

IP SLAs

- Feature History for IP SLA, on page 1
- Cisco IP SLAs, on page 2
- Are IP SLAs restricted by network topology?, on page 2
- Performance metrics for IP SLAs, on page 3
- SNMP and IP SLA, on page 3
- Network performance measurement with Cisco IP SLAs, on page 4
- IP SLA responder and IP SLA control protocol, on page 5
- Response time computation for IP SLAs, on page 6
- Monitoring and storing network performance statistics, on page 6
- Benefits of IP SLA, on page 7
- Configure the IP SLA responder, on page 8

Feature History for IP SLA

This table provides release and platform support information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature Name and Description	Supported Platform
Cisco IOS XE 26.2.1ea	IP SLAs: Cisco IP SLAs feature support has been introduced.	Cisco C9550 Series Smart Switches
Cisco IOS XE 17.18.1	IP SLAs: Cisco IP SLAs is a network enhancement feature that enables proactive network performance measurement by sending test data across the network to monitor performance between multiple locations or paths.	Cisco C9350 Series Smart Switches Cisco C9610 Series Smart Switches

Cisco IP SLAs

Cisco IP Service Level Agreements (IP SLAs) is a network enhancement feature that

- enables proactive network performance measurement
- by sending test data across the network
- to monitor performance between multiple locations or paths.

How does the IP SLA work

Summary

IP SLA operates by creating and sending test traffic across your network, mimicking real user data or application flows. This process helps you measure how well your network is performing in real time.

Workflow

1. IP SLA generates test packets that behave like actual network or application traffic (for example, voice, video, or web traffic).
2. These packets are sent between two network devices (such as routers or switches) or from a Cisco device to another IP-enabled device, like a server or another network appliance.
3. The receiving device (or application server) processes the test packets and may respond, depending on the type of IP SLA operation. The sending device measures how long the process takes, how many packets arrive, and other key performance metrics.
4. By analyzing this test traffic, IP SLA provides real-time data about network health, such as latency, jitter, packet loss, and connectivity.
5. These measurements help network administrators quickly detect, troubleshoot, and resolve network issues. The data is also valuable for analyzing network performance trends and for designing or optimizing network topologies.

Are IP SLAs restricted by network topology?

Cisco IP SLAs operate at the IP layer (Layer 3 of the OSI model), which means they do not rely on the underlying Layer 2 (data link layer) transport technologies such as Ethernet, Frame Relay, or MPLS. As a result, IP SLAs can be set up to send test traffic across any type of network infrastructure.

The benefits of this are:

- End-to-End monitoring

  You can measure network performance from one end of the network to another, regardless of how the devices are connected or what types of underlying connections are used.

- Works across diverse networks

  IP SLAs work over any combination of network types—wired, wireless, WAN, LAN, VPN, etc.

- User experience focus

  Because IP SLAs test traffic takes the same path as actual user traffic, the metrics collected (like delay, packet loss, or jitter) accurately reflect what end users are experiencing.

Performance metrics for IP SLAs

Cisco IP SLAs gather a variety of important network performance metrics. These metrics help network administrators understand how well the network is operating and how users experience network services. Here's what each metric means:

- Round-trip and one-way Delay

  Measures how long it takes for data to travel from the source to the destination and back (round-trip), or just one way. This helps determine if there are delays in the network that could affect applications.

- Directional jitter

  Measures the variation in delay between packets as they travel in one direction. High jitter can cause problems for real-time applications like voice and video.

- Directional packet loss

  Tracks how many packets are lost between the source and destination in one direction. Packet loss can result in poor application performance or dropped calls.

- Packet sequencing (order of arrival)

  Checks whether packets arrive in the same order they were sent. Out-of-order packets can disrupt certain applications, especially voice and video.

- Per-hop path information

  Provides details about each step (hop) a packet takes between source and destination. This can help identify where in the network problems are occurring.

- Directional connectivity

  Verifies if a path is up and reachable in a specific direction, helping to detect outages or unreachable segments.

- Server or website download times

  Measures how long it takes to download content from a server or website, simulating the user's experience when accessing online resources.

Collecting and analyzing these metrics allows network administrators to monitor, troubleshoot, and optimize network performance. By understanding the quality of the network from the perspective of the end user, they can ensure reliable and satisfactory service for critical applications and services.

SNMP and IP SLA

Cisco IP SLAs can send their performance data using a standard protocol called SNMP (Simple Network Management Protocol). Because SNMP is widely supported, this allows IP SLA measurements to be collected

and displayed by many network monitoring and management applications, such as Cisco Prime IPM and other third-party tools.

