A predictive network path tool that uses historical data to recommend optimal paths within Cisco SD-WAN
- ☒ B) A hardware device for real-time network traffic monitoring and analytics
- ☒ C) Provides visibility into network performance, including the public Internet, by working with ThousandEyes
- ☒ D) A set of network management tools that leverage SNMP and flow protocols into a single dashboard
- ☒ E) An antivirus solution that protects networks from cyber threats

## 1.4 Question 2

Which of the following data sources does ThousandEyes WAN Insights use to provide network performance visibility? Select all that apply.

- ☒ A) Historical network data
- ☒ B) Public internet performance data
- ☒ C) SNMP data
- ☒ D) Flow protocol data
- ☒ E) Antivirus data

# 1.5 Cisco Integrations



Cognitive Level: Remember

**Describe the integration between Cisco technologies, such as ThousandEyes, vManage Cisco Catalyst Manager, Webex Control Hub, Meraki, and Endpoint Agent deployment through Secure Client**

## Overview

ThousandEyes integration complements Cisco's network management capabilities, offering deeper insights into performance issues across various environments, from SD-WAN to cloud services.

The way in which ThousandEyes generally integrates with other Cisco technologies is by installing agents on them, usually an endpoint agent or an enterprise agent. In some cases, Cisco devices contain embedded license units to facilitate installation.

Describing these integrations helps you gain a solid foundation of the ThousandEyes platform and its role in enhancing network monitoring within the Cisco ecosystem. This knowledge is useful for those looking to leverage the full potential of their network monitoring solutions.

## Key Concepts

### Cisco Catalyst SD-WAN Integration (formerly vManage)

Cisco Catalyst SD-WAN Manager is a centralized network management platform for SD-WAN deployments. ThousandEyes integration with vManage involves installing ThousandEyes Enterprise agents into Cisco SD-WAN devices and leveraging the ThousandEyes global network of cloud agents.

These Enterprise Agents are managed through the vManage console, allowing network administrators to deploy and configure ThousandEyes tests across the SD-WAN.

The way ThousandEyes Enterprise Agents are installed in vManage is described in the ThousandEyes documentation: [Installing Enterprise Agents on Cisco Routers with vManage](#)

This integration provides visibility into both internal network segments and external paths, including internet and cloud environments.

To review the supported Cisco routers and hardware requirements, see the [Support Matrix](#).

### Webex Control Hub Integration

Cisco's Webex Control Hub is the central interface to manage a Webex organization, add users, assign Webex services, view their usage analytics, and more.




Figure 1.5-1: Webex Control Hub Landing Page

When ThousandEyes integrates with Webex Control Hub, you get direct access to ThousandEyes path visualization data within Webex Control Hub in the Troubleshooting section, selecting a user and looking in Network Path.



Figure 1.5-2: Network Path after integrating ThousandEyes

The Network Path is pulled from ThousandEyes data and can show whether the network path quality is poor during specific sections of the meeting. It can show the hops that the user took to reach the Webex services, loss, latency, jitter, location, device name and device IP address of the end-to-end path and from each hop.




Figure 1.5-3: Detailed Network Path

From Network Path, you can launch ThousandEyes Dashboard and get more details from the endpoint agent.




Figure 1.5-4: Launching ThousandEyes Dashboard from Network Path

The information provided in Network Path helps network administrators identify root causes for negative participant experiences.

Besides Network Path, different observation points can be monitored with this integration since endpoint agents, enterprise agents and cloud agents can be used, therefore different tests can be configured. For example:


- Webex Desktop Client can be monitored if an endpoint agent is installed on the computer (Mac or Windows)
- RoomOS devices registered to Control Hub already have the Endpoint Agent units built into the firmware.
- Performance of Webex Meetings can be monitored with Cloud and Enterprise Agents

## Meraki

The integration between ThousandEyes and Meraki MX devices empowers distributed organizations to monitor external applications and services effectively.

This solution leverages the Meraki Insight (MI) feature which is designed to give Meraki customers an easy way to monitor the performance of web applications and WAN Links on their network and easily identify if any issues are likely caused by the network (LAN or WAN) or the application server. The data used by MI is based on end-user HTTP/S data that are already traversing the MX appliance and does not need synthetic probing.

With the ThousandEyes integration, customers can create customized network and application testing for critical applications inside or outside their infrastructure. For example, customers can monitor their internal DNS server response time and availability and measure the average resolution time for a specific domain.



*Figure 1.5-5: Example of critical applications that can be monitored*

The ThousandEyes Agent is embedded in the MX architecture and does not require external hardware or out-of-band activation. It uses ThousandEyes Enterprise Agents deployed on the MX via Meraki Dashboard as the vantage point monitoring customers' data centers, cloud VPCs/VNETs, branch offices, and other internal or Internet-based network assets.

After the Enterprise Agent is deployed on the MX Appliance, it will register with the Meraki and ThousandEyes backend cloud to download all the test and management configurations using encrypted communication.

The ThousandEyes Test template feature uses predefined protocols, probing intervals, and alerts templates that simplify the monitoring configuration for the target applications. The Enterprise Agent behaves as a service inside the MX Appliance and uses the Appliance hardware and firmware as a base platform to perform the Agent monitoring tasks.

Active Application Monitoring

88 days left of trialBeta

Powered by ThousandEyes

Monitored NetworksSettings

Search by network name

9 matching results

View applicationsAdd tests

| Monitored Networks | Location      | Last contact  | Applications | Enable                              | Remove |
|--------------------|---------------|---------------|--------------|-------------------------------------|--------|
| SC-Office          | California    | Aug 29, 22:26 | View         | <input checked="" type="checkbox"/> |        |
| MX450-HA           | Fremont Area  | Aug 29, 22:26 | View         | <input checked="" type="checkbox"/> |        |
| TE Network 1       | United States | —             | View         | <input checked="" type="checkbox"/> |        |
| TE Network 2       | United States | —             | View         | <input checked="" type="checkbox"/> |        |
| TE Network 3       | United States | —             | View         | <input checked="" type="checkbox"/> |        |
| TE Network 4       | United States | —             | View         | <input checked="" type="checkbox"/> |        |
| TE Network 5       | United States | —             | View         | <input checked="" type="checkbox"/> |        |
| TE Network 7       | United States | —             | View         | <input checked="" type="checkbox"/> |        |
| Child Network 2    | United States | —             | View         | <input checked="" type="checkbox"/> |        |

Figure 1.5-6: Example of monitored applications after ThousandEyes was integrated

### Cisco Secure Client

The ThousandEyes Endpoint Agent can be deployed as a module within the Cisco Secure Client bundle.

These Endpoint Agents can be deployed across devices that connect to the corporate network remotely, allowing them to monitor network conditions from the user's perspective, both before and after the VPN connection is established.

The Cisco Secure Client installer can be deployed using various methods, such as pre-deployment package, a web deployment package, or by loading it directly onto the headend (ASA Firewall, FTD, or ISE server). Additionally, deployment can be facilitated through SecureX and within an XDR environment.

Once an Endpoint Agent is active, it attempts to register with the ThousandEyes service using a secret key and a unique machine ID. Upon successful registration, the Endpoint Agent downloads the latest configuration. When the configuration is up-to-date, the data collection begins with browser-based data and periodic local network probes. The endpoint agent can then measure application performance with scheduled tests, automated session tests, browser sessions, or network access. All of the collected data and metrics are displayed in one expansive view.

### Cisco Secure Access

You can think of Secure Access as your one-stop SASE provider for securing the end-user with multiple layers of defense. The ThousandEyes Endpoint Agent can also be registered with this platform in order to provide visibility of the end user perspective to the Security Service Edge (SSE).




Figure 1.5-7: Cisco Secure Experience Insights

Instead of installing the Endpoint Agent as a standalone application, you can now enable it as a module for environments with existing Cisco Secure Client deployments. Keep in mind that only one of the installers can be used: either the ThousandEyes Endpoint Agent or the Endpoint Agent as a module in the Cisco Secure Client bundle, but not both.

The integration allows IT teams to monitor network performance directly from the endpoint, ensuring that remote users experience consistent connectivity and service quality.

By monitoring network performance on devices using Cisco Secure Client, IT can better diagnose issues related to remote connectivity, including VPN performance.

## DNA

ThousandEyes integration with Cisco DNA Center involves the deployment of ThousandEyes Enterprise Agents throughout the network infrastructure managed by DNA Center.

These agents can be installed on Cisco Catalyst 9000 series, more specifically 9300 and 9400 switches.

The integration enables Cisco DNA Center to leverage the network performance monitoring and analytics capabilities of ThousandEyes, providing IT teams with extended visibility into the network.

Cisco DNA Center uses these agents to collect various types of telemetry data, including Internet and cloud performance metrics, and integrates this data into its assurance and analytics dashboard. This allows network operators to correlate internal network performance data gathered by DNA Center with external network data collected by ThousandEyes, creating a comprehensive view of network health and performance.




Other benefits that this integration provides are:

- **Enhanced Troubleshooting:** Quickly identify and troubleshoot network issues that originate beyond the enterprise perimeter, such as ISP or cloud provider problems.
- **Proactive Monitoring:** Proactively monitor user experience for cloud applications, allowing IT to anticipate performance degradation and resolve issues before users are affected.
- **Data-Driven Decision-Making:** Leverage combined insights from DNA Center and ThousandEyes to make informed decisions regarding network and application performance improvements.

The integration process consists of installing an enterprise agent on a Catalyst 9000 Series switch using DNA center. In general, the process is as follows:

1. Add a new Enterprise Agent and select Cisco Application Hosting tab to download the correct image for Catalyst 9000 switches. Copy the account group token.
2. Log into the DNA Center and navigate to App Hosting for Switches. Use the "New Application" button to upload the Enterprise Agent image.
3. Configure the application.



*Figure 1.5-7: Example of Enterprise Agent in DNA Center*

## Summary

The tables below will help you remember concepts described before.

Table 1.5-1: Integration focus and objectives

| Integration                   | Primary Purpose or Focus                                 |
|-------------------------------|----------------------------------------------------------|
| Cisco Catalyst SD-WAN Manager | Network performance insights across the SD-WAN           |
| Webex Control Hub             | Performance monitoring for Webex services                |
| Meraki                        | Monitoring external applications from SD-WAN sites       |
| Cisco Secure Client           | Network performance data from user devices (VPN clients) |
| Cisco DNA                     | End-to-end network visibility and analytics              |

Table 1.5-2: Agent types and pre-installation status

| Integration                   | Pre-installed Agent Units | Type of Agent Used to Monitor                                         |
|-------------------------------|---------------------------|-----------------------------------------------------------------------|
| Cisco Catalyst SD-WAN Manager | No                        | Enterprise Agents                                                     |
| Webex Control Hub             | No                        | Endpoint, Cloud, Enterprise Agent depending on the observation point. |
| Meraki                        | Yes                       | Enterprise Agents                                                     |
| Cisco Secure Client           | Yes                       | Endpoint Agents                                                       |
| Cisco DNA                     | Yes                       | Enterprise Agents                                                     |

**Notes:**

- **Cisco Catalyst SD-WAN Manager:** ThousandEyes agents are not pre-installed; they must be deployed on the SD-WAN devices.
- **Webex Control Hub:** Performance monitoring is done through Network Path, therefore there are no pre-installed agent units. Except for RoomOS devices that contain embedded Endpoint Agent units.
- **Meraki:** Some Meraki MX models have support for ThousandEyes Enterprise Agents pre-installed, allowing for direct activation and use.
- **Cisco Secure Client:** The Endpoint Agent component of ThousandEyes is integrated as a module within the Secure Client.
- **Cisco DNA:** Select Cisco Catalyst 9000 series switches and other devices managed by DNA Center come with embedded ThousandEyes Enterprise Agent units.

## Resources

- **Cisco Catalyst SD-WAN Manager:** [vManage Integration](#), [Installing Enterprise Agents on Cisco Routers with vManage](#)
- **Webex Control Hub:** [Webex Control Hub Integration](#), [Webex Control Hub Integration Tutorial](#)
- **Meraki:** [Meraki MX ThousandEyes Configuration Guide](#), [ThousandEyes Meraki MX Devices Demo](#)
- **Cisco Secure Client:** [Cisco Secure Client Integration](#), [Cisco Secure Access Experience Insights](#)
- **Cisco DNA:** [Installing Enterprise Agents with DNA Center](#), [Cisco DNA Center Integration](#)

## Sample Questions

### 1.5 Question 1

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 Agent
- ☒ D) Endpoint Agent

### 1.5 Question 2

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

### 1.5 Question 3

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

## 1.5 Question 4

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.

## 1.5 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

# 1.6 Metric Baseline



Cognitive Level: Remember

Describe setting a metric baseline

## Overview

A metric baseline is a reference point that represents the normal or expected performance of a network or system. It is established by collecting and analyzing performance data over a specific period. Baselines are essential for network monitoring and troubleshooting as they allow you to:

- **Assess performance:** Compare current measurements against the baseline to determine if the network is performing as expected.
- **Detect anomalies:** Quickly identify deviations from the baseline that may indicate a performance issue or security threat.
- **Troubleshoot problems:** Use the baseline to understand the magnitude of a problem and how it differs from normal behavior.

## Key Concepts

Establishing a metric baseline involves several key steps:

1. **Define metrics and objectives:** Determine which metrics are crucial for monitoring network health, such as bandwidth utilization, CPU usage, or application response time. Clearly define what you aim to achieve by establishing a baseline for these metrics.
2. **Choose a time period:** Select a timeframe for data collection that captures representative network behavior, including both peak and off-peak traffic patterns. The length of this period depends on the stability and variability of the network environment.
3. **Collect data:** Gather relevant data points from appropriate sources, such as ThousandEyes tests, SNMP monitoring, or network device logs.
4. **Analyze data:** Examine the collected data to identify patterns, trends, and typical performance levels. Use statistical methods to determine average values, percentiles, or standard deviations for each metric.
5. **Establish baseline:** Based on the data analysis, set a reference point that represents the normal or expected performance for each metric.
6. **Monitor and update:** Continuously monitor network metrics over time and update the baseline as network conditions or requirements evolve.
7. **Define alert thresholds:** Set thresholds based on the baseline values to trigger alerts when metrics deviate significantly from the expected range. This allows you to proactively detect and address potential performance issues.

## Resources

- [Getting Started with Dashboards](#)
- [Creating and Editing Alert Rules](#)
- [Dynamic Baselines](#)
- [The Art of Designing ThousandEyes Alert Rules](#)
- [Create a Baseline Dashboard](#)
- [Event Detection](#)

## Sample Questions

### 1.6 Question 1

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

### 1.6 Question 2

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.

# 1.7 Integration Types



Cognitive Level: Apply

Select the integration type, such as API, alerting thresholds, open telemetry, and ITSM for the requested data

## Overview

When an alert is triggered by your networking monitoring system, the relevant data should be presented somewhere that an engineer can see it and take action. When designing and implementing a networking monitoring solution, it is important to consider the different ways that data and alerts can be presented, exported, or integrated into other platforms to appropriately react to network events based on your business needs. This section will examine the different kinds of integrations available for exporting and presenting data from ThousandEyes.

## Types of Integrations

### API

An **Application Programming Interface (API)** allows third-party tools to interact with another application. ThousandEyes offers an API that lets you list and create synthetic tests, configure agents, and perform many other administrative tasks.

APIs can be interacted with in various ways, such as using libraries or packages available for different programming languages or purpose-built third-party tools like Postman. However, designing a system around interacting with APIs can be time-consuming and costly, requiring custom in-house applications or third-party tools.

#### Examples of Prebuilt API Integrations:

- **DNA Center:** Deploys monitoring agents (Enterprise Agents) onto devices managed by DNA Center and presents data for monitored applications.
- **Meraki:** Deploys monitoring agents (Enterprise Agents) onto devices managed by Meraki.
- **Webex Control Hub:** Deploys end-user monitoring agents (Endpoint Agents) onto devices managed by Webex Control Hub and shows network path data alongside relevant call performance data.

| Pros                                                                                              | Cons                                                                                        |
|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Flexible and can be used to create new tests, view data, and complete other administrative tasks. | Requires labor and overhead to maintain custom applications.                                |
| Can interact with many different sources and build custom applications.                           | Can be complicated to set up and is not real-time.                                          |
| Prebuilt integrations with other applications.                                                    | Management and storage of results, like alert states or test configurations, are necessary. |

## ALERT THRESHOLDS

Alert thresholds are essential for notifying about incidents as soon as they occur. They can be configured to trigger notifications when certain conditions are met, with details, including alert status, being queryable via the API. Notifications can be sent via emails to registered platform users or external recipients, through custom webhooks or custom-built integrations.

## Custom Webhooks

Webhooks are HTTP requests sent to a target URL to perform an action on that target server using data from the webhook body. Custom webhooks use a templating format to customize the HTTP request body using webhook variables and logic.

## SERVICENOW

ServiceNow is an IT Service Management (ITSM) tool that can receive webhooks and create service tickets based on the webhook data for engineers to act on.

| Pros                                                | Cons                                                        |
|-----------------------------------------------------|-------------------------------------------------------------|
| Email, webhooks, and alert integration methods.     | No visibility to test metrics before an alert is triggered. |
| Alert state (triggered/cleared) is managed for you. |                                                             |
| Alerts are triggered if certain conditions are met. |                                                             |

## IT Service Management (ITSM)

**IT Service Management (ITSM) Tools** allow a team to manage tasks related to the administration and operation of the IT infrastructure. Most ITSMs can ingest webhooks or query APIs to create service requests, tasks, events, or incidents.



| Pros                                                                            | Cons                                                                             |
|---------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Flexible interface to create and manage service requests, incidents, and tasks. | Requires labor and overhead to manage and maintain.                              |
| Can ingest alerts generated from ThousandEyes and other platforms.              | Usefulness depends on the quality and quantity of the data and tools integrated. |
| Becomes more useful as more data and tools are integrated.                      |                                                                                  |

## OpenTelemetry

**OpenTelemetry** is an observability framework that facilitates streaming real-time telemetry data for collection, monitoring, and reporting.

| Pros                                                                       | Cons                                                                                      |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Standardized framework for streaming data to data visualization platforms. | Requires a platform to ingest, store, and visualize the metrics.                          |
| Real-time streaming of test metrics.                                       | Needs additional components like alerting or visualization tools to make the data useful. |

## Resources

- [ThousandEyes Developer Documentation](#)
- [Getting Started with the ThousandEyes API](#)
- [Automation and IT Ops Integration](#)
- [ThousandEyes Integration Guides](#)
- [Creating and Editing Alert Rules](#)
- [Alert Notifications](#)
- [Custom-Built Integrations](#)
- [OpenTelemetry API](#)
- [Data Observability Backend with OpenTelemetry](#)
- [What is OpenTelemetry?](#)
- [What is ITSM?](#)

# Sample Questions

## 1.7 Question 1

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) OpenTelemetry
- ☒ B) DNA Center Integration
- ☒ C) Custom Webhooks
- ☒ D) Alerts API

## 1.7 Question 2

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

## 1.7 Question 3

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

# 1.8 Cisco Network Assurance Platforms



Cognitive Level: Apply

Select a Cisco network assurance platform based on business requirements, such as application communication, user experience, web, and events

## Overview

To be prepared for this section of the exam, you need to identify the right platform based on a business need. This requires knowing the capabilities, use cases, and focus of each network assurance platform. Below is a high-level overview of the main platforms and their key features.

The network assurance platforms you must know for this section are:

1. ThousandEyes
2. Meraki Insights
3. AppDynamics
4. Catalyst Center (formerly DNA Center)
5. Cisco Catalyst SD-WAN Manager (formerly vManage)

## ThousandEyes

- **Main Features:** Network Intelligence, End-to-End Visibility, Internet and WAN Monitoring, Cloud and SaaS Performance Analysis, VoIP and Video Monitoring, BGP and Route Visualization.
- **Primary Use Cases:** Monitoring network performance and issues across the internet, cloud, and enterprise WANs.
- **Unique Focus:** Digital experience monitoring with visibility into every network layer and service.

## Meraki Insights

- **Main Features:** Network Health Scores, WAN and LAN Monitoring, Application Health Scores, Remote Worker Connectivity, Meraki Device Integration.
- **Primary Use Cases:** Ensuring optimal performance of WAN, LAN, and cloud applications within Meraki-based networks.
- **Unique Focus:** Simplifying the monitoring and management of Meraki networks.

## AppDynamics

- **Main Features:** Application Performance Monitoring (APM), Business Performance Monitoring, End-User Monitoring, Infrastructure Visibility, Network Performance Monitoring, Machine Learning and Analytics.
- **Primary Use Cases:** Monitoring and optimizing application performance across cloud and on-premises environments.
- **Unique Focus:** Application performance management and monitoring.

## Catalyst Center

- **Main Features:** Centralized Management, Automated Troubleshooting, Security and Policy Enforcement, AI/ML Insights, Integration with Cisco Platforms.
- **Primary Use Cases:** Managing and automating Cisco Catalyst switches, routers, and wireless networks.
- **Unique Focus:** Comprehensive network management for Cisco Catalyst infrastructure.

## Cisco Catalyst SD-WAN Manager

- **Main Features:** SD-WAN Network Management, Policy Administration, Zero Touch Provisioning, Historical Performance Monitoring, Security Management, Cloud OnRamp.
- **Primary Use Cases:** Managing and optimizing SD-WAN deployments.
- **Unique Focus:** Centralized management and optimization of SD-WAN infrastructure.

## Resources

For more detailed information on each platform, please refer to the following resources:

- [ThousandEyes Platform Overview](#)
- [Meraki Insights Introduction](#)
- [AppDynamics Overview](#)
- [Cisco Catalyst Center Solution Overview](#)
- [Cisco Catalyst SD-WAN Solution Overview](#)

## Sample Questions

### 1.8 Question 1

You are a network engineer at a multinational corporation responsible for ensuring optimal performance and security across various environments, including remote hybrid workers, branch offices, and cloud services. Select a SaaS-based Network Assurance platform that enables comprehensive monitoring and visibility into hybrid worker activities, internet traffic, and branch office connectivity.

- ☒ A) Meraki Insights
- ☒ B) AppDynamics
- ☒ C) ThousandEyes
- ☒ D) Catalyst Center

## 1.8 Question 2

As a network engineer, you need to select a network assurance platform that provides end-to-end visibility and metrics for remote workers accessing SaaS applications. The solution should monitor the user experience from the endpoint device, through the VPN, across the internet, and to the SaaS provider. Which platform is best suited for this use case?

