

KillTest

質量更高 服務更好



練習題

<https://www.killtest.net>

一年免費更新服務

Exam : 300-445

Title : Designing and
Implementing Enterprise
Network Assurance

Version : DEMO

1.Exhibit:



An engineer works to optimize a website by reducing the page-load time to below 500 ms. The engineer set up a Cisco Thousand Eyes page-load test to baseline the current website performance.

Which action should be recommended to reduce page-load time?

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

Answer: C

Explanation:

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

When the goal is to reduce the page-load time from 1023 ms to below 500 ms, the engineer must target the heaviest components. Lazy loading is a design pattern that defers the initialization of non-critical resources at page load time. Instead of loading all images simultaneously when the user first navigates to the URL, lazy loading ensures that images are only downloaded as they are about to enter the viewport. This significantly reduces the initial DOM load time and the total Page Load Time because the browser does not have to wait for large image files to be fully retrieved before declaring the page "loaded."

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


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

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

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

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

2.Refer to Exhibit:



A network engineer is deploying a Cisco Thousand Eyes agent to monitor the network for a SaaS application without affecting the performance of the employee endpoints.

Which Thousand Eyes agent must be deployed to obtain the network metrics from branch A?

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

Answer: C

Explanation:

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

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

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

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

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

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

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

Introduction to Thousand Eyes

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

3.Refer to the exhibit.

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

Which integration type should be configured between Thousand Eyes and Grafana?

- A. opentelemetry
- B. custom-webhook
- C. push-api
- D. poll-api

Answer: A

Explanation:

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

According to the ENNA architecture guidelines, the Thousand Eyes Streaming API allows users to push granular test metrics (such as network latency, packet loss, and jitter) to external collectors in an OTel-compatible format. In the provided JSON payload, the "type" field is a mandatory parameter that defines the integration protocol. For Grafana, which natively supports Open Telemetry Protocol (OTLP) via its Open Telemetry Collector, the value must be set to "opentelemetry" (Option A). This tells the Thousand Eyes streaming engine to encapsulate the data according to the OTel semantic conventions, ensuring that Grafana can correctly interpret and visualize the metrics without additional custom parsing logic. While other options exist in the Thousand Eyes ecosystem, they do not fit the specific API call shown for this use case:

Custom Webhooks (Option B) are typically used for event-driven alerts and notifications (e.g., sending a POST request when a threshold is breached) rather than continuous high-fidelity metric streaming. Push-api and poll-api (Options C and D) are not valid "type" values within the context of the v7 /stream

endpoint, as the streaming service specifically utilizes the Open Telemetry framework for real-time delivery.

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

Introduction to Thousand Eyes for Open Telemetry

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

4.DRAG DROP

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

Answer Area

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

Answer:

Answer Area

Catalyst Center	AppDynamics
AppDynamics	Catalyst Center
ThousandEyes	ThousandEyes
Meraki	Meraki

Explanation:

AppDynamics
Catalyst Center
Thousand Eyes

Meraki

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

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


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

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

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

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

5.Refer to the exhibit.



An engineer must use Cisco Thousand Eyes testing to monitor their Cisco Catalyst SD-WAN fabric.

Which SD-WAN component is being monitored by Thousand Eyes?

A. underlay

- B. IPsec tunnels
- C. overlay
- D. GRE tunnels

Answer: A

Explanation:

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

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

By monitoring the underlay (Option A), the engineer can identify if high latency or packet loss is caused by a specific hop within the service provider's infrastructure or at a peering point. This "underlay visibility" is critical for troubleshooting SD-WAN performance issues where the overlay may report a tunnel down, but the root cause lies in a BGP routing change or physical fiber cut in the provider network. Thousand Eyes Enterprise Agents, natively integrated into Catalyst 8000 and ISR 4000 platforms, allow for this persistent underlay monitoring without additional hardware.

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

IPsec/GRE Tunnels (Options B & D): These represent the transport mechanisms of the overlay.

Thousand Eyes probes the path under these tunnels to ensure the transport health is sufficient to support the fabric.