- ☒ A) Catalyst Center
- ☒ B) AppDynamics
- ☒ C) Meraki Insights
- ☒ D) ThousandEyes

## 1.8 Question 3

Which network assurance platform is best for providing network visibility and performance across any network, where metrics can be correlated with application-level metrics, including for services in multi-cloud deployments?

- ☒ A) Catalyst Center
- ☒ B) AppDynamics
- ☒ C) Meraki Insights
- ☒ D) ThousandEyes

## 2.1 Enterprise Agent Configuration

 Cognitive Level: Apply

Configure enterprise agent on application servers and network infrastructure devices, including dedicated devices

### Overview

This task assesses your ability to configure and install enterprise agents on specific devices and network infrastructure. Key points to understand:

1. There are no specific "application agents" in ThousandEyes.
2. Enterprise agents can be installed on application servers for monitoring if the operating system is supported.
3. Enterprise agents can be configured on various network infrastructure devices.

#### Preparing for the Exam

The main deployment methods to focus on are:

- Cisco App Hosting with Docker via CLI (for switches and routers)
- DNA Center Switch Agent Configuration
- Meraki MX Enterprise Agent Configuration
- vManage Router Enterprise Agent Configuration
- Nexus switches (app-hosting or guest shell)

Each method has its own specific steps and considerations. Hands-on experience with these different deployment methods will be valuable for the exam.

 **YouTube:** *ThousandEyes Integration with Cisco Networking Platforms*

(Video link: <https://www.youtube.com/watch?v=40slh7N2sII&t=184s>)

### Key Concepts

#### Cisco App Hosting

Cisco App Hosting is a method for installing applications, including ThousandEyes enterprise agents, on supported Cisco devices. This approach leverages containerization technology to run applications securely on network infrastructure.

## CONCEPT AND ARCHITECTURE

Cisco App Hosting uses a separate namespace from the main operating system for security reasons. The ThousandEyes enterprise agent runs within a Docker container, isolated from the device's core functions.

## SUPPORTED DEVICES

You can find a [list of supported devices for Cisco App Hosting](#) in the Cisco App Hosting on the Catalyst 9000 Series Switches White paper.

## Cisco App Hosting Command Lifecycle

When working with Cisco App Hosting, it's crucial to understand the lifecycle of app management commands. These commands follow a logical sequence for installing, activating, running, stopping, deactivating, and uninstalling applications.

### ↓ Deploying an App

- a. Install:** `app-hosting install appid myapp package usbflash1:myapp.tar`
- b. Activate:** `app-hosting activate appid myapp`
- c. Start:** `app-hosting start appid myapp`

```
$ app-hosting install appid myapp package usbflash1:myapp.tar
Installing myapp...
Installation complete.
$ app-hosting activate appid myapp
Activating myapp...
Activation complete.
$ app-hosting start appid myapp
Starting myapp...
myapp is now running.
```

`install > activate > start`

### 🗑 Removing an App

- a. Stop:** `app-hosting stop appid myapp`
- b. Deactivate:** `app-hosting deactivate appid myapp`
- c. Uninstall:** `app-hosting uninstall appid myapp`

```
$ app-hosting stop appid myapp
Stopping myapp...
myapp has been stopped.
$ app-hosting deactivate appid myapp
Deactivating myapp...
Deactivation complete.
$ app-hosting uninstall appid myapp
Uninstalling myapp...
Uninstallation complete.
```

`stop > deactivate > uninstall`

For detailed steps and more information on installing enterprise agents on Cisco routers with Docker, refer to the following resources:

- [Installing Enterprise Agents on Cisco Routers with Docker](#)
- [Cisco IOS XE Programmability Configuration Guide - ThousandEyes Integration](#)

Hands-on experience with Cisco App Hosting and configuring ThousandEyes enterprise agents using this method is highly recommended for exam preparation.

## Deployment Methods

ThousandEyes Enterprise Agents can be installed on various Cisco devices using different methods. Here's an overview of the main deployment options:

## CATALYST SWITCHES

- **App Hosting:** Use Docker containers to run ThousandEyes agents.
- **DNA Center:** Deploy and manage agents through Cisco DNA Center.

For detailed instructions, see:

- [Installing Enterprise Agents with DNA Center](#)
- [Cisco DNA Center User Guide - ThousandEyes Enterprise Agent](#)

## NEXUS SWITCHES

- **App Hosting:** Similar to Catalyst switches, use Docker containers.
- **Guest Shell:** Run the agent in a secure Linux container environment.

## ROUTERS

- **App Hosting:** Deploy agents using Docker containers.
- **vManage:** For SD-WAN environments, deploy agents through vManage.

For router deployments, refer to:

- [Installing Enterprise Agents on Cisco Routers with Docker](#)
- [Installing Enterprise Agents on Cisco Routers with vManage](#)

 **YouTube:** *Configuring ThousandEyes on SD-WAN Devices*

(Video link: <https://www.youtube.com/watch?v=4Vgc-fZ66TE>)

## MERAKI DEVICES

- **Meraki MX:** Security appliances that support ThousandEyes agent installation.

For Meraki deployments, see [Meraki MX ThousandEyes Configuration Guide](#).

 **YouTube:** *Meraki Integration with ThousandEyes*

(Video link: <https://www.youtube.com/watch?v=DsDSJRZW9DM>)

 **YouTube:** *Optimizing Network and App Performance With ThousandEyes on Meraki MX*

(Video link: <https://www.youtube.com/watch?v=LgTBJFuFaZ0&list=PLtqMY6nG16HgNz1sLYWy-pmpesHlgwOIS>)



Each deployment method has its own specific steps and considerations. Hands-on experience with these different deployment options will be valuable for the exam.

## Resources

- [Application Hosting on the Cisco Catalyst 9000 Series Switches White paper](#)
- [Installing Enterprise Agents on Cisco Routers with Docker](#)
- [Meraki MX ThousandEyes Configuration Guide](#)
- [Configuring Test Settings](#)
- [Configuring an Enterprise Agent to Use a Proxy Server](#)
- [Agent Settings](#)

## Sample Questions

### 2.1 Question 1

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

### 2.1 Question 2

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

## 2.1 Question 3

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

## 2.2 Endpoint Agent Deployment

 Cognitive Level: Remember

Describe endpoint agent deployment at scale across the enterprise on end-user devices (Windows, Mac, and Room OS)

### Overview

ThousandEyes Endpoint Agents can be deployed at scale across enterprise environments using various methods tailored to specific operating systems and business scenarios.

### Windows Deployment

**Active Directory Environments:** For enterprises using Active Directory, deploying the ThousandEyes Endpoint Agent and browser extension via [Group Policy Objects \(GPOs\)](#) is typically the most efficient approach. This method ensures centralized management and streamlined updates for domain-joined Windows machines.

**Cloud-Managed Environments:** If your organization leverages cloud-based device management, [Microsoft Intune](#) provides the flexibility to deploy and manage Endpoint Agents effectively across diverse Windows devices, including those in remote or hybrid work setups.

### Cisco Secure Client Deployment

Organizations already utilizing Cisco Secure Client (formerly AnyConnect) across their network can leverage this existing infrastructure for streamlined Endpoint Agent deployment. This offers flexibility in deployment points, including:

- **VPN Headends**
- **Firewalls**
- **Identity Services Engine (ISE)**

This approach can be particularly advantageous for managing deployments across diverse operating systems where Cisco Secure Client is already in place.

### Mac Deployment

For Mac-centric environments or those with a mix of Windows and macOS devices, [Munki](#) is the recommended and fully supported option for large-scale Endpoint Agent deployment.

 While other tools like JAMF can be used for deployment, Munki is the officially supported option. Other deployment methods, including JAMF, are supported on a best-effort basis for any issues faced during deployment.

## RoomOS Deployment

If your organization uses Cisco video conferencing systems, you can easily deploy the ThousandEyes Endpoint Agent to RoomOS devices through the native integration with [Webex Control Hub](#). This integration allows administrators to:

- Activate and manage Endpoint Agents across their organization.
- Enable the integration for all devices or select specific ones.

Deployment is handled within the Webex Control Hub admin console's Devices section.

## Resources

- [Install Endpoint Agent for Windows via Group Policy](#)
- [Guidance for Windows Software Deployment Teams](#)
- [Monitoring Webex Meeting with EPA](#)
- [Secure Client Integration](#)
- [ThousandEyes Deployment Using Cisco Secure Client](#)
- [Deployment ThousandEyes Endpoint Agent Using Cisco Secure Client](#)

## Sample Questions

### 2.2 Question 1

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

### 2.2 Question 2

An administrator has configured a Group Policy Object (GPO) to deploy ThousandEyes Endpoint Agent, but noticed it was not installed on one of the office PCs. What is the most appropriate first troubleshooting 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

## 2.2 Question 3

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


## 2.3 Test Configuration

### Cognitive Level: Apply

Configure tests using tools, such as ThousandEyes and Meraki Insights

## Overview

ThousandEyes uses synthetic tests to monitor network and application performance. These tests run from various vantage points, including Cloud Agents, Enterprise Agents, and Endpoint Agents. This section focuses on configuring Cloud and Enterprise Agent tests.



ThousandEyes tests are organized into layers, similar to the OSI model, to provide a structured view of test results. This layered approach helps correlate information and isolate problems more effectively. For example, you can analyze network-specific metrics in the Network layer view and then correlate them with HTTP metrics from the same test round to determine if a network issue is impacting a web page.



**YouTube: Getting Started with Cloud and Enterprise Agent Test Types**

(Video link: <https://www.youtube.com/watch?v=Qqxd2IGk8P0>)

Primary methods for configuring tests include:

- **Standalone Tests:** Create tests from scratch, configuring all settings manually.
- **Test Templates:** Use pre-configured templates to quickly deploy tests for common use cases.

### Create a new standalone test

- 1 Navigate to **Cloud & Enterprise Agents > Test Settings > Tests**.
- 2 Click **Add New Test**.
- 3 Choose the **Test Layer** and **Test Type**.
- 4 Configure the required settings.
- 5 (Optional) Configure advanced settings.
- 6 Click **Create New Test**.

### Create a new test using a test template

- 1 Navigate to **Cloud & Enterprise Agents > Test Settings > Tests**.
- 2 Click **Add New Test** and select **Add from Template**.
- 3 Choose a template from the **Deploy Template** list.
- 4 Configure the required settings, such as target and agents.
- 5 Click **Deploy**.

For more information on test templates, see the [Test Templates](#) documentation.

Additional resources for test configuration and analysis:

- [Getting Started with Cloud and Enterprise Agent Tests](#)
- [Layers and Test Types](#)
- [Test Result Interpretation](#)
- [Test Settings](#)
- [Using the Path Visualization View](#)
- [Getting Started with Views](#)

Configuration instructions for each test type are linked in the respective sections below.

# BGP Tests

BGP tests operate at the routing layer to monitor BGP prefix reachability and path changes from multiple vantage points on the internet. Key features include:

- Automatically monitor relevant BGP prefixes based on the target of other test types like HTTP Server or Network tests
- Specify a custom prefix and prefix length in CIDR notation to monitor
- Include covered (more specific) prefixes of the configured prefix
- Use ThousandEyes' public BGP monitors which peer with routers all over the internet
- Set up private BGP monitors to peer directly with your own routers for an internal view

## BGP Test Configuration

1. Navigate to **Cloud & Enterprise Agents > Test Settings**
2. Click **Add New Test** > Select **Routing** layer and **BGP**
3. Enter a Test Name and the Target Prefix in CIDR notation (e.g. 192.0.2.0/24)
4. Optionally check **Include covered prefixes** to monitor more specific prefixes
5. Select Public or Private BGP Monitors, or both
6. Configure Alert Rules if desired
7. Click **Create New Test**

The target prefix will now be monitored from the selected BGP vantage points. The BGP Route Visualization view shows the AS path from each monitor to the prefix.

## BGP Test Resources

- [Inside-Out BGP Visibility](#)
- [Using the BGP Route Visualization View](#)

 **YouTube: Configuring BGP Tests**

(Video link: [https://www.youtube.com/watch?v=Tyk-vmxc\\_Fw](https://www.youtube.com/watch?v=Tyk-vmxc_Fw))



# Network Tests



Configure network tests such as TCP/UDP, network characteristics, loss, jitter, and latency

ThousandEyes provides two primary types of network tests: agent-to-server and agent-to-agent. Both test types offer two views:

1. **Overview:** Displays data on packet loss, latency, jitter (mean deviation of latency), path MTU, and bandwidth (agent-to-server and Enterprise Agents only). Agent-to-agent tests also show throughput.
2. **Path Visualization:** Provides a traceroute-like map of each router in the path from agent to target, including IP, MPLS, and routing information about each node and link.

## Agent-to-Server Test Configuration

To configure an Agent-to-Server network test:

1. Navigate to **Cloud & Enterprise Agents > Test Settings**
2. Click **Add New Test** and select **Network Layer** then **Agent-to-Server**
3. Configure the following settings in the Basic Configuration tab:
  - **Target:** Enter a domain name or IP address
  - **Protocol:** Choose TCP or ICMP
  - **Port:** Specify the target port number (for TCP only)
  - **Path Trace Mode:** Select **In Session** to perform path trace within an established TCP session
  - **Agents:** Select Cloud and/or Enterprise Agents to run the test
  - **Alerts:** Enable and configure alert rules as needed
4. Configure Advanced Settings:
  - **Perform bandwidth measurements:** Available for Enterprise Agents only
  - **Perform MTU measurements:** Determines path maximum transmission unit in Path Visualization
  - **Collect BGP data:** Enables BGP Path Visualization View
  - **Transmission Rate:** Option to enforce fixed packet rate
  - **No. of Path Traces:** Adjust the number of path trace packets (1-10)
  - **Ping Payload Size:** Set payload size for End-to-End metric probes (0-1400 bytes)
  - **DSCP selector:** Set the Differentiated Services Code Point for QoS handling

## Agent-to-Agent Test Configuration

To configure an Agent-to-Agent network test:

1. Follow steps 1-2 from the Agent-to-Server configuration
2. In the Basic Configuration tab, set:
  - **Target Agent:** Select a ThousandEyes Enterprise Agent
  - **Agents:** Choose source Enterprise Agents
  - **Direction:** Specify network measurement direction (Source to Target, Target to Source, or Both)
  - **Protocol:** Select TCP or UDP
  - **Enable Throughput:** Option to perform throughput measurements
  - **Path Trace Mode:** Configure as in Agent-to-Server tests
3. In the Advanced Configuration tab, set:
  - **Server Port:** Specify port number on the Target Agent
  - **MSS:** Set Maximum Segment Size (Auto or Manual, 30-1400 bytes)
  - **Collect BGP data:** As in Agent-to-Server tests
  - **Transmission Rate:** As in Agent-to-Server tests
  - **No. of Path Traces:** As in Agent-to-Server tests
  - **Payload Size:** Set packet size for network and throughput measurements
  - **DSCP selector:** As in Agent-to-Server tests

## Network Test Resources

- [Network Test Overview](#)
- [The Dual Origin of Network Test Results](#)

 **YouTube: Configuring Network Tests**

(Video link: <https://www.youtube.com/watch?v=R2UmbzrQzHY>)

# DNS Tests



## Configure DNS tests

ThousandEyes offers three main types of DNS tests:

1. **DNS Server Test:** Measures availability and performance of specific DNS servers
2. **DNS Trace Test:** Traces the full DNS resolution path for a domain
3. **DNSSEC Test:** Verifies DNSSEC signatures and the chain of trust

## DNS Test Configuration

To configure a DNS Server test:

1. Navigate to **Cloud & Enterprise Agents > Test Settings**
2. Click **Add New Test** and select **DNS Layer** then **DNS Server**
3. Configure the following settings:
  - **Domain:** Enter the domain name to query
  - **Record type:** Select the DNS record type (A, AAAA, CNAME, MX, NS, TXT, etc.)
  - **DNS Servers:** Specify target DNS servers or use **Lookup Servers** to populate with authoritative nameservers
  - **Interval:** Set test frequency
  - **Agents:** Select Cloud and/or Enterprise Agents to run the test
  - **Alerts:** Enable and configure alert rules as needed
4. Advanced Settings:
  - **Send Recursive Queries:** Option to send recursive queries to target servers
  - **Network measurements:** Enable associated Network test for target server(s)

## DNS Test Resources

- [DNS Tests](#)



**YouTube: Configuring DNS Tests**

(Video link: <https://www.youtube.com/watch?v=uHP1jc6-2fl>)

# Voice Tests



## Configure voice tests

ThousandEyes provides two types of voice tests:

1. RTP Stream Test: Evaluates VoIP call quality and performance
2. SIP Server Test: Assesses SIP server performance and availability

## RTP Stream Test Configuration

To configure an RTP Stream test:

1. Navigate to **Cloud & Enterprise Agents > Test Settings**
2. Click **Add New Test** and select **Voice Layer** then **RTP Stream**
3. Basic Settings:
  - **Target:** Select an Enterprise Agent or Cloud Agent as the endpoint
  - **Interval:** Set test frequency
  - **Agents:** Choose source agents
4. Advanced Settings:
  - **Server Port:** Specify port for incoming RTP sessions
  - **Codec:** Select codec name and associated bit rate
  - **Duration:** Set test duration in seconds
  - **De-jitter Buffer Size:** Configure buffer size to mitigate delay variations
  - **Collect BGP data:** Enable BGP Path Visualization
  - **No. of Path Traces:** Set number of path trace packets (1-10)
  - **DSCP:** Set Differentiated Services Code Point for traffic prioritization

## SIP Server Test Configuration

To configure a SIP Server test:

1. Follow steps 1-2 from the RTP Stream configuration
2. Basic Settings:
  - **SIP Server:** Enter domain name or IP address of the SIP server
  - **SIP Proxy:** Enable and configure SIP Proxy if needed
  - **Protocol:** Choose TCP or UDP
  - **Port:** Specify SIP service listening port
3. Advanced Settings:
  - **Perform SIP Register:** Enable/disable SIP registration
  - **User:** Set username for SIP registration
  - **Auth User:** Configure alternative username for authentication
  - **Password:** Provide authentication password
  - **Desired status code:** Specify SIP status code for successful test
  - **Verify Headers:** Configure header verification using literal text or POSIX regex

## Voice Test Resources

- [Voice Tests](#)

 **YouTube: Configuring Voice Tests**

(Video link: <https://www.youtube.com/watch?v=vPbHf7Weydo>)

# Web Tests



## Configure web tests

ThousandEyes offers three main types of web tests:

1. **HTTP Server Test:** Measures basic availability and response time of web servers
2. **Page Load Test:** Evaluates full page load experience, including all page elements
3. **Transaction Test:** Simulates multi-step user interactions like logins or checkouts

## Web Test Configuration

To configure a Page Load test:

1. Navigate to **Cloud & Enterprise Agents > Test Settings**
2. Click **Add New Test** and select **Web** as the Layer and **Page Load** as the Test Type
3. Basic Configuration:
  - **URL:** Enter the target URL, domain name, or IP address
  - **Interval:** Set test frequency
  - **Agents:** Select Cloud and/or Enterprise Agents to run the test
4. Advanced Settings:
  - **HTTP Server Timing:** Configure timeout and target time for view
  - **Page Load Timing:** Set timeout and target time for view
  - **Proxy Settings:** Configure direct, agent proxy, or specific proxy
  - **HTTP Authentication:** Set up Basic, NTLMv2, Kerberos, or OAuth authentication
  - **HTTP Request:** Configure HTTP version, window size, SSL version, and more
  - **Custom Headers:** Add custom HTTP headers
  - **HTTP Response:** Set desired status code, verify content, and limit download size
  - **Client Certificate:** Configure client-side SSL certificates

## Web Test Resources

- [Web Layer Tests](#)
- [When to Use a Page Load Test](#)



**YouTube: Configuring HTTP Server Tests**

(Video link: [https://www.youtube.com/watch?v=BTxWMG\\_WTV0](https://www.youtube.com/watch?v=BTxWMG_WTV0))

### **YouTube: Configuring Page Load Tests**

(Video link: <https://www.youtube.com/watch?v=L4bTqPvM1ME>)

### **YouTube: Using Web Test Views**


(Video link: <https://www.youtube.com/watch?v=Tgko5nrRYIo>)

## Sample Questions

### 2.3 Question 1

Refer to the exhibit. Which setting should be used for this network Agent to Server test to prevent firewalls from detecting the test traffic as malicious?

- ☒ A) Path Trace Mode: In Session
- ☒ B) Protocol: TCP
- ☒ C) Port: 80
- ☒ D) Probing Mode: Force SYN



The screenshot displays the configuration interface for a Network Agent to Server Test. The settings are as follows:

- Target:** 10.10.10.10
- Protocol:** TCP (selected from a dropdown menu)
- Port:** 80
- Probing Mode:** Three buttons are shown: "Prefer SACK" (selected), "Force SACK", and "Force SYN".
- Path Trace Mode:** A checkbox labeled "In Session" is currently unchecked.
- Interval:** 2 minutes (selected from a dropdown menu)
- Agents:** Select agent(s) (selected from a dropdown menu)
- Alerts:** A checkbox labeled "Enable" is checked. Below it, a dropdown menu shows "1 of 3 alert rules selected", and a link "Edit Alert Rules" is present.

*Exhibit 2.3-1: Network Agent to Server Test Configuration*

## 2.3 Question 2

Refer to the exhibit. 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

The screenshot displays the 'DNS Server Test Configuration' interface. It features two tabs: 'Basic Configuration' (active) and 'Advanced Settings'. Under 'Basic Configuration', the following settings are visible:

- Domain:** A text input field containing 'internal-domain.com', followed by two dropdown menus set to 'IN' and 'A'.
- Interval:** A dropdown menu set to '10 minutes'.
- Agents:** A dropdown menu showing '1 of 8 selected'.
- DNS Servers:** A list containing 'ns47.domaincontrol.com.' with a close icon, followed by a '1 more' button and another close icon. A 'Lookup Servers' button with a magnifying glass icon is also present.
- Alerts:** A checkbox labeled 'Enable' which is checked, followed by a dropdown menu showing '2 of 5 alert rules selected' and a link to 'Edit Alert Rules'.


*Exhibit 2.3-2: DNS Server Test Configuration*

## 2.3 Question 3

Refer to the exhibit. An engineer is trying to configure a Page Load test and is trying to assign the "east1-agent-1" to run it. What is the reason?

- ☒ 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





*Exhibit 2.3-3: Page Load Test Configuration*

## 2.3 Question 4

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

## 2.3 Question 5

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

## 2.4 Endpoint Agent Tests

 Cognitive Level: Apply

Configure endpoint agent tests in ThousandEyes

### Overview

ThousandEyes Endpoint Agent is an application installed on end-user devices to collect network and application data. It enables IT teams to assess application performance, network connectivity, and system health from the end-user perspective.

### Endpoint Agent Test Types

Endpoint Agents perform two main types of tests to measure and monitor application performance:

- **Synthetic Tests**
  - Scheduled Tests: These tests are executed by Endpoint Agents at regular, predefined intervals without requiring user interaction. They can be configured to monitor performance on both the network (agent-to-server tests) and web (HTTP server tests) layers.
  - Dynamic Tests: Building upon Scheduled Tests, Dynamic Tests allow Endpoint Agents to automatically generate tests for IP addresses and ports detected when an application initiates a connection.
- **Real User Tests:** When users access monitored domains from monitored networks, Endpoint Agents collect and record performance metrics related to in-browser and network performance.

 **YouTube: Getting Started: Endpoint Agent Test Types**

(Video link: <https://www.youtube.com/watch?v=b0B9w1oi5sg>)

In addition to these test types, Endpoint Agents also include a Network Access feature that monitors and records the performance of the Endpoint's network components, such as physical wired or wireless connections, gateways, VPNs, proxies, and DNS servers.

### Synthetic Tests

 **YouTube: Configuring Endpoint Agent Test Types**

(Video link: <https://www.youtube.com/watch?v=UDjltEM2e0>)

## SCHEDULED TESTS

Scheduled tests run at regular intervals from Endpoint Agents, similar to tests run by Cloud and Enterprise Agents. They provide consistent performance data for network and web applications, establishing performance baselines. There are two types of scheduled tests available:

- HTTP Server Tests
- Network (Agent-to-Server) Tests

To configure a scheduled test:

1. Navigate to **Endpoint Agents > Test Settings** in the ThousandEyes interface.
2. Click the **Create New Test** button.
3. Select either **HTTP Server** or **Network** as the test type.
4. Configure the basic settings:
  - Test Name
  - Target (URL or IP address)
  - Interval (how often the test should run)
  - Agent Label (can be used to assign the test to specific Endpoint Agents)
5. Set any advanced options as needed.
6. Click **Create New Test** to save and activate the test.

For HTTP Server tests, you can configure options like the HTTP method, request headers, and authentication if required. Network tests allow you to choose between ICMP or TCP for probing the target.

Scheduled tests measure response time, page load time (HTTP tests), packet loss, latency, and availability. They run automatically without user interaction, providing continuous baseline performance data from end-user devices.

## DYNAMIC TESTS

Dynamic tests, previously called Automated Session Tests, automatically monitor the performance of collaboration applications like Webex, Zoom, and Microsoft Teams. These tests identify and test remote targets based on observed network connections to dynamic remote servers.

To set up a dynamic test:

1. Go to **Endpoint Agents > Test Settings > Synthetic Tests**.
2. Click **Monitor Application**.
3. Select the application you want to monitor (e.g., Webex, Zoom) or create a custom template.
4. Configure the global settings for the application.
5. Review the included tests (which are pre-configured by default).
6. Optionally, adjust individual test configurations if needed.
7. Specify the test interval and any alert rules.
8. Assign the test to a label that contains the desired Endpoint Agents.
9. Click **Review** and then **Deploy Now** to activate the dynamic test.

Dynamic tests measure latency, packet loss, jitter, and TCP connection success. The Path Visualization feature shows

the hop-by-hop path from the Endpoint Agent to the application servers.

Dynamic tests are particularly useful for monitoring applications with changing infrastructure, as they adapt to new connection targets automatically.

 **YouTube: Automated Session Testing Simplifies End User Monitoring**

(Video link: <https://www.youtube.com/watch?v=wXeUlswg434>)

## Real User Tests

Real User Tests, previously called Browser Sessions, capture actual user interactions with websites and web applications. These tests use the ThousandEyes browser extension to collect performance data when users access specified domains.

Real User Tests require two components:

1. **Endpoint Agent:** Core software installed on the user's device.
2. **Browser Extension:** Add-on for Chrome and Chromium-based browsers for in-browser metric collection.

To set up a Real User Test:

1. Navigate to **Endpoint Agents > Test Settings > Real User Tests**.
2. Click **Add New Monitored Domain Set**.
3. Configure the following settings:
  - Domain Set Name
  - Monitored Domains (list of domains to monitor)
  - Excluded Subdomains (optional)
  - Agents (select which Endpoint Agents will participate)
4. Click **Add New Monitored Domain Set** to save the configuration.

The browser extension must also be installed on the Endpoint Agents within the assigned label.

Real User Tests collect page load time, response time, network connectivity metrics, and in-browser performance metrics. Once configured, the Endpoint Agents will automatically collect data when users visit the specified domains, providing insights into real-world performance and user experience.

For more details, refer to the [Real User Tests](#) documentation.

## Network Access

The Network Access layer provides insights into the local network environment, including gateway devices, DNS servers, and VPN servers. It consists of two parts:

1. **Network Topology:** Visualizes devices in use by Endpoint Agents for Real User Tests
2. **Wireless:** Displays information about wireless data (e.g., SSID, signal strength, channel)

Network Access data is continuously collected while an Endpoint Agent is online and accessing websites defined in the

monitored domain set.

## Endpoint Agent Views

ThousandEyes provides several views for examining test results. Each view corresponds to a specific test type and offers unique insights:

### Scheduled Tests

- **Documentation:** [Endpoint Agent Scheduled Tests View](#)
- **Purpose:** Monitor performance of applications or network paths at pre-defined intervals.
- **Data Presentation:**
  - **Timeline:** Up to 30 days of test data.
  - **Map:** Geographical display of agent locations and metrics.
  - **Table:** Detailed breakdowns of metrics.
  - **Path Visualization:** Network path from Endpoint Agent to target.
- **Use Case:** Ensure consistent performance quality over time.

### Dynamic Tests

- **Documentation:** [Endpoint Agent Dynamic Tests View](#)
- **Purpose:** Provide real-time performance data based on user interactions.
- **Data Presentation:**
  - **Timeline:** Up to 30 days of data.
  - **Tabs:** Detailed metrics in Map and Table views.
  - **Path Visualization:** Network topology analysis.
- **Use Case:** Diagnose sporadic issues not captured in scheduled tests.

### Real User Tests

- **Documentation:** [Endpoint Agent Real User Tests View](#)
- **Purpose:** Capture actual user experience and interactions.
- **Data Presentation:**
  - **Timeline:** User interactions over time.
  - **Pages Tab:** List of visited sites with experience scores.
  - **Sessions Tab:** Waterfall view of user actions on a page.
  - **Map:** Geographic analysis of metrics.
  - **Table:** Detailed metric analysis.
  - **Path Visualization:** Network path examination.
- **Use Case:** Understand end-user experience and identify user-affecting problems.

### Local Networks

- **Documentation:** [Endpoint Agent Local Networks View](#)
- **Purpose:** Provide insights into local network environment.
- **Data Presentation:**
  - **Network Topology:** Visual representation of network devices and performance.
  - **Path Visualization:** Network path details.
  - **Performance Metrics:** Various network performance data.
- **Use Case:** Diagnose VPN, DNS, and other local network-related issues.

 **YouTube: Using Endpoint Agent Test Views**

(Video link: [https://www.youtube.com/watch?v=D-M\\_hNg6Z5A](https://www.youtube.com/watch?v=D-M_hNg6Z5A))

# Endpoint Agent Labels

Endpoint Agent tests are assigned to agents dynamically using labels due to the frequently changing nature of local workstations.

To configure Endpoint Agent labels:

1. Navigate to **Endpoint Agents > Agent Settings > Agent Labels**
2. Click **Add New Label**
3. Define the label name and choose a color
4. Use the **Filter** drop-down lists to configure a filter for the label
5. Save changes

## Using Wildcards

You can use the **wildcard** option while configuring a label for Hostname, SSID, and Username. Example:

DOMAINNAME\*

 **YouTube: Using Endpoint Agent Labels**

(Video link: <https://www.youtube.com/watch?v=BAwJ9-j5EgM>)

# Proxy Configuration for Endpoint Agent Tests

In enterprise environments, it's common to route traffic through proxy servers. ThousandEyes allows you to configure proxy settings specifically for Endpoint Agent scheduled tests:

1. Go to **Endpoint Agents > Agent Settings > Proxy Settings**.
2. Click **Add New Proxy Config**.
3. Choose between Static or PAC file configuration.
4. Enter the necessary proxy details (host, port, authentication if required).
5. Optionally, configure a bypass list for direct connections to specific destinations.

These proxy settings apply only to scheduled tests and do not affect the agent's connection to ThousandEyes services.

## Resources

- [Endpoint Agent Overview](#)
- [Endpoint Agent Test Settings](#)
- [Getting Started with Endpoint Agents](#)
- [Real User Tests](#)
- [Single Agent View](#)

# Sample Questions

## 2.4 Question 1

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




Exhibit 2.4-1: Endpoint Agent Test Configuration

## 2.4 Question 2

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

## 2.4 Question 3

What type of endpoint agent test will gather browser activity?

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

## 2.4 Question 4

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

## 2.5 Synthetic Web Tests



Cognitive Level: Remember

Describe the purpose, implementation, and limitations of synthetic web tests

### Overview

ThousandEyes transaction tests are synthetic web tests that simulate multi-step user journeys through web applications. They go beyond single page load tests by interacting with the target application to test complete workflows. This makes transaction tests ideal for monitoring key business processes in web apps, such as:

- Logging in and accessing data in SaaS applications
- Searching for a product, adding it to the cart, and checking out on an ecommerce site
- Scheduling a meeting and joining it in a web conferencing app

Some key features and concepts of ThousandEyes transaction tests include:

- The ThousandEyes Recorder IDE for recording user actions and generating test scripts
- Custom markers to measure the duration of specific steps within the overall transaction
- Screenshots captured at key points for visual validation and troubleshooting
- Integration with APIs to incorporate data from other systems into the test

It's important to note that while transaction tests can make API calls as part of the user journey, dedicated API tests are more suitable for solely testing APIs directly.

Transaction tests should run at an appropriate frequency to catch issues quickly without overloading the application. Intervals of 5-15 minutes are common. Very frequent tests may require special tuning.

### Hands-on Activities

#### Activity 1: Get started with the IDE recorder

1. Go to <https://docs.thousandeyes.com/product-documentation/getting-started/getting-started-with-transactions#getting-started-with-the-recorder-ide> and follow the instructions to setup the IDE recorder.
2. Choose a user flow that you can record, for example, the login process to office 365 or any other web application that requires a login process (avoid MFA in the getting started process).
3. Set the transaction script.
4. Observe the test results.

For a visual guide on configuring transaction tests, watch the following video:

 **YouTube: Configuring Transaction Tests**

(Video link: <https://www.youtube.com/watch?v=SmCJCeTsyp8>)

## Activity 2: Create a script using a template

1. Go to the Github Transaction script repository: <https://github.com/thousandeyes/transaction-scripting-examples>
2. Select a template and adapt it to your needs.
3. Set the transaction script and assign it to an agent.
4. Observe the test results.

Check out the following video for a walkthrough on using transaction test views:

 **YouTube: Using Transaction Test Views**

(Video link: [https://www.youtube.com/watch?v=H\\_axZ97Pn5E](https://www.youtube.com/watch?v=H_axZ97Pn5E))

## Resources

- [Metrics from Synthetic Tests](#)
- [Managing Synthetic Tests](#)
- [Monitoring an Application Using Synthetic Tests](#)
- [Configuration Options for Synthetic Tests](#)
- [Transaction Scripting Examples on GitHub](#)
- [ThousandEyes Recorder Documentation](#)

## Sample Questions

### 2.5 Question 1

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 > One Drive - 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

Create New

```
1 import { By, Key, until } from 'selenium-webdriver';
2 import { driver, markers, credentials, downloads, transaction, test } from 'thousandeyes';
3
4 runScript();
5
6 async function runScript() {
7
8   const settings = test.getSettings();
9
10  // Load page
11  await driver.get(settings.url);
12  await driver.takeScreenshot();
13
14  // Search
15  markers.start('SearchForWebdriver');
16  await driver.findElement(By.name('q')).sendKeys('webdriver', Key.RETURN);
17  await driver.takeScreenshot();
18  markers.stop('SearchForWebdriver');
19
20  // Wait for full page load
21  await driver.wait(until.titleIs('webdriver - Google Search'), 1000);
22  await driver.takeScreenshot();
23 };
24
```

Install ThousandEyes Recorder on your machine

*Exhibit 2.5-1: Transaction Script Example*

## 2.5 Question 2

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

## 2.5 Question 3

True or False: ThousandEyes Synthetic Tests eliminate the need for any real user monitoring on your online learning platform.

- ☒ A) True
- ☒ B) False

## 2.6 Web Authentication Methods

 Cognitive Level: Apply

Implement common web authentication methods when testing web applications.

### Overview

Enterprise web applications typically require authentication for access. Implementing and testing various authentication methods is crucial for security and functionality. Comprehensive web application testing should support multiple authentication mechanisms to effectively evaluate HTTP requests, page loads, transactions, and APIs.

#### Preparing for the Exam

While the exam focuses on ThousandEyes-specific authentication knowledge, a general understanding of web authentication principles is expected. This includes familiarity with protocols like SAML, OAuth, and OpenID Connect, as well as concepts such as token-based authentication and session management.

### Key Concepts

#### Single Sign-On (SSO)

Enables users to access multiple applications with one set of login credentials.

#### Multi-Factor Authentication (MFA)

Requires users to provide two or more verification factors for enhanced security.

#### Security Assertion Markup Language (SAML)

A standard for exchanging authentication and authorization data between an identity provider and a service provider, enabling SSO.

#### Basic Authentication

Transmits usernames and passwords in plain text, encoded with Base64.

#### NTLM Authentication

A suite of Microsoft security protocols for user authentication.

#### Kerberos Authentication

Uses secret-key cryptography and a trusted third party for secure authentication.

### OAuth Authentication

An open standard for access delegation, granting limited access to user data without sharing passwords.

### OpenID Connect

An authentication layer on top of OAuth 2.0 for verifying user identity and obtaining basic profile information.

## Authentication Methods and Test Types

The relationship between authentication methods and test types in ThousandEyes is important to understand for effective monitoring. Different test types support various authentication methods:

- **HTTP Tests:** Support Basic, NTLM, Kerberos, and OAuth authentication methods. For two-step OAuth scenarios, custom headers can be configured to accommodate more complex authentication schemes.
- **Page Load Tests:** Support the same authentication methods as HTTP Tests, including Basic, NTLM, Kerberos, and OAuth. Note that SSO and MFA are not directly supported.
- **Transaction Tests:** Offer the most flexibility, allowing scripted interactions to handle complex authentication flows, including multi-step processes and dynamic challenges.
- **API Tests:** Authentication settings are configured in the API Step Builder for each step, which is different from other test types. API Tests support Basic (RFC 7617) and Bearer (RFC 6750) HTTP authentication schemes.

## Preferred Authentication Methods in Different Scenarios

The choice of authentication method often depends on the specific requirements of the application and the security needs of the organization:

1. **Single Sign-On (SSO):** Preferred in enterprise environments with multiple applications, as it improves user experience and centralizes access control. ThousandEyes supports SSO testing, which is crucial for ensuring seamless access across integrated systems. For more information on implementing SSO in ThousandEyes tests, check out the guide on [Transaction Test SSO Support](#).
2. **Multi-Factor Authentication (MFA):** Recommended for applications handling sensitive data or requiring high security. While MFA enhances security, it can complicate automated testing. ThousandEyes provides ways to handle some forms of MFA in transaction tests, as detailed in the [Working with Secure Credentials](#) documentation.
3. **API Key Authentication:** Often used for machine-to-machine communication and is well-suited for API tests. ThousandEyes API tests can be configured to use API keys, as shown in the [API Test Use Cases](#) documentation.
4. **OAuth and OpenID Connect:** Preferred for applications that need to access user data from third-party services without handling passwords directly. These methods are particularly useful in transaction tests simulating user interactions across multiple services.

## ThousandEyes-Specific Implementations

**Note:** When implementing authentication in ThousandEyes, always refer to the most up-to-date documentation, as features and supported methods may evolve. The [Test Settings for Page Load and Transaction Tests](#) provides comprehensive information on configuring authentication for different test types.

Platform-specific considerations:

- **Credential Management:** ThousandEyes provides a secure credential store for managing authentication information used in tests. This feature is unique to the platform and crucial for maintaining security in automated testing scenarios.
- **Agent-Specific Settings:** Some authentication methods, like Kerberos, require specific configuration on ThousandEyes agents. These settings are managed through the ThousandEyes interface and may differ from standard implementations.
- **Custom Scripting:** For complex authentication flows, ThousandEyes transaction tests allow custom JavaScript to handle unique scenarios that may not be covered by standard authentication methods.

To learn more about configuring these ThousandEyes-specific features, consult the [Working with Test Settings](#) guide.

## Hands-on Activities

### Activity 1: Explore Authentication methods available for web tests

1. If you haven't, sign up for a ThousandEyes trial here: <https://www.thousandeyes.com/certificationsignup>.
2. Go to **Cloud & Enterprise Agent > Test settings > Start monitoring > Start with a single test**.
3. Select **Web**.
4. All authentication settings will be available in the **Advanced Settings** tab.
5. What are the schemes available for HTTP Server, Page Load and Transaction Tests?
6. What are the parameters required by each different authentication method (Basic, NTLM, Kerberos, OAuth)?
7. Configure an HTTP Server test towards the ThousandEyes API (<https://developer.cisco.com/docs/thousandeyes/overview/>).
  - For example, the agents endpoint will return a list of ThousandEyes Agents available: URL:  
`https://api.thousandeyes.com/v7/agents`.
  - In **Advanced Settings > Scheme: None > Custom Headers > Root Request > Authorization: Bearer your-oauth2-bearer-token**
  - Go back to the **Basic Configuration** settings and after selecting the interval and agents click on **Create New Test**.
8. In **Views**, select your test name and review the results of the test for web and network layers and for each of the metrics available as well as the path visualization view.



HTTP AUTHENTICATION

Scheme

None
Basic
NTLM
Kerberos

HTTP REQUEST

HTTP Version \*

Prefer HTTP/2
HTTP/1.1 only

Follow Redirects \*
☒ Enable

Device Emulation

Desktop Large
1440 x 900 px

User Agent

Default

Agent ID Strategy

Add x-thousandeyes-agent He...

Custom Headers

Root Request

Authorization: Bearer your-oauth2-bearer-token

Figure 2.6-1: Web Authentication Methods Activity

## Resources

- [Transaction Test SSO Support](#)
- [Working with Test Settings](#)
- [Test Settings Page Load Transaction](#)
- [Working with Secure Credentials](#)
- [API Test Use Cases](#)

## Sample Questions

### 2.6 Question 1

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. What must be done to meet the requirements? (Select 2)

- ✓ 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

HTTP AUTHENTICATION

Scheme

None

Basic

NTLM

Kerberos

OAuth

HTTP REQUEST

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...

*Exhibit 2.6-1: HTTP Authentication Options*

## 2.6 Question 2

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.

## 2.6 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) 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.

## 3.1 Network Issue Diagnosis

 Cognitive Level: Analyze

Diagnose network issues, such as packet loss, congestion, routing, and jitter using collected data

### Overview

Troubleshooting network issues requires a methodical approach and strong analytical skills. This section focuses on diagnosing common network problems like packet loss, congestion, routing issues, and jitter using data collected from various sources like ThousandEyes, Meraki Dashboard, Catalyst Center, and SD-WAN Manager.

Understanding the data collected, the metrics, and their impact on user experience is crucial for effective diagnosis. For example, high latency can lead to slow application responsiveness, while jitter can disrupt voice and video calls.

When facing an issue, start by determining if it's application or network-related. Then, leverage collected data to identify the problem area and its potential causes.

### Key Concepts

This section defines key network metrics, explores their common causes, and explains their impact on user experience. Understanding these concepts is crucial for effective network diagnostics.

#### Packet Loss

Packet loss occurs when one or more data packets traveling across a network fail to reach their destination. This compromises the integrity of the transmitted data and can lead to various issues depending on the application and protocol being used. Packet loss is typically measured as a percentage of packets lost relative to the total number of packets sent.

#### COMMON CAUSES OF PACKET LOSS

- **Network Congestion:** When network traffic exceeds the capacity of network devices, packets may be dropped. This is common in scenarios with high-bandwidth applications like video streaming or large file transfers.
- **Transmission Errors:** Problems with the physical media, such as signal degradation, noise, or interference, can corrupt packets and lead to loss. The impact of these factors varies depending on the type of transmission media used (e.g., fiber optic cables are less susceptible to interference than copper wires).
- **Device Misconfiguration:** Incorrect settings on network devices, including firewalls, routers, and switches, can cause packets to be dropped. This could involve misconfigured access control lists (ACLs), Quality of Service (QoS) policies, or routing rules.
- **Routing Changes:** When network routes change, packets may be lost if they are no longer directed to a valid destination. This can happen during network maintenance, outages, or configuration updates.
- **Hardware Failures and Software Bugs:** Malfunctions in network hardware (e.g., faulty network interface cards) or software bugs in network device operating systems can also contribute to packet loss.

## IMPACT OF PACKET LOSS

- **Choppy audio and video during streaming or video calls**
- **Frozen video frames**
- **Interrupted streaming services**
- **Failed file downloads**
- **Overall reduction in user productivity**

## Latency

Network latency refers to the time it takes for a packet to travel from one point in the network to another. It's the delay experienced by data as it traverses the network and is commonly measured in milliseconds.

### FACTORS CONTRIBUTING TO LATENCY

- **Signal Propagation Delay:** This inherent delay depends on the distance the signal must travel and the speed of transmission through the chosen medium. Signals travel faster through fiber optic cables compared to copper wires. Geographic distance significantly influences latency as signals take longer to travel over longer distances.
- **Network Device Processing:** Routers and switches introduce a small delay as they process packet headers and determine the appropriate forwarding path. The complexity of the device's configuration and the volume of traffic it handles can impact processing time.
- **Traffic Load and Queuing:** Packets may experience delays while waiting in queues due to high traffic load or Quality of Service (QoS) prioritization. QoS mechanisms can prioritize certain types of traffic (e.g., voice over data), which can delay other packets.
- **Inefficient Routing:** Each hop between network devices adds to the overall latency. Poorly designed routing paths, such as those with unnecessary hops or congested links, contribute to higher latency.

### IMPACT OF LATENCY

- **Sluggish Application Responsiveness:** Applications become slow to respond to user actions, impacting productivity and accuracy.
- **Decreased Voice and Video Quality:** Calls suffer from delays, making conversations difficult. This can lead to users talking over each other or experiencing long silences.
- **Impact on Business Operations:** In time-critical environments, such as financial trading, high latency can cause lost revenue due to delayed transactions and slow reaction to market changes.

## Jitter

Jitter is the variability in the time it takes for data packets to be forwarded from their source to their destination. Ideally, packets should arrive at a consistent pace. However, when a network experiences high jitter, packets can arrive out of order.

## FACTORS CONTRIBUTING TO JITTER

- **Network Congestion:** When too many data streams overwhelm the available bandwidth of a network, packets can be delayed, causing jitter. This is especially noticeable in real-time applications like video conferencing and VoIP calls.
- **Bursty Traffic:** Sudden spikes in network traffic can cause temporary jitter as the network adjusts to the changing load. This can occur during peak usage periods or when large data transfers are initiated.
- **Wireless Interference:** In wireless networks, interference from other signals can cause delays in packet delivery, leading to jitter. Interference from other wireless devices, microwave ovens, or physical obstructions can contribute to this.
- **Hardware and Software Issues:** Faulty or outdated networking hardware can lead to inconsistent packet delivery times, resulting in jitter. This might include issues with network interface cards, drivers, or firmware on network devices.

## IMPACT OF JITTER

High jitter primarily impacts voice and video communications, making conversations confusing and difficult to understand. Users may experience:

- **Distorted audio**
- **Choppy video**
- **Dropped calls**

## Routing Issues

Border Gateway Protocol (BGP) is the standard routing protocol used to exchange routing information across the internet. It enables the sharing of reachability information between autonomous systems (AS), which are groups of networks operated under a single administrative organization. BGP allows networks to learn about available routes to different parts of the internet and make decisions about how to forward traffic.

When troubleshooting network issues, analyzing the reachability to prefixes of interest is crucial. Reachability must be considered from both sides:

1. How you can route to the target.
2. How the target can route back to you.

## COMMON ROUTING ISSUES

- **Reachability Issues:** This occurs when there is no correct route to a prefix. This could happen if the originating AS stops advertising the prefix due to a failure or misconfiguration, or if a malicious AS falsely advertises the prefix.
- **Convergence and Unstable Conditions:** This happens when a route continually toggles between available and unavailable states, or when the route information changes frequently due to underlying infrastructure changes. This can lead to intermittent connectivity issues as routes change and traffic is rerouted.

# Resources

- [ThousandEyes Metrics: What Do Your Results Mean?](#)
- [Using the BGP Route Visualization View](#)
- [Viewing Data](#)
- [Monitoring BGP Routes with ThousandEyes](#)
- [Finding the Root Cause of Loss and Latency in Internet-Facing Applications](#)
- [4 Real BGP Troubleshooting Scenarios](#)
- [ThousandEyes BGP Monitors](#)
- [BGP Route Leak](#)

## Sample Questions

### 3.1 Question 1

Users at a remote corporate site (site 30 or "s30") 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?




Exhibit 3.1-1: Before Issue




Exhibit 3.1-2: After Issue

View in [ThousandEyes](#)

- ✓ 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.

### 3.1 Question 2

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






Exhibit 3.1-3: Before Incident #1



Exhibit 3.1-4: Before Incident #2




Exhibit 3.1-5: During Incident #1




Exhibit 3.1-6: During Incident #2

View in [ThousandEyes](#)

- ✓ A) Involve the Voice team as the RTP test does not return any relevant results for the agent located on site 20 “s20”
- ✓ 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.

## 3.2 End-Device Network Issue Diagnosis

### Cognitive Level: Analyze

Diagnose end-device network issues, such as issues with a default gateway, local network, DNS server, proxy, VPN gateway, wireless, and real-time streaming using collected data

### Overview

Maintaining a reliable and secure network requires diagnosing network issues on traditional end-devices, such as computers and smartphones. ENNA v1.0 exam emphasizes the endpoint agent, its capabilities, metrics, and troubleshooting techniques to identify and resolve network problems impacting applications from the end-user's perspective. These issues can affect users working from various locations, including the office and home. While IP-connected devices such as sensors, cameras, and IoT devices are not within the current scope (v1.0), they may be included in future versions of the exam.

### Key Concepts

#### Diagnosing End-Device Network Issues

This section focuses on analyzing data to troubleshoot network problems originating from end devices. We'll cover how to leverage information from ThousandEyes Endpoint Agents to identify and resolve these issues.

We can categorize the issues into three main areas:

1. Local network problems
2. Web performance problems
3. Broader network problems (beyond the local network)

#### LOCAL NETWORK PROBLEMS

Local networks can be problematic for anyone, but especially for remote workers. Endpoint Agents provide a unique perspective into the network conditions experienced by remote workers and devices.

They gather crucial data, including:

- ICMP ping and traceroute information to network elements (proxies, VPNs, gateways) over the past 24 hours.
- Network profiles capturing connection details.

Data is collected in focused bursts every 5 minutes.

Potential Causes:

- Gateway malfunctions
- Bandwidth limitations
- Endpoint performance bottlenecks (CPU, memory)
- Poor Wi-Fi signal strength
- General Wi-Fi connectivity problems

Troubleshooting:

- Verify endpoint network settings
- Examine network interface details (speed, errors)
- Use network diagnostic tools (ping, traceroute) within ThousandEyes to gain historical context and path visualization
- Analyze logs and error messages from end devices to identify errors, configuration problems, or hardware failures

## **WEB PERFORMANCE ISSUES**

Potential Causes:

- Browser cache/cookie issues
- Web application errors or outages
- API limitations or failures
- DNS server problems

Troubleshooting:

- Within ThousandEyes, correlate web performance metrics with underlying network conditions
- Analyze HTTP error messages for clues
- Clear browser cache and cookies, and verify firewall/proxy configurations

## **BROADER NETWORK PROBLEMS (BEYOND THE LOCAL NETWORK)**

Potential Causes:

- ISP outages
- Network congestion
- Physical infrastructure issues (fiber cuts, power outages)
- DNS server problems
- Routing anomalies

Troubleshooting:

- In ThousandEyes, pinpoint the time the issue began
- Use the network visualization to isolate problematic nodes by filtering for an affected agent
- Differentiate between end-to-end packet loss (at the agent) and forwarding loss occurring along the network path

## ThousandEyes Endpoint Agent Metrics

Familiarize yourself with the metrics collected by ThousandEyes Endpoint Agents. This knowledge is essential for:

- Problem Context: Understanding the type and scope of data collected helps you interpret the situation.
- Root Cause Identification: By correlating metrics across different layers (e.g., HTTP response time against network latency), you can pinpoint whether the issue stems from the network, application, or endpoint itself.
- Troubleshooting Remote Workers: Determine if a remote worker's issues are due to local network congestion, Wi-Fi signal strength, device resource constraints, or problems with the ISP or wider internet.

The table below provides an overview of the metrics and views available for data collected by Endpoint Agents.

Table 1: Endpoint Agent Metrics

| Metric Type | Description                                                                          | Relevant Views                                                  |
|-------------|--------------------------------------------------------------------------------------|-----------------------------------------------------------------|
| Network     | Packet loss, latency, jitter, connection failures                                    | Real User Tests, Scheduled Tests, Dynamic Tests                 |
| System      | CPU load, memory usage                                                               | Local Networks, Real User Tests, Scheduled Tests, Dynamic Tests |
| VPN         | VPN loss, VPN latency                                                                | Local Networks, Real User Tests, Scheduled Tests, Dynamic Tests |
| DNS         | DNS server loss, DNS server latency, domain resolution time                          | Local Networks                                                  |
| Gateway     | Gateway loss, gateway latency                                                        | Local Networks                                                  |
| Proxy       | Proxy loss, proxy latency                                                            | Local Networks                                                  |
| Wireless    | Signal quality, throughput, retransmission rate, roaming events, channel swap events | Local Networks                                                  |
| HTTP        | Availability, response time, throughput                                              | Scheduled Tests                                                 |

Refer to our official product documentation for a more detailed breakdown: [Data Collected by Endpoint Agent](#)

### Additional things to note

- **Test Capabilities:** Know the types of tests Endpoint Agents can perform and their limitations.
- **Test Prioritization:** Understand how to prioritize tests, especially in resource-constrained environments.
- **Authentication and Security:** Be comfortable configuring HTTP authentication and SSL verification settings for tests.
- **Network Visualization:** Become proficient in interpreting the symbols and data displayed in ThousandEyes' network visualization views.

## Resources

- [Troubleshooting Endpoint Agent issues](#)
- [Endpoint Agent Views Reference](#)
- [Data Collected by Endpoint Agent](#)
- [What is jitter? An article on Meraki portal](#)
- [WiFi and LAN Monitoring](#)
- [VPN Monitoring](#)

## Sample Questions

### 3.2 Question 1

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 ?




Exhibit 3.2-1




Exhibit 3.2-2






Exhibit 3.2-3




Exhibit 3.2-4

- ☒ 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.

## 3.2 Question 2

The Endpoint stopped appearing online after it was moved to another network.

|                          |                                             |                                                                                   |                                                                                   |                                                                                   |         |                                                                                   |                     |                |               |               |              |           |
|--------------------------|---------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|---------|-----------------------------------------------------------------------------------|---------------------|----------------|---------------|---------------|--------------|-----------|
| <input type="checkbox"/> | OATROSHC-WIN10-<br>OATROSHC-WIN10-loatrosch |  |  |  | 1.191.3 |  | Almere, Netherlands | 37 minutes ago | 173.38.220.55 | 192.168.2.139 | 113 days ago | Advantage |
|--------------------------|---------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|---------|-----------------------------------------------------------------------------------|---------------------|----------------|---------------|---------------|--------------|-----------|

Exhibit 3.2-5

The customer reviewed the endpoint logs but did not identify anything suspicious.

```
2024-05-07 13:20:12.452 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:206 - RECONNECT is currently scheduled or running, ignoring schedule attempt for CONNECT
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService@DefaultAgentSettingsService.cpp:496 - Checkin failed with error code: 1
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService@DefaultAgentSettingsService.cpp:558 - Calling connect on persistent connection
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService@DefaultAgentSettingsService.cpp:576 - Will consider polling for scheduled tests in 30s
2024-05-07 13:20:14.468 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:492 - Scheduling checkin in 10m
2024-05-07 13:20:14.468 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.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.DefaultPersistentConnection@DefaultPersistentConnection.cpp:206 - RECONNECT is currently scheduled or running, ignoring schedule attempt for CONNECT
2024-05-07 13:20:24.753 DEBUG [3068.8248] system.EtwTraceConsumer@DefaultEtwTraceConsumer.cpp:108 - Flushed ETW trace buffers to file. path="C:\ProgramData\ThousandEyes\Endpoint Agent\etwtraces\ThousandEyes Endpoint Agent.etl", size=320kB
2024-05-07 13:20:24.765 DEBUG [3068.8248] net.SocketMonitoring@SocketMonitoringService.cpp:335 - Processed ETW event batch: published=2, stale=213, deduped=0, duplicatesDueToCoalesceInterval=0, skipped=41
2024-05-07 13:20:24.812 DEBUG [3068.7712] scheduler.task.statscommitter@UsageStatsCommitter.cpp:43 - Starting UsageStatsCommitter task
2024-05-07 13:20:24.813 DEBUG [3068.7712] tests.scheduledtestsmetric@DefaultScheduledTestsMetricsService.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.statscommitter@UsageStatsCommitter.cpp:95 - Updated usage stats object with 69 entries.
2024-05-07 13:20:25.109 DEBUG [3068.2148] scheduler.simpleDefaultTaskScheduler.cpp:173 - Scheduled task: te:agent:UsageStatsCommitter@0x1, delay=60000 ms
2024-05-07 13:20:28.498 DEBUG [3068.2148] net.WinHttpClient@WinHttpClient.cpp:2115 - Request to: wss://c1.eb.thousandeyes.com/relay/connect timed out: 12002: The operation timed out
2024-05-07 13:20:28.499 DEBUG [3068.8248] net.WinHttpClient@WinHttpClient.cpp:2115 - Request to: https://c1.eb.thousandeyes.com/status.json timed out: 12002: The operation timed out
2024-05-07 13:20:28.499 DEBUG [3068.4352] net.DefaultWebSocketStompConnection@DefaultWebSocketStompConnection.cpp:135 - WebSocket connection wait: 4210ms
2024-05-07 13:20:28.499 ERROR [3068.4352] net.DefaultWebSocketStompConnection@DefaultWebSocketStompConnection.cpp:148 - asyncOpen (WebSocket): Input/output error [generic:5]
2024-05-07 13:20:28.499 ERROR [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:317 - asyncOpen: Input/output error [generic:5]
2024-05-07 13:20:28.499 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:485 - Running post-run checks for RECONNECT, next=NONE, retry=true
2024-05-07 13:20:28.499 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 13:20:28.499 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:308 - Set schedule delay=1500ms, jitter=323ms, nextDelay=2250ms
2024-05-07 13:20:28.499 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@CheckinService.cpp: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 13:20:28.771 DEBUG [3068.8072] api.userSpaceProxyClient.cpp:50 - Relaying HTTP request https://c1.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@DefaultWebSocketStompConnection.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
```

Exhibit 3.2-6

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?**

- ☒ 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.

## 3.3 Web Application Performance Issues

 Cognitive Level: Analyze

Diagnose web application performance issues using collected data such as browser waterfalls

### Overview

This section explores concepts and tools that provide insights into web application performance problems.

### Key Concepts

To effectively troubleshoot issues, it's essential to distinguish between application-related and network-related issues. This knowledge equips us to optimize both network performance and application functionality, ensuring a seamless user experience.

### Waterfall Charts

A waterfall chart provides a visual representation of how a browser interacts with web page objects during the loading process. It displays the timeline of each object's download, including HTML, CSS, JavaScript files, images, and other resources.

The chart starts with the initial request to load the webpage and shows the sequential loading of objects over time. Each object's bar length represents the time taken to download it, while the horizontal axis represents time.

By analyzing a waterfall chart, developers can identify bottlenecks, such as slow-loading resources or dependencies, and optimize the webpage's performance. This insight helps improve user experience by reducing page load times and enhancing overall responsiveness.

### Indicators for Analysis

Key indicators to look out for when analyzing performance:

- **Response Time:** The time it takes for a web server to respond to a user's request.
- **Load Time:** The time it takes for all elements of a web page to fully load and display to the user.
- **Availability:** Whether the web application is accessible to users or experiencing downtime.
- **HTTP Status Codes:** These codes provide insights into the success or failure of web requests, helping to pinpoint potential issues (e.g., server errors, user errors).

ThousandEyes provides the views to correlate these indicators with network data to provide a comprehensive understanding of web application performance bottlenecks.

# Document Object Model (DOM)

The DOM represents an HTML document's structure as a tree of objects called nodes. Nodes have parent-child relationships, with some embedded in the main HTML (e.g., text, scripts) and others referencing external resources.


Learn more about waterfall charts and the DOM model in the [ThousandEyes documentation](#).

## Using Waterfall Chart Metrics for Troubleshooting

Among the various test types offered by ThousandEyes, only Page Load and Transaction tests generate a waterfall chart, providing a detailed visualization of how a browser interacts with web page objects during the loading process.

This section provides a simplified approach to troubleshooting web performance issues using waterfall chart metrics from ThousandEyes Page Load and Transaction tests. By understanding these metrics and following the proposed troubleshooting path, you can gather evidence to identify the problem and find a solution.

Start your troubleshooting process by asking questions based on the specific metric that shows abnormal change, as this can negatively impact user experience. Begin at the network layer and work your way up to the application layer, using the provided questions as a guide. After answering these questions, determine the appropriate next steps. Remember that each metric can be directly affected by common issues.



### Connect

**Metric meaning:** The time to establish a TCP handshake with the Target Server.

#### Questions to Ask:

- Is the TCP handshake established correctly?
- What do the network metrics collected by the agent indicate?

#### Possible Causes:

- Routing problems (path updates, router misconfiguration)
- Packet loss
- Internet outage
- Application server outage

## DNS

**Metric meaning:** The duration to resolve a domain record to an IP address. By default, BrowserBot does not cache DNS records at startup.

### Questions to Ask:

- Is DNS resolving properly?
- How long is it taking to resolve?

### Steps to Take:

1. Identify the DNS servers configured for the agent and who manages them.
2. Test if the DNS server is responding by directly running queries to it with dig or nslookup.
3. Test if the DNS is reachable.

### Possible Causes:

- DNS server outage
- DNS server configuration problem
- Network path to the DNS server affected
- DNS Hijack

## SSL

**Metric meaning:** The duration of SSL/TLS negotiation.

### Questions to Ask:

- Is SSL/TLS negotiation completing successfully?
- Are there any SSL/TLS errors?

### Steps to Take:

1. Collect information about the SSL/TLS handshake errors.
2. Analyze if the trust relationship is established between server and user (Does the agent trust the CA that issued the certificate that the server is presenting?).
3. Check if the certificate is presenting the full root chain.
4. Verify if the server is using a self-signed certificate.
5. If the certificate is not from a well-known CA, ensure the CA certificate is installed on the agent.

### Possible Causes:

- Application server configuration problem
- Expired or invalid certificates

## Send

**Metric meaning:** The duration in which the browser successfully sends a request to the server. Also known as Time to First Byte.

**Questions to Ask:**

- Is the browser sending the request correctly?
- How much time did it take?

**Steps to Take:**

- Collect information about HTTP code errors to determine if this is a user (400 errors) or a server error (500 errors).

**Possible Causes:**

- Incorrect proxy settings or misconfigured network settings
- Browser Malformed Requests
- Cross-Origin Resource Sharing (CORS) Errors
- Ad Blockers or Browser Extensions

## Wait

**Metric meaning:** The duration between the completion of a browser's SEND request and receipt of the first byte of a server's response.

**Questions to Ask:**

- How much time did it take to hear back from the server?
- What is the server's performance?

**Steps to Take:**

1. Correlate the web metrics to the network metrics such as latency and packet loss.
2. Identify the nodes in the path to the destination (e.g., CDNs, load balancers, firewalls).

**Possible Causes:**

- Server processing time
- Misconfigured or suboptimal server settings
- Resource starvation (high CPU, memory, or disk I/O usage)
- Network latency (physical distance between client and server)
- Bandwidth throttling or limitations
- CDN issues
- Ineffective server-side caching
- Client-side issues (slow DNS resolution, misconfigured network settings, outdated hardware)

## Receive

**Metric meaning:** The time between the first byte of the server response to the last byte of the data payload.

**Questions to Ask:**

- How much time did the server take to respond?
- What is the time to last byte or content download time?

**Steps to Take:**

1. Analyze payload sizes using browser development tools or network analysis tools.
2. Collect a HAR file replicating the problem.
3. Monitor server metrics for CPU, memory, and I/O to identify any resource bottlenecks.
4. Profile application performance using APM tools to identify slow-running code, especially code that generates the response (outside the scope of ThousandEyes, where AppDynamics would be best suited).
5. Optimize content delivery by implementing or improving the use of a CDN and ensuring effective caching.
6. Review network performance and correlate with waterfall metrics using ThousandEyes path visualization.
7. Test across different networks to identify if the problem resides on the internet, a specific ISP AS, or the hosting network.
8. Enable compression (gzip or Brotli) on the server.
9. Investigate third-party services if the application relies on APIs or services (ThousandEyes has a test for APIs that can be monitored replicating the application flow and how it interacts with APIs, helping to root cause the problem).

**Possible Causes:**

- Large payloads (not compressed or optimized)
- Server performance problems (slow content generation)
- Limited network bandwidth
- Network congestion
- CDN performance issues
- Server resource limitations

**Blocked**

**Metric meaning:** The time that a browser waits for an already established connection to become available. Web browsers are designed to allow a maximum number of concurrent connections per domain. Blocking time means that the browser is waiting for other requests to complete and represents the time that is spent before a request is sent because other requests are being handled.

**Questions to Ask:**

- Are there any requests in a blocked state?
- How are requests being queued?

**Steps to Take:**

1. Use browser developer tools to see how requests are being queued and to identify any patterns in the blocked time (e.g., a specific file or domain).
2. Use ThousandEyes path visualization and layered views to correlate web metrics to the network.
3. Review rate limiting configurations to ensure they are appropriate for your traffic levels.
4. Optimize page load by reducing the number of initial concurrent requests (combine files, use sprites, defer non-critical requests, implement lazy loading).

**Possible Causes:**

- Server overload
- Too many concurrent requests
- Rate limiting
- DDoS protection
- Browser throttling
- Limited browser resources

## Aggregate Metrics

- DOM Load Time: Transaction time from the beginning of the first object load to the end of the final object load.
- Page Load Time: The time from the initial request to when the page is fully rendered. Redirect time is taken into account when determining total page load time.

## Resources

- [Monitoring Core Web Vitals Metrics](#)
- [Navigating Waterfall Charts for Page Load and Transaction Tests](#)
- [Using Transaction Test Views](#)
- [Using the Page Load View](#)
- [How to Generate a HAR File in your browser](#)
- [What Information Is Transmitted in a Page Load or Transaction Test?](#)
- [HTTP response code definitions](#)

## Sample Questions

Your decision should be based exclusively on the exhibits presented.

### 3.3 Question 1

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




Exhibit 3.3-1



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

Request Headers
Response Headers

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  
www-authenticate: Basic realm="Fake Realm"

Exhibit 3.3-2

- ✓ 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

3.3 Question 2

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




Exhibit 3.3-3

Request HeadersResponse 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 (X11; Linux x86\_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/118.0.5993.70 Safari/537.36  
x-thousandeyes-agent: yes

Request HeadersResponse 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

Exhibit 3.3-4

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

## 3.4 Security Issue Identification

### Analyze

Identify security issues such as DDoS attacks, DNS hijacking, BGP hijacking, and route leaking affecting network performance

## Overview

For a Network Assurance engineer, identifying the fingerprint parameters of security attacks such as DDoS, DNS hijacking, route leaking, and BGP hijacks is crucial. While reading and understanding the nature of these attacks is a good starting point, hands-on experience and the ability to interpret network and application protocol metrics shown in exhibits or outputs are essential to meet this requirement.

## Key Concepts

### DDoS Attack

A Distributed Denial of Service (DDoS) attack's primary objective is to render a service unavailable by denying service to users. One of the most effective ways to achieve this is to generate numerous rogue requests from different locations, making it difficult for real users' requests to be responded to.

Common DDoS attack types include:

- **Volumetric Floods:** Generating traffic to overwhelm bandwidth and resources (e.g., TCP floods, UDP floods).
- **Protocol Attacks:** Exploiting Layer 3 or Layer 4 weaknesses to consume servers' processing capacity (e.g., SYN Flood).
- **Application Attacks:** Generating traffic requests to consume all the server's processing capabilities (e.g., HTTP flood, DNS attacks).
- **Reflection/Amplification Attacks:** Exploiting open DNS resolvers or other vulnerable services to amplify traffic and overwhelm the target (e.g., NTP amplification, DNS amplification).

Monitor DDoS attack patterns or symptoms:

- Look for high latency links and packet loss from several locations (Cloud agents might be best to provide this visibility).

#### Resources:

- [Monitoring DDoS Attacks and Mitigation](#)
- [DDoS Monitoring](#)

### DNS Hijack

A DNS Hijack is a security attack that aims to redirect DNS queries to a rogue DNS server. Attackers may use techniques such as cache poisoning, rogue DNS servers, or man-in-the-middle attacks to achieve this.

**Symptoms:**

- Packet loss
- NS queries resolving to rogue name servers
- Unexpected redirection to malicious websites

**Monitoring strategy:**

- Monitor your name servers for any unauthorized changes.
- Monitor for query errors and increased resolution time of queries.
- Compare DNS records to known legitimate IP addresses or domain names.

**Mitigation during an attack:**

- Flush DNS cache or encourage network operators to do so.
- Implement DNSSEC to ensure the authenticity and integrity of DNS data.

**Resources:**

- [Tips for Instrumenting DNS Alerts](#)

## BGP Hijack

BGP Route Hijacking, also known as prefix hijacking, route hijacking, or IP hijacking, is the illegitimate takeover of groups of IP addresses by corrupting Internet routing tables maintained using the Border Gateway Protocol (BGP). By maliciously manipulating BGP IP prefixes, an attacker (IP hijacker) can reroute traffic to intercept or modify it. This type of attack is successful because BGP ingests the announced IP address prefixes, which are presumed to be owned by the announcing peer.

Attackers may announce more specific prefixes or claim shorter paths to attract traffic. They often target unused prefixes to avoid immediate detection by legitimate owners.

**Symptoms:**

- **BGP Path Changes:** Observe and analyze changes in the AS Path at a specific monitoring point.
- **Availability Drop:** Traffic redirection due to rogue announcements can reduce availability.
- **Packet Loss:** Monitor and document packet loss incidents during the attack.

**Resources:**

- [What is BGP Route Hijacking?](#)
- [Anatomy of a BGP Hijack on Amazon's Route 53 DNS Service](#)
- [Best Practices to Combat Route Leaks and Hijacks](#)

## BGP Route Leak

According to the IETF RFC, a BGP route leak is defined as "the propagation of routing announcement(s) beyond their intended scope." This means an announcement from an Autonomous System (AS) of a learned BGP route to another AS violates the intended policies, potentially causing traffic to be misdirected.

BGP's trust-based nature makes it vulnerable to route leaks. Attackers exploit this to propagate routes, causing issues like traffic blackholing, performance degradation, and increased latency. Using a more specific prefix can make the leaked route more preferable, increasing the impact.

#### Symptoms:

- Packet loss
- BGP path changes
- Increased latency
- Traffic blackholing

#### Resources:

- [Webinar | Detecting Hijacks and Leaks](#)
- [BGP Route Leak](#)
- [RFC 7908](#)

## Case Studies

Review these additional case studies to strengthen your skills in identifying security issues:

-  **YouTube:** Analysis of Amazon Route 53 BGP Hijack  
  
(Video link: <https://www.youtube.com/watch?v=YXm4GJMUIP0>)
- [Twitter Outage Analysis](#)
- [Suspicious Route Against A Root DNS Prefix](#)
- [Akamai Prolexic Routed Outage Analysis](#)
- [Analyzing the Wikipedia DDoS Attack](#)
- [Craigslist DNS Hijack: Charting the Effects](#)

## Sample Questions

### 3.4 Question 1

In real-life applications using ThousandEyes, you can switch between various views. However, for the exam, you will be limited to up to three exhibits. When reviewing answer options, remember to

- Analyze using only the provided exhibits.
- Choose the answer that can be confirmed with the information given.

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




Exhibit 3.4-1: Los Angeles before the Outage




Exhibit 3.4-2: Los Angeles during the Outage

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

#### Hint

- Analyze the details and contrast the provided exhibits to accurately identify potential network issues. Note any changes in Autonomous System (AS) numbers, which are crucial for determining the cause of network problems.
- If there are multiple agents visible in the path visualization view showing packet or forwarding loss, focus on one agent and compare its path against subsequent exhibits to determine the root cause.

### 3.4 Question 2

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

# 4.1 Network Condition Alert Rules



Cognitive Level: Apply

Configure alert rules based on network conditions, such as TCP protocol behavior, congestion, error counters, performance, throughput, state of BGP routing table, internet insights, MPLS, VPN, NetFlow, SNMP, and syslog

## Overview

Alert rules can be configured in ThousandEyes to monitor various network conditions and metrics from Cloud and Enterprise Agents, Endpoint Agents, BGP, devices, and Internet Insights. Alerts can be set up to notify when thresholds are exceeded for metrics like packet loss, latency, jitter, page load time, throughput, BGP reachability, device interface status, and more.

## Key Concepts

### ThousandEyes Alert Rules

- All alert rules have four sections: Description, Settings, Notifications, and Alert Conditions.
- Settings configure the "big picture" of what test data will trigger the alert.
- Notifications determine who/what systems get notified when an alert triggers.
- Alert Conditions specify when the alert should trigger based on global and location criteria.
- Default alert rules are automatically added to new tests but can be disabled.
- Custom alert rules are recommended to match specific requirements and reduce noise.

### Alert Structure

All alert rules have four sections:

1. **Description:** The alert type (data source/test type) and the alert name.
2. **Settings:** Selection of tests that will trigger this alert.
3. **Notifications:** Recipients and systems to be notified when an alert is triggered.
4. **Alert Conditions:** Criteria for when this alert should trigger.



## Add New Alert Rule

Alert Type
Web
HTTP Server

Rule Name

### Compatible Test Types

Web
HTTP Server
Page Load
Transaction

### Settings

Tests
0 of 29 test(s) selected

Agents
All agents

Severity
Info
Minor
Major
Critical

### Notifications

EMAIL
Send emails to
Select Email Address

☒ Send an email when the alert clears

Add message

### ALERT CONDITIONS

New! Quantile dynamic baselines: minimize false positives by relying less on outliers to determine the baseline. Try it today.
Learn More

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

Error
is present

Cancel
Create New Alert Rule

Figure 4.1-1: Alert Structure

## SETTINGS

Alert settings options change depending on the alert category selected. For example, Cloud and Enterprise Agents alerts have different settings than Endpoint Agents to reflect the supported test types and collected conditions.

There are 5 alert categories available:

| Cloud and Enterprise Agents                                                                     | Endpoint Agents                                                                                         | BGP Routing                                                                           | Devices                                                     | Internet Insights                                                                                                   |
|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>- Agents</li> <li>- Severity</li> <li>- Tests</li> </ul> | <ul style="list-style-type: none"> <li>- Agents</li> <li>- Severity</li> <li>- Visited Sites</li> </ul> | <ul style="list-style-type: none"> <li>- Monitors</li> <li>- Prefix Length</li> </ul> | <ul style="list-style-type: none"> <li>- Devices</li> </ul> | <ul style="list-style-type: none"> <li>- Affected Tests</li> <li>- Catalog Providers</li> <li>- Severity</li> </ul> |

## Alert Conditions

Alert conditions have two sections:

- Global condition
- Location condition

When the global condition is met, any agent that meets the location condition in a test round will be included in the alert as "active".

### GLOBAL CONDITIONS

This section describes how to apply the global section of the Alert Conditions.

The figure shows a configuration interface for alert rule conditions. It consists of a main sequence of dropdown menus and input fields: "All" (dropdown), "conditions are met by", "any of" (dropdown), "1" (input), "agent" (dropdown), "1" (input), "of", "1" (input), and "time in a row:". Below the "All" dropdown is a dropdown labeled "Any". Below the "any of" dropdown is a dropdown labeled "the same". Below the "agent" dropdown is a dropdown labeled "% of agents".

Figure 4.1-2: Alert Rule Conditions

The Global section has the following format:

```
<All>/<Any> conditions are met by <any of>/<the same> ### <monitor>/<% of monitors>/<agent>/<% of agents> # of  
# times in a row
```

The `<All>/<Any> conditions` option sets how many individual location alert conditions are required to continue evaluating the Global section. For example, if "All" is selected, the alert will only trigger when all conditions are met. If "Any" is selected, the alert will trigger if any condition is met.

For alert rules that need more than one test round to trigger, the `<any of>/<the same>` section sets if the agents or monitors being evaluated must be the same each impacted test round. Setting `<the same>` allows you to catch specific use cases.

The `<monitor>/<% of monitors>/<agent>/<% of agents>` section allows you to choose a count or percentage of agents or monitors needed for the test to trigger. Using a percentage is best when you have multiple tests with varying numbers of agents or monitors.

When using a percentage, the percentage of agents or monitors is truncated, not rounded up. So if you have 14.7% of agents meeting the alert conditions and have set the "% of agents" to 15%, the alert will not trigger.

The `# of # times in a row` sets how many test rounds the alert rule will look at. When the two numbers are the same (1 of 1, 4 of 4, etc.), all specified test rounds must meet the location conditions for the alert to trigger. Think of this as a sliding window of test rounds that must contain the first number of rounds meeting the location conditions.

### LOCATION CONDITIONS

Location alert conditions are where you set the specific metrics on which an alert becomes active. You can set any number of metrics for an alert, though bear in mind that the more metrics you set, the less likely it is an alert will activate.

Location alert conditions are configured by choosing at least one metric (the test characteristic against which you're measuring change) and one operator (the type of measure). Depending on the metric, other configurable options include threshold values and units.

A location alert is included within a global alert when a single alert trigger meets the location alert conditions for at least one round, regardless of the thresholds set for the global alert.

It's important to note that location alerts trigger and clear independently from the global alert. If you see multiple location alerts triggered under a global alert, you cannot assume that all the listed location alerts met the initial alert criteria from a per-round basis.

For more on global and location alert conditions, see the [ThousandEyes documentation](#).

### Example Alert Conditions

Understanding the alerting capabilities of a network assurance platform like ThousandEyes is highly important. This section summarizes common events to alert on for each alert category. The first column lists the event type, and the second contains the alert condition configuration.

#### NETWORK TESTS ALERT CONDITIONS

Use this configuration to monitor network metrics such as packet loss, latency, jitter, bandwidth, and throughput.

| Event                        | Condition                                    |
|------------------------------|----------------------------------------------|
| High Latency in Asia-Pacific | Latency $\geq$ 180 ms                        |
| High Network Packet Loss     | Packet Loss $\geq$ <code>threshold_%</code>  |
| High Network Jitter          | Jitter $\geq$ <code>threshold_ms</code>      |
| QoS Marking Change           | Any hop not in DSCP <code>dscp_value</code>  |
| Network Loop Detected        | Path length $>$ <code>max_path_length</code> |

#### BGP ROUTING ALERT CONDITIONS

Use this configuration to monitor AS path, route reachability, and route updates.

| Event                     | Condition                                                                                      |
|---------------------------|------------------------------------------------------------------------------------------------|
| Route Flaps               | Path changes > 1 & reachability < 100%                                                         |
| Prefix Hijack             | BGP ASN not in <code>expected_asn_list</code>                                                  |
| DDoS Mitigation Activated | BGP ASN in <code>mitigation_asn_list</code> or prefix not in <code>expected_prefix_list</code> |
| Upstream Provider Change  | BGP HOP# from origin not in <code>expected_hop_list</code>                                     |

## DNS & WEB TESTS ALERT CONDITIONS

Use this configuration to monitor web server response time, wait time, load time, transaction duration, and/or DNS response.

| Event                       | Condition                                                                    |
|-----------------------------|------------------------------------------------------------------------------|
| Slow DNS Resolution         | Response time > 20 ms                                                        |
| DNS Mapping Change/Spoofing | Mapping not in <code>expected_ip_address</code>                              |
| Slow Transaction            | Duration > <code>threshold_ms</code>                                         |
| Embed URL Not Working       | Any component domain in <code>domain_list</code> & component load incomplete |
| Slow Throughput             | Throughput < <code>threshold_kbps</code> Kbps                                |

## INTERNET INSIGHTS ALERT CONDITIONS

Use this configuration to monitor affected applications, outage error types, locations for application outages, affected domains, ASNs, locations, and interfaces for network outages.

| Event                          | Condition                                                                |
|--------------------------------|--------------------------------------------------------------------------|
| Google Workspace App Outage    | Affected App in Google Workspace                                         |
| Application Outage due to DNS  | Affected app in <code>app_list</code> & Outage Error Type in DNS         |
| CDN Network Outage in US       | Locations in United States & affected domain in <code>domain_list</code> |
| Network Outage Services Impact | Affected tests count & location in <code>location_list</code>            |

## Resources

- [ThousandEyes Alerts](#)
- [ThousandEyes Event Detection](#)
- [ThousandEyes Metrics: What Do Your Results Mean?](#)

## Sample Questions

### 4.1 Question 1

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

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

### 4.1 Question 2

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

AS origin AS 1234

BGP

BGP Origin: (ASN not in [1234])

Settings

Notifications

GENERAL

Rule NameAS origin AS 1234

Tests1 of 29 test(s) selected

Prefix LengthIPv4 /10 - /32 And IPv6 /32 - /128

MonitorsAll monitors

☒ Alert on covered prefix data

SeverityInfoMinorMajorCritical

ALERT CONDITIONS

All conditions are met byany of1monitor1of1time in a row:

BGP Originmeets all of:

ASNnot in1234

Collapse All

Exhibit 4.1-1

## 4.1 Question 3

Refer to the exhibit. The alert rule is set up as shown, but didn't trigger. Why?

- ☒ A) Alert conditions weren't met and won't trigger with current setup
- ☒ B) Alert needs two consecutive agent failures to trigger
- ☒ C) Response code is set up incorrectly
- ☒ D) All of the above

Settings

Notifications

GENERAL

Rule Name

Response code 4XX

Tests

1 of 21 test(s) selected

Agents

All agents

Severity

Info

Minor

Major

Critical

ALERT CONDITIONS

All conditions are met by

any of

2

agents

1

of

1

time in a row:

Response Code

is

any error (≥ 400 or no response)

Exhibit 4.1-2




Exhibit 4.1-3

## 4.1 Question 4

Refer to the exhibit. 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 ≥ Static 500ms
- ✓ D) Response Code is server error(5XX)

### Add New Alert Rule

Alert Type

Web

HTTP Server

Rule Name

ENNA-HTTP-Alert

Settings

Notifications

Tests

0 of 29 test(s) selected

Agents

All agents

Severity

Info

Minor

Major

Critical

#### ALERT CONDITIONS

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

?

?

?

—

+

Exhibit 4.1-4



# 4.2 End-User Experience Alert Rules

 Cognitive Level: Apply

Configure alert rules that affect the end-user experience, such as CPU utilization, connectivity types (wired to wireless, Wi-Fi), browser behavior, and VPN

## Overview

End-user experience metrics are mainly found in the Endpoint Agents alert rule set. Some relevant metrics are also available in Cloud and Enterprise Agent Web tests (Page Load and Transaction). These metrics, while not specific to end users, provide valuable insights into user experience. Transaction tests can execute an entire user flow, such as a login process, and repeat the flow each time the test runs. These results closely mirror the end-user experience.

Alert rule categories, structure, and condition concepts are covered in section 4.1.

Key areas monitored with ThousandEyes alert rules include:

- Local system performance
- Network tests and path trace
- Local network connectivity
- Application (HTTP scheduled tests) monitoring
- Browser experience

Meraki and Catalyst Center also offer options for configuring alerts related to end-user experience, such as monitoring VPN tunnels, wireless connectivity, and endpoint security events.

## Example Alert Conditions

It is crucial for a network assurance engineer to understand the alerting capabilities of a specific platform, such as ThousandEyes Endpoint Agents. This section summarizes commonly alerted events and their configurations for each alert category. The tables below list the event types and their corresponding alert conditions.

### Local System Performance Alert Conditions

Use this configuration to monitor the endpoint's CPU utilization or memory load. These can be found in the Real User Tests > Endpoint and Scheduled Tests > Endpoint End-to-End (server) alert types. It's important to alert on system-related metrics that might impact the overall user experience.

| Event                | Condition           |
|----------------------|---------------------|
| Endpoint performance | CPU utilization ≥ % |
| Endpoint performance | Memory load ≥ %     |

## NETWORK TESTS AND PATH TRACE ALERT CONDITIONS

Use this configuration to monitor endpoint network performance metrics and the path to the destination. Alert conditions for metrics such as packet loss, latency, and jitter are available in all Endpoint Alert types. Path trace conditions are specific to Scheduled Tests > Endpoint Path Trace alert types.

| Event                     | Condition                 |
|---------------------------|---------------------------|
| End-to-End packet loss    | $\geq$ _%                 |
| End-to-End packet latency | $\geq$ _ ms               |
| Path monitoring           | Any Hop ASN in/not in ASN |
| Path monitoring           | Path length > #           |

## LOCAL NETWORK CONNECTIVITY ALERT CONDITIONS

Use this configuration to monitor local network metrics. This alert condition is found within Real User Tests. Metrics include Connection Failures Count, Content Time, Error Count, Experience Score, Jitter, Latency, Packet Loss, Page Load Time, Response Time, and Signal Quality. Note that not all wireless metrics in the test views are available for alert rules.

| Event           | Condition               |
|-----------------|-------------------------|
| Wifi monitoring | Signal Quality $\leq$ % |

### Application Monitoring Alert Conditions

Use this configuration to monitor web pages or SaaS platforms relevant to end users. Application and network metrics can be correlated for root-cause analysis.

| Event                  | Condition                 |
|------------------------|---------------------------|
| Application Monitoring | Response Time $\geq$ ms   |
| Application Monitoring | Page Load Time $\geq$ %   |
| Application Monitoring | Error count $\geq$ #      |
| Network Monitoring     | Packet Loss Time $\geq$ % |
| Network Monitoring     | Latency $\geq$ ms         |

## Browser Experience Alert Conditions

Use this configuration to monitor metrics related to user browsing experience. These conditions are available for the Real User Tests alert type. For a comprehensive monitoring strategy, scheduled tests alert type can also be configured for other browsing experience metrics. Metrics from scheduled tests are derived from active monitoring, while metrics from real user tests come directly from the end user's browser interaction, representing passive monitoring.

| Event                     | Condition                                                | Threshold    |
|---------------------------|----------------------------------------------------------|--------------|
| Connection Failures Count | Count Greater than (>) or Less than (<)                  | Number       |
| Content Time              | Number in milliseconds Greater than (>) or Less than (<) | Milliseconds |
| Error Count               | Greater than (>) or Less than (<)                        | Number       |
| Experience Score          | Greater than (>) or Less than (<)                        | Percentage   |
| Jitter                    | Greater than (>) or Less than (<)                        | Milliseconds |
| Latency                   | Greater than (>) or Less than (<)                        | Milliseconds |
| Packet Loss               | Percentage Greater than (>) or Less than (<)             | Percentage   |
| Page Load Time            | Greater than (>) or Less than (<)                        | Milliseconds |
| Response Time             | Greater than (>) or Less than (<)                        | Milliseconds |
| Signal Quality            | Greater than (>) or Less than (<)                        | Percentage   |

### Preparing for the Exam

- Configure alerts to notify on end-user experience
- Understand the different metrics collected by the Endpoint Agent and how these metrics are obtained
- Understand the difference between "All" and "Any" global configuration settings

## Resources

- [ThousandEyes Alerts](#)
- [ThousandEyes Event Detection](#)
- [ThousandEyes Metrics: What Do Your Results Mean?](#)

# Sample Questions

## 4.2 Question 1

Refer to the exhibit. A network engineer is tasked with configuring an alert that will trigger if the Endpoint Agent path ASN changes on a specific hop. What is the alert type and condition needed to meet the requirement?

- ✓ A) Scheduled tests, Hop#
- ✓ B) Real User Test, Hop#
- ✓ C) Scheduled tests, Any Hop
- ✓ D) Real User Test, Path Length

Add New Alert Rule

Alert Type

Endpoint Path Trace

Alert Rule Name

Settings

Notifications

Tests

0 of 29 test(s) selected

Agents

All agents

Severity

Info

Minor

Major

Critical

ALERT CONDITIONS

All conditions are met by at least 1 agent and 30 % of agents 1 of 1 time in a row:

1 from the source meets all of:

Delay

≥

100

ms

Collapse All

Exhibit A

## 4.2 Question 2

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  $\geq$  %

## 4.3 Dashboard Configuration

### Cognitive Level: Apply

Select deliverables or metrics such as dashboard and alerts for IT operations, production support, app/dev teams, and executives

## Overview

This section guides you through configuring effective dashboards. The key is selecting metrics that not only align with specific business requirements but also provide actionable insights.

Start by identifying your target audience and setting clear expectations for the dashboard's purpose. To ensure accurate interpretation of the data, provide clear explanations of the source for each metric.

A well-designed dashboard acts as a powerful tool for decision-making. By tailoring metrics and visualizations to specific business objectives, you can empower users to make better decisions and drive positive outcomes.

### Preparing for the Exam

- Review the [Getting Started with Dashboards](#) guide. This resource provides step-by-step instructions to build a dashboard in 5-10 minutes and covers fundamental key concepts about dashboards.
- Ensure you understand the different dashboard display settings, including relative vs. fixed-time intervals, global time override, and local widget time settings.
- Create dashboards and familiarize yourself with the various data sources, categories, metrics, and measures available for widget configuration. Remember that each test type (category) will offer a different list of metrics depending on the data source.
- Practice working with measures to understand how they affect the presentation of your data in widgets.

## Key Concepts

### Metrics

ThousandEyes tests provide a range of metrics, which can be viewed in the Views menu for Cloud and Enterprise Agents. Users can select specific metrics using a drop-down filter list within the interface.

Different test types offer distinct sets of metrics. For example:


Page Load Test metrics include:

- Page Load Time
- Errors
- Timeouts
- Completion

HTTP Test metrics include:

- Availability
- Response Time
- Throughput

Agents running Page Load tests collect metrics specific to page load and HTTP, as well as network metrics such as loss, latency, and jitter. All these metrics are accessible in views and can be utilized in dashboard construction.



ThousandEyes Views: Page Load Test Results

## Dashboards

ThousandEyes dashboards provide customized live views of Enterprise & Cloud Agent tests, Endpoint Agent tests, device layer data, and Internet Insights. These dashboards allow users to visualize and monitor key performance indicators in real-time.

To illustrate, consider an "API Health Overview Dashboard." This dashboard could utilize Number widgets (from the data summary type) to display the average (mean) values for critical API metrics, such as:

- API Transaction Time
- API Call Time
- DNS Time

By configuring these widgets to use the "Cloud and Enterprise Agents" data source, the "Web-API" category, and the relevant metrics, users can create a concise view of API performance.

## WIDGETS

Widgets are customizable visual elements used to display data on ThousandEyes dashboards. They come in various types, including:

- Live status (agent status, tests, alert lists)
- Breakdown (stacked bar, grouped bar, and pie charts)
- Data summary (tables, multi-metric tables, numbers)
- Time series (line charts, stacked area charts, box and whisker plots)
- Maps


## EMBEDDED WIDGETS

**Embedded widgets** allow you to display ThousandEyes data visualizations on external web pages, making the information accessible to a broader audience without requiring direct platform access.

## Troubleshooting with Dashboard Drill Down

Dashboards serve as valuable tools for troubleshooting network issues. The API Health Overview Dashboard, for instance, provides a comprehensive view of API performance metrics, allowing engineers to monitor and troubleshoot efficiently. Let's walk through a troubleshooting scenario using this dashboard.






*API Health Overview Dashboard*


The API Health (Mean) Widgets are configured using Number tiles from the data summary widget type. They use Cloud and Enterprise Agents as the data source, focusing on Web-API category metrics such as API transaction time, call time, and DNS time, with mean as the measure.

These widgets are ideal for monitoring API performance based on average values for transaction, call, and DNS time for specific tests.




Widget Configuration


In our scenario, let's say the engineer finds that the 1.57s API transaction time is high compared to how their API usually performs. To investigate further, the engineer clicks on the mean transaction time widget, which brings up a window with links to the tests associated with the widget.




Upon clicking on the widget, the engineer is taken to the exact point in time for the metric shown in the widget. Here, they can contrast API transaction time with other metrics such as API completion, packet loss, latency, and jitter to correlate this particular event.



From this page, the engineer can also switch from the API layer view to the Agent to Server network layer view.



In this view, the engineer can look at the path visualization, which can help root-cause network performance issues affecting the API Transaction time. For example, in the image above, notice the yellow color for the agent. This indicates that for this test round, the agent had an 8% packet loss. You can see a peak for both end-to-end packet loss and API transaction time.



By expanding the nodes in the path visualization, the engineer can find links with latency, further aiding in the root-cause analysis of network performance issues affecting the API Transaction time.

This drill-down capability of ThousandEyes dashboards is a powerful troubleshooting tool, allowing engineers to move from high-level overviews to detailed, specific data points quickly and efficiently. For a deeper dive into these features, refer to the [Troubleshooting with Dashboard Drill Down](#) documentation.

## Resources

- [ThousandEyes Dashboards](#)
- [ThousandEyes Dashboard Widgets](#)
- [ThousandEyes Metrics: What Do Your Results Mean?](#)
- [Data Collected by Endpoint Agent](#)
- [Proxy Metrics in HTTP Server Tests](#)
- [Troubleshooting with Dashboard Drill Down](#)

## Sample Questions

The following sample questions require you to analyze data presented in two ThousandEyes dashboards used to monitor its application service at <https://thousandeyes.com>:

- [Executive Dashboard](#): provides a high-level overview of application performance
- [IT Operations Dashboard](#): offers granular insights for troubleshooting and performance optimization.

Refer to the data in these dashboards to answer the questions below.

### 4.3 Question 1

Which type of test are we using for these dashboards?

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

### 4.3 Question 2

Which type of widgets were used in the executive dashboard? (Select all that apply)

- ☒ A) Agent status
- ☒ B) Map
- ☒ C) Line
- ☒ D) Number
- ☒ E) Color Grid

### 4.3 Question 3

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

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

### 4.3 Question 4

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

### 4.3 Question 5

In the executive dashboard, what is the page completion time for the Mexico City agent?

- ☒ A) 100%
- ☒ B) 83.4%
- ☒ C) 15.2%
- ☒ D) 99.67%

### 4.3 Question 6

In the executive dashboard, what is the total error count for ThousandEyes web page in the last 15 days?

- ☒ A) 520
- ☒ B) 1.58
- ☒ C) 4610
- ☒ D) 4805

### 4.3 Question 7

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

### 4.3 Question 8

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



## 4.4 Alert Configuration and Functionality



Cognitive Level: Analyze

Validate alert configuration and functionality

### Overview

Validating alert rules is a critical step to ensure they are triggering as intended based on the configured settings and conditions. This involves a methodical approach to test the individual alert conditions as well as the overall alert rule behavior.

Generally speaking, the key aspects to focus on when validating alerts include:

- Verifying the tests, agents, monitors, and devices the alert rule applies to
- Confirming that alert conditions trigger as expected based on the configured metric thresholds
- Ensuring alert notifications are sent to the right recipients/systems with the correct information
- Tuning thresholds to reduce alert noise while still catching important events
- Testing alerts under different scenarios to validate the configuration is optimized

The most recurrent question while validating an alert is related to its sensitivity:

- Why did my alert did not trigger?

On the other hand, there might be scenarios where the alert has become too noisy:

- Why is my alert triggering so often?

We seek to address these two questions as they are relevant to the exam and real life.

### Key Concepts

#### Sensitivity

As discussed in 4.1, the global setting is the first logic evaluated in order to trigger the alert. This configuration defines the overall sensitivity of the alert.

#### How many rounds to trigger on?

Lets start with the most basic setting of the "Global" section to have the alert trigger: if any 1 round of 1 meets the local conditions. This means that anytime the conditions are met, the alert triggers. On the surface this looks great. But what if the test is running into "noise"? In the below example, 3 rounds in 4 the conditions are met and an alert is triggered.

##### Round 1:

- Alert conditions are met
- Alert is triggered

**Round 2:**

- Alert conditions are met
- Alert status is unchanged (triggered)

**Round 3:**

- Alert conditions are not met
- Alert is cleared

**Round 4:**

- Alert conditions are met
- Alert is triggered again.

In this situation having the alert trigger on any single round would result in two extra notifications. Round 3 is clear, and round 4 retriggers.

While in some cases these multiple rounds of alerts are wanted, usually the desired result here would be to have this alert stay triggered, thus saving two notifications which need to be evaluated and acted upon. These additional notifications for the same trigger are the "noise" mentioned earlier.

So how do we desensitize the alert rule a small amount to reduce that noise?

The best place to start would be to increase the threshold of testing: instead of 1 of 1 test rounds, what happens if we go to any 1 of 2 test rounds?

**Round 1:**

- Alert conditions are met
- Alert is triggered

**Round 2:**

- Alert conditions are met
- Alert status is unchanged (triggered)

**Round 3:**

- Alert conditions are not met
- Alert status is unchanged (triggered) as the previous test round keeps the alert triggered.

**Round 4:**

- Alert conditions are met
- Alert status is unchanged (triggered).

In this example you would need two test rounds in a row that did not meet the alert conditions for an alert to clear.

Using this logic you can tailor your alert rule to ensure that you are receiving notifications precisely how often you want them, avoiding alert-fatigue.

## Any or The Same agents?

In many situations all you care about is that some agents have met the alert conditions for a test round. But if you work in an environment where that creates a lot of noise, an alternative to consider would be to stipulate that you want the **same** agent to be the one used for the alert conditions to be met.

### ANY (2 OF 2 TEST ROUNDS)

#### Round 1:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are not met
- Alert is not triggered

#### Round 2:

- New York, NY:
  - Alert conditions are met
- Los Angeles, CA:
  - Alert conditions are not met
- Alert is not triggered (only 1 of 2 test rounds)

#### Round 3:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are met
- Alert triggered

#### Round 4:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are met
- Alert status is unchanged (triggered) as the previous test round keeps the alert triggered.

#### Round 5:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are not met
- Alert cleared (no agents met the alert conditions this test round breaking 2 of 2)

## THE SAME (2 OF 2 TEST ROUNDS)

### Round 1:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are not met
- Alert is not triggered

### Round 2:

- New York, NY:
  - Alert conditions are met
- Los Angeles, CA:
  - Alert conditions are not met
- Alert is not triggered (only 1 of 2 test rounds)

### Round 3:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are met
- Alert is not triggered (different agents met the conditions)

### Round 4:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are met
- Alert triggered (Los Angeles, CA met the conditions twice in a row)

### Round 5:

- New York, NY:
  - Alert conditions are not met
- Los Angeles, CA:
  - Alert conditions are not met
- Alert cleared (no agents met the alert conditions this test round breaking 2 of 2)

## Percent of agents or number of agents?

While specifying the exact number of agents needed to trigger an alert is often very convenient, it is not exactly scalable. If the alert rule is dedicated to a specific test, this is fine. But as the number of tests grow, typically the level at which you want the alert rule to trigger also changes. Enter 'Percent of agents' as a solution.

A word of caution about using percent of agents: it will always use whole numbers. In the case of a partial agent, the number of agents is adjusted to the closest whole number.

As an example: if a rule is set with 10% of agents as a threshold, and you have 18 agents, the rule will only trigger if 2 agents meet the criteria, as there is no such thing as 1.8 agents. This one fact has caused a surprising number of support cases.

**Hint:** Many find it easier to make an alert too sensitive and adjust the sensitivity down. While you will receive more "noise" at first, you will not miss any alerts you really want while fine-tuning the alert sensitivity to your preference.

## Alert Rule Validation

### Validate the individual alert conditions

The most specific you can be in an Alert Rule are the individual local conditions. You need to know when and if any given condition will trigger for any test. Without this granular understanding, building complex alert rules can be a very frustrating process.

Learn what each of the Alert Metrics represents and what information it uses.

### Validate the set of conditions

In the Global section you can select either "Any" or "All of" the alert conditions for triggering the alert.

#### ANY

Just as the name implies "Any" means that even just one of the specified local alert conditions can trigger the local part of the alert rule.

This is great if you want a single alert rule to catch multiple issues. The best practice when utilizing this method is to group like metrics together. So one alert rule might cover network metrics, while another monitors page load and transaction test status.

Think of the "Any" as a logical "OR" for the alert conditions.

#### ALL OF

Your other option is "All of", allowing you to build some fairly complex alert condition sets.

One of the primary uses of the "All of" option is to lower the "noise" of your alerts. The more complex the condition set is the less likely your alert is to trigger. This effectively desensitizes the alert rule.

Think of the "All of" as a logical "AND" for the alert conditions.

**Hint:** A test alert rule can be used to validate each of the alert conditions separately. A Global setting any one agent one of one time in a row, the most sensitive global condition is great for this.



You will want to consider disabling notifications for the test alert rule to prevent spamming the rest of your team.

## Resources

- [ThousandEyes Alerts](#)
- [ThousandEyes Event Detection](#)

## Sample Questions

### 4.4 Question 1

An alert rule for a Web - HTTP Server test is not triggering when the HTTP response code is 500 Internal Server Error. The alert conditions are configured with "Response Code" set to "any error ( $\geq 400$  or no response)". What could be causing the alert to not fire?

- ☒ A) The alert rule is disabled
- ☒ B) The test is not enabled on any Enterprise Agents
- ☒ C) The alert rule's "Settings" section does not have the correct test selected
- ☒ D) The HTTP server is returning a 200 OK response code

### 4.4 Question 2

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?

- ☒ 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

# 4.5 Network Capacity Planning

## Cognitive Level: Analyze

Recommend optimization for network capacity planning, such as topology and configuration changes, and QoS based on data interpretation

## Overview

Network Assurance engineers must be proactive in monitoring network capacity and performance to ensure a consistently high-quality user experience. As organizations evolve, the ability to gather and interpret relevant data becomes increasingly important for making informed optimization decisions. This section will strengthen your understanding of key concepts and tools related to capacity planning and optimization, preparing you to effectively analyze data and recommend appropriate actions.

## Key Concepts

### Capacity Planning

Capacity planning is a crucial aspect of network assurance. It involves forecasting bandwidth requirements based on user needs and growth trends. This proactive approach ensures that users consistently receive the best possible digital experience over the network.

ThousandEyes WAN Insights can assist with capacity planning by integrating with Cisco Catalyst SD-WAN Manager. This integration obtains the maximum bandwidth for each circuit, providing valuable data for optimization decisions. For more information, refer to the product documentation on [How Capacity Planning works](#) with WAN Insights.

Capacity planning is essential for all traffic types but is particularly critical for applications sensitive to bandwidth limitations and network issues, such as video conferencing and Voice over IP (VoIP) calls.

### Quality of Service (QoS)

Quality of Service is a key feature that can be leveraged alongside a robust capacity planning strategy. QoS helps prioritize network traffic to ensure optimal performance for critical applications.

To understand QoS, consider this analogy:

Imagine your network as a highway full of vehicles, where each vehicle represents a chunk of data. QoS acts like a traffic management system, ensuring that high-priority vehicles (critical data) don't get stuck in congestion.

### IMPORTANCE OF QOS

Different network activities have varying requirements for optimal performance:

- VoIP calls require low latency to avoid awkward pauses.
- Video streaming needs a steady data stream to prevent buffering.

QoS ensures these activities run smoothly, even during peak network usage.

## HOW QOS WORKS

1. Data packets are tagged with priority levels.
2. Network devices read these tags and sort data into different queues.
3. Higher priority data is processed first.
4. Some QoS implementations reserve bandwidth for critical applications, similar to a dedicated lane on a highway.

## Data Interpretation for Optimization

To make informed optimization recommendations, you'll need to analyze data from various sources, including:

- Telemetry
- SNMP (Simple Network Management Protocol)
- CLI (Command Line Interface) outputs
- Syslog messages
- NetFlow

These data sources provide insights into network performance, utilization, and potential bottlenecks. By interpreting this data, you can identify areas for improvement and suggest appropriate optimizations.

## Resources

- [QoS Design Principles and Best Practices](#)
- [Enterprise QoS Solution Reference Network Design Guide](#)
- [ThousandEyes Capacity Planning](#)
- [ThousandEyes WAN Insights](#)

## Sample Questions

### 4.5 Question 1

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



## 4.5 Question 2

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

## 4.5 Question 3


CLI outputs show that a router's egress queue for voice traffic is consistently full, leading to increased latency. Based on this data, which optimization would you recommend?

- ☒ A) Increase the queue size for voice traffic
- ☒ B) Implement traffic shaping on non-voice traffic
- ☒ C) Disable QoS on the router to allow all traffic equal priority
- ☒ D) Replace the router with a higher-capacity model

## 4.5 Question 4

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



4.5-1: Capacity Planning Results

# Domain 1 Answer Key

## Agent Types

### 1.1 Question 1

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

- ☐ A) Synthetic Agent
- ☒ B) Enterprise Agent
- ☐ C) Cloud Agent
- ☐ D) Endpoint Agent

#### Explanation

The correct answer is **B) Enterprise Agent** because it:

- Meets all requirements, especially collecting the path from the internal network to the cloud provider
- Can be installed on access switches, as implicitly asked in the question

Incorrect answers:

- A) This answer is a distractor, as all ThousandEyes agents are synthetic agents
- C) Would not provide insights from the internal network
- D) Good for internal network insights but cannot be installed on access switches

### 1.1 Question 2

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

### Explanation

The correct answer is **D) Cloud Agent** because:

- They are pre-deployed by ThousandEyes and immediately available
- They provide diverse geographical perspectives without requiring infrastructure ownership
- There are 200+ observation points worldwide

Incorrect answers:

- A) This answer is a distractor, as all ThousandEyes agents are synthetic agents
- B) This answer choice could be a distractor if the reader misinterprets the data center as the observation point rather than the target
- C) Would need to be installed on machines where customers are connecting, which is a less practical solution for this scenario

## 1.1 Question 3

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

- ☐ A) Synthetic Agent
- ☐ B) Enterprise Agent
- ☐ C) Cloud Agent
- ☒ **D) Endpoint Agent**

### Explanation

The correct answer is **D) Endpoint Agent** because:

- It is directly installed on the user's workstation which provides individual insights on any problem the web application and SaaS platforms may experience

Incorrect answers:

- A) This answer is a distractor, as all ThousandEyes agents are synthetic agents
- B) Cannot be deployed to the user's machine
- C) You can't select the location of the agent, as they are pre-deployed, making it difficult to use for a specific user location to measure end-user experience.

# Agent Location

## 1.2 Question 1

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

### Explanation

The correct answer is **A) Install an agent on the overlay network** because:

- It proactively measures and monitors SD-WAN performance and routing policy validation which is the overlay

Incorrect answers:

- B) Installing an agent on the DMZ does not guarantee being able to run tests on the Overlay network
- C) The LAN is the wrong location as this is an internal network, not part of the SD-WAN tunnel
- D) This is if you want to gain hop-by-hop visibility into the network underlay, not measure SD-WAN

## 1.2 Question 2

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



## Explanation



This question tests the understanding of various agent types and their deployment methods in the context of monitoring VPN-connected users.

Correct Answers:

- C) Deploying Enterprise Agents on the VPN concentrator provides visibility into the VPN tunnel itself. This helps identify if issues stem from the VPN infrastructure or beyond.
- D) Deploying Endpoint Agents on employee machines offers a true representation of the end-user experience, encompassing everything from the user's device, home network, and the VPN connection to the application.

Incorrect Answers:

- A) While helpful for general internet performance insights, Cloud Agents won't have visibility into the VPN tunnel.
- B) AppDynamics focuses on application performance monitoring (APM), providing server-side insights. While valuable, it doesn't replace network-level monitoring for VPN users.
- E) Data center-based tests lack visibility into the remote employee's network path, including the VPN connection.

# Active and Passive Monitoring

## 1.3 Question 1

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

### Explanation

The correct answer is **C) Sending a continuous ping from one office to another to measure latency** because it generates and injects synthetic traffic into the network to directly measure performance metrics such as latency, allowing for proactive testing of network performance.

Incorrect Answers:

- A) Analyzing SNMP data is passive monitoring
- B) Capturing packets is passive monitoring
- D) Collecting NetFlow records is passive monitoring

## 1.3 Question 2

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**
- ☐ 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



## Explanation



The correct answer is **B) Passive monitoring can provide real-time data on network performance without adding traffic to the network** because it highlights a key advantage of passive monitoring, which is the ability to monitor the network using actual traffic without introducing additional load or overhead

Incorrect answers:

- A) This is an active monitoring technique, as this will generate synthetic traffic and can potentially add load to the network
- C) Passive monitoring allows for the generation of test traffic to simulate user behavior: This is an active monitoring technique, as this will generate synthetic traffic and can potentially add load to the network
- D) This is an active monitoring technique, as this will generate synthetic traffic and can potentially add load to the network



# WAN Insights

## 1.4 Question 1

What is ThousandEyes WAN Insights, and how does it complement Cisco's SD-WAN network infrastructure?  
Select all that apply.

- ☒ **A) A predictive network path tool that uses historical data to recommend optimal paths within Cisco SD-WAN**
- ☐ B) A hardware device for real-time network traffic monitoring and analytics
- ☒ **C) Provides visibility into network performance, including the public Internet, by working with ThousandEyes**
- ☐ D) A set of network management tools that leverage SNMP and flow protocols into a single dashboard
- ☐ E) An antivirus solution that protects networks from cyber threats

### Explanation

The correct answer is **A) A predictive network path tool that uses historical data to recommend optimal paths within Cisco SD-WAN** and **C) Provides visibility into network performance, including the public Internet, by working with ThousandEyes** because they are both ThousandEyes WAN Insights

Incorrect answers:

- B) WAN Insights is not a hardware offering
- D) WAN Insights is not a network management tool
- E) WAN Insights is not an antivirus solution

## 1.4 Question 2

Which of the following data sources does ThousandEyes WAN Insights use to provide network performance visibility? Select all that apply.

- ☒ **A) Historical network data**
- ☐ B) A WAN-dedicated ISP that offers high-speed connectivity
- ☐ C) SNMP data
- ☒ **D) Flow protocol data**
- ☐ E) Antivirus data



## Explanation



The correct answer is **A) Historical network data** and **D) Flow protocol data**

Incorrect answers:

- B) WAN Insights is an integration between ThousandEyes and Cisco Catalyst SD-WAN Manager
- C) Passive monitoring data such as flow records from the SD-WAN fabric but SNMP is not leveraged
- E) WAN Insights is not a solution focused on security

# Cisco Integrations

## 1.5 Question 1

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 Agent**
- ☐ D) Endpoint Agent

### Explanation

The correct answer is **C) Enterprise Agent** because Enterprise Agents are deployed within the Cisco SD-WAN devices to provide detailed network performance data and insights, which can be managed through the vManage console.

Incorrect answers:

- A) Cloud Agents live on the ThousandEyes managed infrastructure
- B) ThousandEyes has a browser extension but is not considered an agent and is used to monitor user browser sessions from the Endpoint Agent
- D) Endpoint Agents cannot be deployed to Cisco SD-WAN devices

## 1.5 Question 2

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

### Explanation

The correct answer is **B) To monitor external applications and services from SD-WAN sites** because the integration of ThousandEyes with Meraki enables distributed organizations to effectively monitor external applications and services from their SD-WAN sites.

Incorrect answers:

- A) This is not a supported deployment type for the Meraki integration
- C) Monitoring services and applications can be viewed as a means to meet security and compliance standards
- D) The primary function is to monitor the remote employee connecting through these branches, enforcing access policies is outside the scope of the Meraki-ThousandEyes integration

## 1.5 Question 3

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

### Explanation

The correct answer is **C) Network performance data from the user's device** because when integrated with Cisco Secure Client, ThousandEyes collects network performance data directly from the user's device to monitor and diagnose issues related to remote connectivity, including VPN performance.

Incorrect answers:

- A) While network hardware configurations are important, ThousandEyes primarily focuses on network performance metrics rather than hardware configuration details.
- B) ThousandEyes is a network intelligence platform that monitors network performance. It does not focus on user activity and behavior which would be more relevant to user monitoring.
- D) ThousandEyes doesn't use data related to secure web gateway performance when integrated with Cisco Secure Client.

## 1.5 Question 4

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.**

#### Explanation

The correct answer is **D) It allows for quick identification and resolution of performance issues.** because the integration of ThousandEyes with Cisco technologies provides comprehensive monitoring and visibility into network performance, enabling IT teams to quickly identify and troubleshoot issues.

Incorrect answers:

- A) While ThousandEyes can complement network management by providing visibility into network performance, it is not specifically a network management tool for configuring and managing Cisco routers and switches.
- B) ThousandEyes does provide visibility into networks within and outside the network engineer's control which helps identify security attacks such as DDoS, BGP Hijacks and Route Leaks quickly but it does not automate the prevention of such attacks.
- C) While automation may be a feature of some Cisco technologies, the integration of ThousandEyes itself does not automate network configuration changes based on user feedback.

## 1.5 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



## Explanation



The correct answer is **B) Network Path** because this network visualization pane displays what is known as path visualization in ThousandEyes, where the hop-by-hop network nodes from the device to the Webex server are shown.

Incorrect answers:

- A) In this section you can add or manage all the devices for your organization that are assigned to a user or part of a place.
- C) This is the section for user management.
- D) In this section you can configure and customize organization-wide settings for Webex.

# Metric Baseline

## 1.6 Question 1

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

### Explanation

The correct answer is **C) SNMP data collected from the routers** because SNMP is specifically designed for collecting device-level metrics like CPU utilization, memory usage, and interface statistics.

Incorrect answers:

- A) DNS resolution is an unrelated metric to device performance.
- B) HTTP is an unrelated metric to device performance.
- D) Path visualization is a graphical representation of the traces sent from the agent to the destination.

## 1.6 Question 2

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.



## Explanation



The correct answer is **C) Capturing both peak and off-peak traffic patterns for a representative view** because it's essential to capture the full spectrum of network behavior, including both peak and off-peak traffic patterns. This provides a more accurate and representative understanding of how the interface is utilized under different load conditions.

Incorrect answers:

- A) Choosing a time period with the lowest network traffic volume would not provide a representative baseline for interface utilization. It is important to capture data that reflects normal operating conditions, including both high and low traffic periods.
- B) Taking only the organization's financial year would not be an accurate representation to establishing a baseline for network interface utilization. The focus should be on capturing typical network usage patterns rather than aligning with financial cycles.
- D) While it may be convenient to limit the amount of data for analysis, doing so could result in an incomplete or skewed baseline that does not accurately reflect the network link's utilization. It is important to collect enough data over a sufficient time period to ensure the baseline is representative of actual usage.



# Integration Types

## 1.7 Question 1

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 their ITSM and generate an incident so it can be responded to appropriately. Select a valid built-in integration that could be used:

- ☒ **A) ServiceNow Integration**
- ☐ B) DNA Center Integration
- ☐ C) Custom Webhooks
- ☐ D) Alerts API

### Explanation

The correct answer is **A) ServiceNow Integration** because it is the only built-in option available.

Important note:

While option C) Custom Webhooks is another route to create alerts if ITSM is not already available as a built-in option, but the scenario does not mention that the ITSM is not available, making this the second-best choice.

## 1.7 Question 2

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

### Explanation

The correct answer is **B) OpenTelemetry** because:

- OpenTelemetry is used for streaming metrics from tests in real-time to an observability platform. These streams can be processed to create helpful graphs, charts, tables, or other visualizations.

Incorrect answers:

- A) While the API is extremely flexible and can give data to be presented in any method you see fit, it is not streamed in real-time. The API is queried only as frequently as it is configured to be called.
- C) The DNA Center Integration leverages the ThousandEyes API to present relevant data to be displayed within the DNA Center interface.
- D) Alert Thresholds are used to send alert notifications when their conditions are met. While these notifications are sent shortly after the alert thresholds are reached, they are not a real-time stream of the data produced by the tests.

## 1.7 Question 3

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**

### Explanation

The correct answer are **A) ServiceNow, B) PagerDuty, E) AWS, F) AppDynamics, and H) Slack**

Incorrect answers:

- C) MS Teams: Supported through custom webhooks, but not natively integrated at the time of writing
- D) Splunk: Alert notifications are possible via email A separate Splunk integration exists for data collection with OTEL, but not for alerts
- G) Webex: While a Webex Control Hub integration exists, it doesn't support alert notifications

# Cisco Network Assurance Platforms Selection

## 1.8 Question 1

You are a network engineer at a multinational corporation responsible for ensuring optimal performance and security across various environments, including remote hybrid workers, branch offices, and cloud services. Select a SaaS-based Network Assurance platform that enables comprehensive monitoring and visibility into hybrid worker activities, internet traffic, and branch office connectivity.

- ☐ A) Meraki Insights
- ☐ B) AppDynamics
- ☒ C) ThousandEyes
- ☐ D) Catalyst Center

### Explanation

The correct answer is **C) ThousandEyes** because:

- ThousandEyes provides Internet/ISP path visibility to applications and services wherever they are hosted. Endpoint Agents cover hybrid workers, while Enterprise Agents cover branch offices and cloud services.

Incorrect answers:

- A) Provides valuable insights for Meraki networks but may not comprehensively cover hybrid workers, internet traffic, and branch offices if different hardware/vendors are used.
- B) Focuses on application performance monitoring, not network assurance
- D) An on-premises solution, not SaaS-based

## 1.8 Question 2

As a network engineer, you need to select a network assurance platform that provides end-to-end visibility and metrics for remote workers accessing SaaS applications. The solution should monitor the user experience from the endpoint device, through the VPN, across the internet, and to the SaaS provider. Which platform is best suited for this use case?

- ☐ A) Catalyst Center
- ☐ B) AppDynamics
- ☐ C) Meraki Insights
- ☒ D) ThousandEyes

### Explanation

The correct answer is **D) ThousandEyes** because:

- ThousandEyes Endpoint Agents can monitor the user experience from the remote worker's device, providing metrics on WiFi, VPN, and system performance.
- Tests can measure network and application metrics from the endpoint across the internet to the SaaS provider, giving end-to-end visibility.

Incorrect answers:

- A) Focuses on managing on-premises Cisco networks and lacks the SaaS and internet visibility needed here.
- B) Is application-centric and would require integration with ThousandEyes for the network visibility aspects.
- C) Is valuable for Meraki-based networks but may not cover the end-to-end visibility from the remote endpoint to the SaaS provider if other network components are involved.

## 1.8 Question 3

Which network assurance platform is best for providing network visibility and performance across any network, where metrics can be correlated with application-level metrics, including for services in multi-cloud deployments?

- ☐ A) Catalyst Center
- ☐ B) AppDynamics
- ☐ C) Meraki Insights
- ☒ **D) ThousandEyes**



## Explanation



The correct answer is **D) ThousandEyes** because it's a network intelligence platform specifically designed for comprehensive network visibility. It monitors any network and correlates network metrics with application performance, making it ideal for campus networks and multi-cloud deployments.

Incorrect answers:

- A) Catalyst Center: Focuses on network management and automation, not network performance visibility
- B) AppDynamics: Primarily an application performance management (APM) tool. While it offers some network visibility, it's not its core strength
- C) Meraki Insights: Provides network insights but is often tied to Meraki hardware and may not be as comprehensive as ThousandEyes, especially in multi-cloud environments

# Domain 2 Answer Key

## Enterprise Agent Configuration

### 2.1 Question 1

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

#### Explanation

Cisco DNA Center can be used to install ThousandEyes Enterprise Agents on Cisco Catalyst 9300 and 9400 Series switches, allowing your IT team to easily monitor performance and quickly identify issues with critical services that your users rely on.

Additionally, Application Hosting with Docker via CLI is another great option for the installation of Enterprise agents.

### 2.1 Question 2

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

### Explanation

The ThousandEyes - Meraki integration allows users to install ThousandEyes Enterprise Agents on supported Meraki MX switches, providing 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.

## 2.1 Question 3

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

### Explanation

The correct answer is **C) Execute the full agent lifecycle.**

After the application is up and running, the agent (ThousandEyes-agent) process connects to the controller in the cloud environment. To apply configuration changes, you need to stop and deactivate the container before modifying the configuration. The full lifecycle (stop, deactivate, activate, start) ensures that the changes are properly applied.

Why the other options are incorrect:

- A) Simply restarting the container is insufficient; the app hosting application needs to be deactivated before making configuration changes.
- B) While reinstalling with the correct settings would work, it's an unnecessarily lengthy process. The existing agent can be reconfigured.
- D) The app hosting application must be stopped and deactivated; it will not automatically pick up the new configuration.

# Endpoint Agent Deployment

## 2.2 Question 1

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

## 2.2 Question 2

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

### Explanation

Group Policy Objects (GPOs) only apply to computers within the designated domain. Before troubleshooting further, verify the PC is part of the correct domain.

The other options are incorrect because:

- A) EPA installation via GPO isn't dependent on specific user logins.
- C) Rebooting the PC won't restart server components or affect GPO deployment.
- D) Restarting the server is irrelevant if the PC isn't a member of the domain where the GPO is deployed.

## 2.2 Question 3

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

#### Explanation

Option B is the most efficient method. Options A, C, and D are slower and less efficient processes for enterprise-scale deployment.

# Test Configuration

## 2.3 Question 1

Refer to the exhibit. Which setting should be used for this network Agent to Server test to prevent firewalls from detecting the test traffic as malicious?

- ☒ A) Path Trace Mode: In Session
- ☐ B) Protocol: TCP
- ☐ C) Port: 80
- ☐ D) Probing Mode: Force SYN

Target: 10.10.10.10

Protocol: TCP Port: 80

Probing Mode: **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](#)

*Exhibit 2.3-1: Network Agent to Server Test Configuration*

## Explanation

The **Path Trace Mode: In Session** setting establishes a TCP connection with the target and uses the same connection for sending path traces, preventing firewalls from flagging the test traffic as potentially malicious.

- B is incorrect: Protocol TCP only specifies the packet protocol and doesn't address firewall concerns.
- C is incorrect: The port number (80) is irrelevant to preventing firewall detection of test traffic.
- D is incorrect: Probing Mode: Force SYN is a fallback mechanism to Prefer SACK and doesn't help with the requirement.

**Note:** In the image, the "Path Trace Mode: In Session" option is unchecked. Enabling this option would be the correct action to take.

## 2.3 Question 2

Refer to the exhibit. 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

Basic Configuration

Advanced Settings

Domain

internal-domain.com

IN

A

Interval

10 minutes

Agents

1 of 8 selected

DNS Servers

ns47.domaincontrol.com. x

1 more x

Q Lookup Servers

Alerts

☒ Enable

2 of 5 alert rules selected

Edit Alert Rules

Exhibit 2.3-2: DNS Server Test Configuration

## Explanation

Correct: **B) Create a DNS Server test monitoring the AAAA record** and **D) Create a DNS Server test monitoring the NS record**.

- The **AAAA record** resolves to the IPv6 IP address, and the DNS Server test is the only test that provides the resolution time metric.
- The same applies to the **NS record**, which provides name server resolution times.

Incorrect:

- **A) Create a DNS Server test monitoring the A record:** The A record resolves to an IPv4 address, not IPv6.
- **C) Create a DNS Trace test monitoring the ANY record:** The ANY query type retrieves all available records, but the resolution time metric is not available in the DNS Trace test.
- **E) Create a DNS Trace test monitoring the NS record:** The resolution time metric is not available in the DNS Trace test.

## 2.3 Question 3

Refer to the exhibit. An engineer is trying to configure a Page Load test and is trying to assign the "east1-agent-1" to run it. What is the reason?

- ☐ 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**




Exhibit 2.3-3: Page Load Test Configuration

### Explanation

The correct answer is **D) The agent does not support Page load tests.**

To run Page Load, Transaction, and API tests, the agent must have BrowserBot installed. If it doesn't, it will be unavailable when selecting agents in the test settings.

Incorrect:

- **A & B) The agent is not running/disabled:** If the agent is not running or disabled, it will show the label "(disabled)" in addition to being grayed out.
- **C) The agent is still registering:** The agent would not show up on the list of agents if it was still registering to the platform.

## 2.3 Question 4

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

#### Explanation

The **Transaction test type** is the most useful for this scenario. This type of test can mimic user interactions with a website, allowing you to identify precisely where the latency occurs. Since the problem is related to performance while logging in or adding items to the cart, a transaction test can be configured to follow a user's journey and pinpoint the source of the issue.

## 2.3 Question 5

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

#### Explanation

The answer is **Enterprise Agent** and **Agent-to-agent test type**. This combination is ideal because it allows for the monitoring of specific network metrics between the branch offices and the data center. The Enterprise Agent is the best choice for this scenario as it provides detailed insights into network performance, while the agent-to-agent test type is the most suitable for measuring communication and network performance between two specific locations.

# Endpoint Agent Tests

## 2.4 Question 1

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




Exhibit 2.4-1: Endpoint Agent Test Configuration

## Explanation

The correct answers are **A) Create a new custom application monitor** and **C) Add a new scheduled test to the monitor**.

To monitor a web application from the employee's point of view using an Endpoint Agent, you need to:

1. Create a custom application monitor to define the specific web application you want to track.
2. Add a scheduled test to this monitor to regularly check the application's performance.

Incorrect options:

- B) Create a new google suite monitor: The Google suite application monitor is a pre-defined template with 3 Scheduled HTTP server tests to monitor meet.google.com, mail.google.com, and docs.google.com, plus 1 Scheduled Network test to Google's DNS. This is not suitable for monitoring a custom web application.
- D) Add a new dynamic test to the monitor: Dynamic tests are specifically designed to monitor traffic for certain applications like Microsoft Teams, Webex, or Zoom. They are not appropriate for monitoring a custom web application.
- E) Add a new test template: All monitor applications are already templates. Adding a new template is not an action that would directly contribute to monitoring the web application.

## 2.4 Question 2

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



### Explanation

The correct answer is **A) \*-corporate**. Here's why:

- **Wildcard Usage:** ThousandEyes supports the use of wildcards (like `*`) in label filters. The `*` symbol represents any string of characters.
- **Matching the Pattern:** You want to match any agent name ( `agentname` ) followed by a hyphen and the word "corporate" ( `-corporate` ). The filter `*-corporate` achieves this by using the wildcard to match any characters before "-corporate".

**Let's break down why the other options are incorrect:**

- B) `agentname-*`: This would match any agent with the prefix "agentname-" followed by any characters, not specifically those connected to the corporate network.
- C) `agent*corporate`: This filter is too broad; it would match any agent with "agent" somewhere in the name, followed by any characters, and ending with "corporate".

## 2.4 Question 3

What type of endpoint agent test will gather browser activity?

- ☐ A) Scheduled tests
- ☐ B) Dynamic tests
- ☒ **C) Real user tests**
- ☐ D) Network Access tests

### Explanation

The correct answer is **C) Real user tests**.

Real user tests in ThousandEyes utilize a browser plugin to capture user interactions and performance metrics as users navigate websites within defined monitored domain sets. This provides insights into actual user experience with web applications.

## 2.4 Question 4

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

#### Explanation

The correct answer is **C) Scheduled, dynamic, and real user tests**.

ThousandEyes provides multiple options for monitoring applications like Microsoft Teams:


- **Scheduled tests** establish a performance baseline by running HTTP server or network tests at regular intervals.
- **Dynamic tests** automatically trigger tests to remote servers when specific applications (like Teams) initiate network connections.
- **Real user tests** capture the actual user experience with Teams during live sessions using the browser plugin.

# Synthetic Web Tests

## 2.5 Question 1

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>One Drive - 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



```
1 import { By, Key, until } from 'selenium-webdriver';
2 import { driver, markers, credentials, downloads, transaction, test } from 'thousandeyes';
3
4 runScript();
5
6 async function runScript() {
7
8   const settings = test.getSettings();
9
10  // Load page
11  await driver.get(settings.url);
12  await driver.takeScreenshot();
13
14  // Search
15  markers.start('SearchForWebdriver');
16  await driver.findElement(By.name('q')).sendKeys('webdriver', Key.RETURN);
17  await driver.takeScreenshot();
18  markers.stop('SearchForWebdriver');
19
20  // Wait for full page load
21  await driver.wait(until.titleIs('webdriver - Google Search'), 1000);
22  await driver.takeScreenshot();
23 };
24
```

Install ThousandEyes Recorder on your machine

Exhibit 2.5-1: Transaction Script Example

### Explanation

All are valid answers, hence D is the right answer.

Further explanations:

- A) There is a script sample specific to this use case.
- B) The recorder IDE can help you get started by recording the user workflow on your personal PC and translate it to JavaScript.
- C) This Github repository is maintained and owned by ThousandEyes. Useful script samples can be found to match your use case.

## 2.5 Question 2

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

### Explanation

The correct answer is **B) Transaction monitoring**.

While the other options provide valuable network-level insights, **transaction monitoring** is specifically designed to simulate and measure multi-step user workflows within a web application. For an e-commerce checkout process, transaction tests can simulate actions such as adding items to a cart, entering payment information, and completing the purchase. This provides a comprehensive understanding of performance and user experience throughout the entire checkout flow.

## 2.5 Question 3

True or False: ThousandEyes Synthetic Tests eliminate the need for any real user monitoring on your online learning platform.

- ☐ A) True
- ☒ **B) False**



### Explanation



False. While Synthetic tests provide valuable insights, they cannot capture the full range of real user behavior. Consider using Real User Monitoring (RUM) tools in conjunction with Synthetic Tests for a more comprehensive picture.

# Web Authentication

## 2.6 Question 1

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. What must be done to meet the requirements? (Select 2)

- ☐ 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

HTTP AUTHENTICATION

Scheme

NoneBasicNTLMKerberosOAuth

HTTP REQUEST

HTTP Version

Prefer HTTP/2HTTP/1.1 only

Request Method

GETPOST

POST Body

Follow Redirects

☒ Enable

User Agent

Default

Custom Headers

Root Request

Enter HTTP headers...

—

+

Exhibit 2.6-1: HTTP Authentication Options

### Explanation

The correct answers are **D) Parameters are not supported by HTTP server OAuth authentication; use a Transaction script instead** and **E) Parameters are not supported by HTTP server OAuth authentication; use an API test instead**.

The question describes an OAuth flow that requires sending parameters in the initial token request. This is not supported by the HTTP Server test type.

- **Basic** and **NTLM** authentication only use username/password credentials.
- While the HTTP Server test's **OAuth** configuration can handle token-based authentication, it doesn't allow for parameters in the token request.

**Transaction scripts** and **API tests** provide the necessary flexibility to handle this flow. Transaction scripts allow you to script the entire process, while API tests are specifically designed for interacting with APIs and support parameters in requests.

## 2.6 Question 2

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.**

### Explanation

The correct answer is **D) Exclude the MFA step from the transaction test and focus only on the SAML login**.

ThousandEyes transaction tests are not designed to interact with external authentication mechanisms like mobile app-based OTPs. The most practical approach in this scenario is to exclude the MFA step from the test and focus on monitoring the SAML-based SSO login process. This provides valuable insights into the performance and availability of the SSO system without the complexities of handling dynamic OTPs.

## 2.6 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) 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.

### Explanation

The correct answers are **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**, and **D) Contact the web application vendor to report the issue and inquire about possible server-side problems**.

When troubleshooting intermittent failures in a transaction test using Basic Authentication, it's essential to take a multi-faceted approach:

- **Credential Verification:** Ensure the credentials used in the test are accurate by manually logging in from different locations. This rules out any typos or location-specific access issues.
- **Waterfall Chart and Response Code Analysis:** Examine the ThousandEyes waterfall charts and HTTP response codes for clues. Slow-loading resources, server errors, or authentication-related errors can point to the root cause.
- **Vendor Communication:** If the issue persists after verifying credentials and analyzing ThousandEyes data, contact the web application vendor. They may have insights into server-side problems or recent changes that could be affecting the test.



# Domain 3 Answer Key

## Network Issues

### 3.1 Question 1

Users at a remote corporate site (site 30 or "s30") 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?




Exhibit 3.1-1: Before Issue




Exhibit 3.1-2: After Issue

View in [ThousandEyes](#)

- ☐ 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.

### Explanation

From the figure, we can observe that the spike in latency is caused by the link between devices 10.84.30.1 and 10.87.10.253. Comparing the discovered network path between prior the incident and during the incident, we can confirm that no routing changes occurred as traffic always goes through these nodes. This is likely a network congestion or traffic load condition.

A valid next step is to review the bandwidth utilization and QoS settings at this site, to identify any possible network congestion conditions.

## 3.1 Question 2

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




Exhibit 3.1-3: Before Incident #1



Exhibit 3.1-4: Before Incident #2




Exhibit 3.1-5: During Incident #1




Exhibit 3.1-6: During Incident #2

View in [ThousandEyes](#)

- ☐ A) Involve the Voice team as the RTP test does not return any relevant results for the agent located on site 20 “s20”
- ☒ **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.

#### Explanation

From the figures, we can observe that in normal conditions, traffic is forwarded from node 10.87.7.51 to 10.84.50.53. During the incident, there is forwarding loss observed at node 10.87.7.51. A Cisco ThousandEyes Enterprise Agent will display the “-” character in the Table view when it is unable to complete measurements for a test.

This doesn't invalidate the test; it shows that data collection from one target agent wasn't completed. Instead of discarding the test, we should focus on other layers. Other agents measure expected MOS to this target, except the one at site 20, there may be issues specific to the target agent at site 20 that need attention.

# End-device Issues

## 3.2 Question 1

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 ?




Exhibit 3.2-1




Exhibit 3.2-2




Exhibit 3.2-3




Exhibit 3.2-4

- ☐ 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.

Let's break down why the correct answer is the most indicative of a DNS issue that happened:

- A) The test target stopped responding - the error says "The host name could not be resolved", the test target response happens after it is resolved. This is not a correct answer.
- B) The FQDN of the test target is non-existent – this answer is not correct because the FQDN exists in the previous test round.
- D) The DNS settings on the endpoint are incorrect – this answer is also incorrect because in the previous round the end-point was using the same DNS settings and the issue did not happen.
- **Correct answer: C)** The DNS servers assigned to the endpoint are unreachable. This is likely because the endpoint is utilizing external DNS servers (8.8.8.8, 8.8.4.4), but based on the exhibit, external tracing is absent, meaning external resources are unreachable.

## 3.2 Question 2

The Endpoint stopped appearing online after it was moved to another network.



Exhibit 3.2-5

The customer reviewed the endpoint logs but did not identify anything suspicious.

```

2024-05-07 13:20:12.452 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:206 - RECONNECT is currently scheduled or running, ignoring schedule attempt for CONNECT
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService@DefaultAgentSettingsService.cpp:496 - Checkin failed with error code: 1
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService@DefaultAgentSettingsService.cpp:558 - Calling connect on persistent connection
2024-05-07 13:20:14.468 DEBUG [3068.8116] agent.services.AgentSettingsService@DefaultAgentSettingsService.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 10m
2024-05-07 13:20:14.468 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.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.DefaultPersistentConnection@DefaultPersistentConnection.cpp:206 - RECONNECT is currently scheduled or running, ignoring schedule attempt for CONNECT
2024-05-07 13:20:24.753 DEBUG [3068.8248] system.EtwTraceConsumer@DefaultEtwTraceConsumer.cpp:108 - Flushed ETW trace buffers to file. path="C:\ProgramData\ThousandEyes\Endpoint Agent\etwTraces\ThousandEyes Endpoint Agent.etl", size=320kB
2024-05-07 13:20:24.765 DEBUG [3068.8248] net.SocketMonitoring@MinSocketMonitoringService.cpp:335 - Processed ETW event batch: published=2, stale=213, deduped=0, duplicatesDueToCoalesceInterval=0, skipped=41
2024-05-07 13:20:24.812 DEBUG [3068.7712] scheduler.task.statscommitter@UsageStatsCommitter.cpp:43 - Starting UsageStatsCommitter task
2024-05-07 13:20:24.812 DEBUG [3068.7712] tests.scheduledtestsmetric@DefaultScheduledTestsMetricsService.cpp:54 - End metric tracking period at: 2024-May-07 13:20:24.812000
2024-05-07 13:20:25.109 DEBUG [3068.2148] scheduler.task.statscommitter@UsageStatsCommitter.cpp:95 - Updated usage stats object with 69 entries.
2024-05-07 13:20:25.109 DEBUG [3068.2148] scheduler.simpler@DefaultTaskScheduler.cpp:173 - Scheduled task: te:agent:UsageStatsCommitter/Dx1, delay=60000 ms
2024-05-07 13:20:28.499 DEBUG [3068.8248] net.MinHttpClient@MinHttpClient.cpp:2115 - Request to: wss://cl.eb.thousandeyes.com/relay/connect timed out: 12002: The operation timed out
2024-05-07 13:20:28.499 DEBUG [3068.8248] net.MinHttpClient@MinHttpClient.cpp:2115 - Request to: https://cl.eb.thousandeyes.com/status.json timed out: 12002: The operation timed out
2024-05-07 13:20:28.499 DEBUG [3068.4352] net.DefaultWebSocketStompConnection@DefaultWebSocketStompConnection.cpp:135 - WebSocket connection wait: 4210ms
2024-05-07 13:20:28.499 ERROR [3068.4352] net.DefaultWebSocketStompConnection@DefaultWebSocketStompConnection.cpp:148 - asyncOpen (WebSocket): Input/output error [generic:5]
2024-05-07 13:20:28.499 ERROR [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:317 - asyncOpen: Input/output error [generic:5]
2024-05-07 13:20:28.499 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:485 - Running post-run checks for RECONNECT, next=NONE, retry=true
2024-05-07 13:20:28.499 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 13:20:28.499 DEBUG [3068.4352] agent.DefaultPersistentConnection@DefaultPersistentConnection.cpp:300 - Set schedule delay=1500ms, jitter=32ms, nextDelay=2250ms
2024-05-07 13:20:28.499 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@CheckinService.cpp: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 13:20:28.771 DEBUG [3068.8072] api.userSpaceProxyClient@UserSpaceProxyClient.cpp:58 - 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@DefaultWebSocketStompConnection.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

```

Exhibit 3.2-6

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?**



- ✓ 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.**

### Explanation

According to the log message, there was a timeout when attempting to connect to c1.eb.thousandeyes.com:

```
2024-05-07 13:20:28.498 DEBUG [3068.2148] net.WinHttpClient@WinHttpClient.cpp:2115 - Request
to: wss://c1.eb.thousandeyes.com/relay/connect timed out: 12002: The operation timed out
2024-05-07 13:20:28.499 DEBUG [3068.8240] net.WinHttpClient@WinHttpClient.cpp:2115 - Request
to: https://c1.eb.thousandeyes.com/status.json timed out: 12002: The operation timed out
```

Since the Endpoint was relocated to the new network with static IP assignment, troubleshooting should commence with verifying the accuracy of the IP credentials and ensuring connectivity with ThousandEyes.

- Option A is incorrect because the new endpoint should function properly in the new network like other endpoints.
- Option B is incorrect because reinstalling the software does not address IP issues.
- Option C is incorrect because incorrect IP credentials manually entered will not resolve automatically.

# Web Application Performance

## 3.3 Question 1

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




Exhibit 3.3-1



Exhibit 3.3-2

- ☐ 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

## Explanation

This one is a tricky one, as it requires you to have a basic understanding of HTTP response codes (see Resources), let's have a look at each potential answer.

- **A) Modify the firewall rules to allow connections to the target domain** is incorrect because our exhibit shows we are getting a response from the target server
- **C) Change the HTTP request method to PATCH** is incorrect because nothing in the response header indicates that the request method is incorrect (that would be code `405 Method Not Allowed`)
- **D) Modify the target URL to an available API endpoint** is incorrect because nothing in the response points to an unavailable API endpoint

Finally, if we check the response code we are getting from the server, 401, we will find it to be a response to Unauthorized requests. Further inspection into the request headers will confirm the issue, as no Authentication header is being sent, thus, **Modify the authentication credentials** is the answer.

### 3.3 Question 2

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




Exhibit 3.3-3

| Request Headers                                                                                                                                                                                                                                                                                                                                                                           | Response Headers                                                                                                                                                                                                          |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre> :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 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/118.0.5993.70 Safari/537.36 x-thousandeyes-agent: yes </pre> | <pre> 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 </pre> |

Exhibit 3.3-4

- ☐ 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

### Explanation

This one will also leverage your knowledge of HTTP response codes, albeit with a twist.

- **There is a network connectivity problem preventing us from reaching the target URL** is incorrect because our exhibit shows we are getting a response from the target server
- **One of the DOM elements cannot be found in the server** is incorrect because even though the waterfall chart is marking component 81 with an issue, we don't see any 404 Not Found response
- **The request timed out waiting for the server to respond** is incorrect because the exhibit shows the server answering promptly to almost all component requests

Finally, if we check the response code we are getting from the server, we will find it to be a `302 Found` redirect. The fact that each redirect is leading to another redirect (multiple 302 responses in a row) points to a misconfiguration on the server that is causing a loop within the app. Thus, the last option is correct:

# Security Issues

## 3.4 Question 1

In real-life applications using ThousandEyes, you can switch between various views. However, for the exam, you will be limited to up to three exhibits. When reviewing answer options, remember to

- Analyze using only the provided exhibits.
- Choose the answer that can be confirmed with the information given.

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




Exhibit 3.4-1: Los Angeles before the Outage




Exhibit 3.4-2: Los Angeles during the Outage

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

#### Hint

- Analyze the details and contrast the provided exhibits to accurately identify potential network issues. Note any changes in Autonomous System (AS) numbers, which are crucial for determining the cause of network problems.
- If there are multiple agents visible in the path visualization view showing packet or forwarding loss, focus on one agent and compare its path against subsequent exhibits to determine the root cause.

### Explanation

While the exhibits clearly depict a network issue with significant packet loss, pinpointing the exact cause as a BGP hijack requires careful analysis. Let's break down why the correct answer is the most indicative of a BGP hijack:

- **Availability Drop:** Although a drop in availability is a symptom of the problem, it doesn't specifically point to BGP hijacking. Various network issues could cause availability drops.
- **HTTP Server response delay:** Similar to availability drop, this is a symptom of the problem, likely caused by the packet loss, but it doesn't explicitly indicate BGP hijacking.
- **Packet Loss:** Again, this is a clear symptom shown in the exhibits but doesn't directly confirm BGP hijacking as the cause.

However, the change in AS path from AS 16509 to AS 10297 is a strong indicator of BGP hijacking. This suggests that the route to the destination was illegitimately taken over by another AS, causing traffic to be misdirected and resulting in packet loss.

The shift in the AS path provides the most concrete evidence supporting the possibility of a BGP hijack in this scenario.

## 3.4 Question 2

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



## Explanation



The exhibits show a significant packet loss issue occurring at a specific point in the network path. The Path Visualization highlights a node within AS 10297 as the source of 100% forwarding loss for multiple agent locations. This suggests a problem beyond the local network and points towards a potential BGP routing issue, specifically a BGP hijack.

Thus, the correct answer is B) Reach out to the ISP to report the suspected BGP hijacking incident

Here's why the other options are not the best next steps:

- **Contact the internal network team to investigate potential misconfigurations on the local routers** is incorrect because the issue appears to be external to the local network, as multiple geographically dispersed agents are affected, and the packet loss originates from AS 10297.
- **Implement traffic filtering rules on the firewall to block traffic originating from AS 10297** is incorrect as implementing firewall rules would not address the root cause, which is likely a routing issue outside of the local network's control.
- **Restart the DNS server to refresh its cache and potentially resolve the observed issue** is incorrect because restarting the DNS server is unlikely to resolve a BGP hijacking issue, as the problem lies within the routing of traffic rather than the DNS server itself.



# Domain 4 Answer Key

## Network Condition Alert Rules

### 4.1 Question 1

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

- ☒ **A) Response Time**
- ☐ B) BGP Reachability
- ☒ **C) Error Type**
- ☐ D) Interface Throughput

#### Explanation

For Endpoint Agent HTTP Server tests, valid alert rule metrics include Response Time and Error Type. BGP Reachability and Interface Throughput metrics are not applicable to this test type. Response Time measures time-to-first-byte while Error Type allows alerting on specific HTTP errors.

### 4.1 Question 2

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

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](#)

Exhibit 4.1-1

## 4.1 Question 3

Refer to the exhibit. The alert rule is set up as shown, but didn't trigger. Why?

- ☒ A) Alert conditions weren't met and won't trigger with current setup
- ☐ B) Alert needs two consecutive agent failures to trigger
- ☐ C) Response code is set up incorrectly
- ☐ D) All of the above

Settings

Notifications

GENERAL

Rule Name

Response code 4XX

Tests

1 of 21 test(s) selected

Agents

All agents

Severity

Info

Minor

Major

Critical

ALERT CONDITIONS

All conditions are met by

any of

2

agents

1

of

1

time in a row:

Response Code

is

any error (≥ 400 or no response)

Exhibit 4.1-2




Exhibit 4.1-3



## Explanation



The correct answer is option **A) Alert conditions weren't met and won't trigger with current setup** because the condition requires 2 agents to generate an error.

Incorrect Options:

- B) Alert will trigger when the condition is met for a single test round
- C) Response code is correctly configured and should trigger for the HTTP code 401 shown in the exhibits
- D) Does not apply

## 4.1 Question 4

Refer to the exhibit. 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)

## Add New Alert Rule

Alert Type
Web
HTTP Server

Rule Name
ENNA-HTTP-Alert

Settings
Notifications

Tests
0 of 29 test(s) selected

Agents
All agents

Severity
Info
Minor
Major
Critical

### ALERT CONDITIONS

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

?

?

?

-

+

Exhibit 4.1-4

### Explanation

The correct answer is **D) Response Code is server error(5XX)**. This is the most specific and relevant condition for the given scenario.

Incorrect Options:

- A) HTTP Error: Too broad, capturing any error in the HTTP process (HTTP, Receive, Wait, SSL, etc.). While it includes HTTP server errors, it's not specific enough.
- B) Time to First Byte: Measures the duration between completing a browser request and receiving the first byte of the server's response. Not related to server error codes.
- C) Response Time: Measures overall request-response duration. Not directly related to HTTP response codes.

# End-User Experience Alert Rules

## 4.2 Question 1

Refer to the exhibit. A network engineer is tasked with configuring an alert that will trigger if the Endpoint Agent path ASN changes on a specific hop. What is the alert type and condition needed to meet the requirement?

- ☒ A) Scheduled tests, Hop#
- ☐ B) Real User Test, Hop#
- ☐ C) Scheduled tests, Any Hop
- ☐ D) Real User Test, Path Length

Add New Alert Rule

Alert Type

Endpoint Path Trace

Alert Rule Name

Settings

Notifications

Tests

0 of 29 test(s) selected

Agents

All agents

Severity

Info

Minor

Major

Critical

ALERT CONDITIONS

All conditions are met by at least 1 agent and 30 % of agents 1 of 1 time in a row:

1 from the source meets all of:

Delay

≥

100

ms

Collapse All

Exhibit 4.2-1

### Explanation

Option **A) Scheduled tests, Hop#** is correct because:

- Scheduled tests are suitable for monitoring consistent changes in the network path, like an ASN change on a specific hop. Real User Tests are better suited for monitoring real-time user experience fluctuations.
- Hop# is the appropriate condition as the requirement explicitly mentions monitoring a specific hop for the ASN change.

#### **Incorrect Answers:**

- B) Real User Test, Hop#: Real User Tests are not the best choice for this scenario as they are designed for monitoring real-time user experience, not network path changes.
- C) Scheduled tests, Any Hop: This option would trigger the alert if the ASN changes on any hop, not just a specific one.
- D) Real User Test, Path Length: This option would trigger the alert if the total number of hops in the path changes, not if the ASN changes on a specific hop.

## 4.2 Question 2

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  $\geq$  %

### Explanation

Option **C) Real User Tests > Endpoint, CPU utilization  $\geq$  %** is correct because:

- This combination directly targets the user's device (Endpoint) where the suspected CPU issue resides.
- It utilizes Real User Tests, which gather real-time performance data during user activity, providing the most accurate representation of the issue's impact.
- It specifically monitors for CPU utilization exceeding a defined threshold ( $\geq$  %), allowing for alerts to be triggered when CPU usage reaches problematic levels.

## Dashboard Configuration

The following sample questions require you to analyze data presented in two ThousandEyes dashboards used to monitor its application service at <https://thousandeyes.com>:

- **Executive Dashboard:** This dashboard ([link](#)) provides a high-level overview of application performance.
- **IT Operations Dashboard:** This dashboard ([link](#)) offers granular insights for troubleshooting and performance optimization.

Refer to the data in these dashboards to answer the questions below.


### 4.3 Question 1

Which type of test are we using for these dashboards?

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



Observe the widgets on both dashboards to determine the test type.




## 4.3 Question 2

Which type of widgets were used in the executive dashboard? (Select all that apply)

- ☐ A) Agent status
- ☒ B) Map
- ☐ C) Line
- ☒ D) Number
- ☐ E) Color Grid

## Explanation

The executive dashboard uses the **map** and **number** widgets.




### 4.3 Question 3

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

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

## Explanation

The Mexico City Mexico (TelMex) Agent displays a Connect time of 0.51 ms.



4.3 Question 4

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



Explanation



The alert rule is displayed at the beginning of the dashboard, indicating the trigger reason.

| Total Alerts 10 of 27 • Active Alerts 1 |                      |                 |                     |          |
|-----------------------------------------|----------------------|-----------------|---------------------|----------|
| Alert Source                            | Alert Rule           | Alert Type      | Start Time (CST)    | Duration |
| • ThousandEyes web Page                 | TE Page Load Latency | Web - Page Load | 05/09/2024 15:12 PM | 18s      |
| ThousandEyes web Page                   | TE Page Load Latency | Web - Page Load | 05/09/2024 14:12 PM | 1m 0s    |
| ThousandEyes web Page                   | TE Page Load Latency | Web - Page Load | 05/09/2024 14:14 PM | 1m 0s    |


4.3 Question 5

In the executive dashboard, what is the page completion time for the Mexico City agent?

- ☐ A) 100%
- ☐ B) 83.4%
- ☐ C) 15.2%
- ☒ D) 99.67%

## Explanation

Move your cursor over the map widget to the Mexico City agent to view the page completion time.



### 4.3 Question 6

In the executive dashboard, what is the total error count for ThousandEyes web page in the last 15 days?

- ☐ A) 520
- ☐ B) 1.58
- ☒ C) 4610
- ☐ D) 4805

## Explanation

The number widget displays the total **TE error count** for the last 15 days.

### 4.3 Question 7

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



### Explanation



Place your mouse over the latest metrics for page load time and DOM load time, then subtract the DOM load time from the page load time ( $1479.6 - 1354.3 = 125.3$  ms).

...

## 4.3 Question 8

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

# Alert Configuration and Functionality

## 4.4 Question 1

An alert rule for a Web - HTTP Server test is not triggering when the HTTP response code is 500 Internal Server Error. The alert conditions are configured with "Response Code" set to "any error ( $\geq 400$  or no response)". What could be causing the alert to not fire?

- ☐ A) The alert rule is disabled
- ☐ B) The test is not enabled on any Enterprise Agents
- ☒ C) The alert rule's "Settings" section does not have the correct test selected
- ☐ D) The HTTP server is returning a 200 OK response code

### Explanation

The correct answer is:

- C) The alert rule's "Settings" section does not have the correct Web - HTTP Server test selected. This means the alert conditions are not being evaluated against the test's data, so the 500 errors are not triggering the alert.

Incorrect options:

- A) A disabled alert rule would prevent alerts, but this would be obvious in the Alert Rules page.
- B) If the test was not assigned to any agents, it would never generate data to trigger alerts. However, this would likely be noticed when viewing the test.
- D) A 200 OK response would not trigger the " $\geq 400$ " alert condition, but the question states 500 errors are occurring, so this is not the issue.

## 4.4 Question 2

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



## Explanation



To make a CPU utilization Endpoint Agent alert less sensitive and reduce noise, the correct options are:

- A) Increasing the number/percentage of agents that must exceed the CPU threshold will prevent a single agent from triggering the alert.
- C) Requiring more rounds of data to be above the threshold (e.g. 2 of 3 rounds instead of 1 of 1) will filter out brief CPU spikes.

The incorrect options that would not reduce alert sensitivity are:

- B) Lowering the CPU utilization percentage would make the alert more sensitive and trigger more frequently.
- D) Enabling the alert on more agents would potentially trigger it more often, not less.

# Network Capacity Planning

## 4.5 Question 1

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

### Explanation

The correct answer is C) Tune the existing QoS configuration to prioritize voice and video traffic.

- This option directly addresses the observed issues (delay and jitter spikes) during peak hours.
- It aligns with the scope of the exam, which includes QoS tuning but not complete redesigns.
- This solution is targeted and likely more cost-effective than other options.

Incorrect options:

- A) A complete QoS redesign is out of scope for this exam and may be unnecessary.
- B) Increasing bandwidth on all links is a costly solution that may not specifically address the voice and video traffic issues.
- D) Replacing all network hardware is an extreme and costly solution that may not directly solve the problem.

## 4.5 Question 2

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



### Explanation

The correct answer is B) Change the access point to a different, less congested channel.

- This directly addresses the high channel utilization issue.
- Reducing channel congestion can decrease retransmissions and improve overall voice call quality.
- This solution is a targeted optimization based on the SNMP data provided.

Incorrect options:

- A) Increasing transmit power may exacerbate interference issues and doesn't address channel congestion.
- C) Disabling all non-voice traffic is an extreme measure that could negatively impact other necessary network functions.
- D) Strict admission control for all clients doesn't specifically target the channel utilization issue and may be too restrictive.

## 4.5 Question 3

CLI outputs show that a router's egress queue for voice traffic is consistently full, leading to increased latency. Based on this data, which optimization would you recommend?

- ☐ A) Increase the queue size for voice traffic
- ☒ **B) Implement traffic shaping on non-voice traffic**
- ☐ C) Disable QoS on the router to allow all traffic equal priority
- ☐ D) Replace the router with a higher-capacity model

## Explanation

The correct answer is B) Implement traffic shaping on non-voice traffic.

- This solution addresses the root cause by managing non-voice traffic to prevent it from overwhelming the voice queue.
- It's a targeted optimization that can reduce latency for voice traffic without major hardware changes.
- This approach aligns with QoS tuning, which is within the scope of the exam.


Incorrect options:

- A) Increasing queue size may delay packets further and doesn't address the underlying issue of queue saturation.
- C) Disabling QoS would likely worsen the situation for voice traffic, which requires prioritization.
- D) Replacing the router is an expensive solution that may not be necessary if the issue can be resolved through configuration changes.

## 4.5 Question 4

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



4.5-1: Capacity Planning Results

## Explanation

This is a tricky question. Is there really an issue with the data being presented or perhaps there is something misconfigured on the platform?

Our exhibit shows that our highest consumption, although marked at 97%, is merely 48Mbps, certainly not enough to be making use of the entire 1Gbps connection from the ISP, so option A would be incorrect. Even though our top traffic is indeed HTTP, nothing in the exhibit indicates that pruning some specific HTTP traffic could fix how data is being presented, so option C is incorrect. The exhibit also fails to provide any reason as to how changing business hours could provide some benefit in this case, so option D is incorrect.

Finally, if we gather all the data we have available: ISP connection is 1Gbps and capacity planning marks 48Mbps as 97% of max capacity, we can reach the conclusion that the max capacity for this interface is misconfigured; it should be set to 1Gbps instead of the value it currently has, thus, option B is correct.