SNMP with IP SLA provides these benefits:

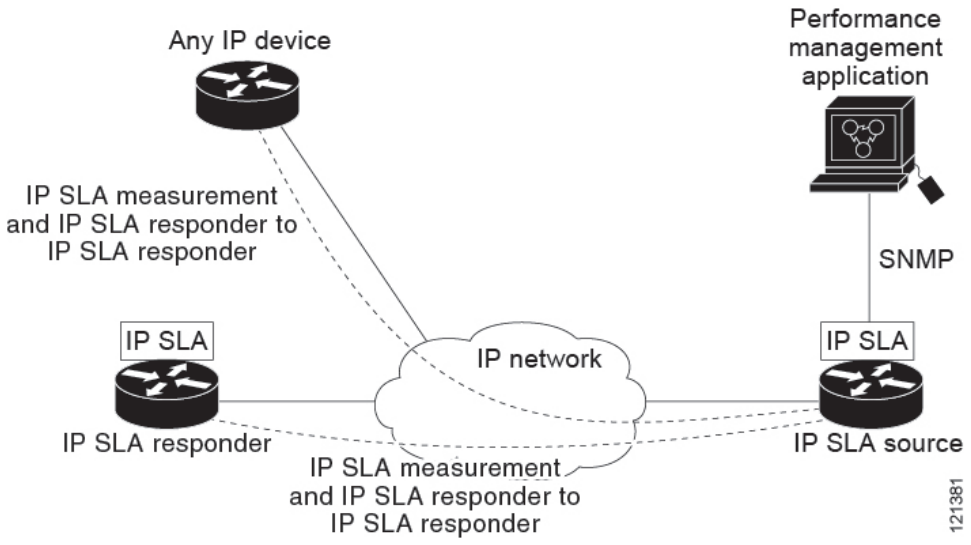
- Network administrators can see all the IP SLA data in one place, rather than checking each device individually. Efficient Troubleshooting: With all performance data available centrally, it's faster and easier to identify and resolve network problems.
- Management tools can analyze trends, create reports, and display performance metrics graphically, making it easier to spot issues and understand overall network health.
- With all performance data available centrally, it's faster and easier to identify and resolve network problems.

Network performance measurement with Cisco IP SLAs

Cisco IP SLAs help you monitor how well your network is performing across any part of your network—from the core (central routers and switches), to the distribution layer, and out to the network edge (branch offices or remote sites). Unlike traditional monitoring that might require extra hardware devices ("probes") placed throughout the network, IP SLAs are built into Cisco devices, so no additional equipment is needed.

The following figure shows how IP SLAs begin when the source device sends a generated packet to the destination device. After the destination device receives the packet, depending on the type of IP SLAs operation, it responds with time-stamp information for the source to make the calculation on performance metrics. An IP SLAs operation performs a network measurement from the source device to a destination in the network using a specific protocol such as UDP.

Figure 1: IP SLAs Operations



IP SLA responder and IP SLA control protocol

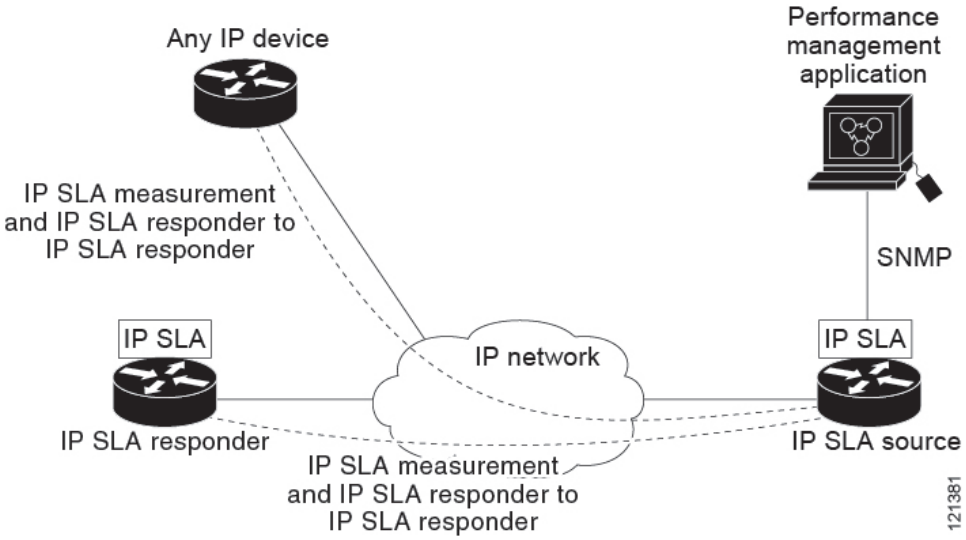
The IP SLA responder is a special software feature built into Cisco devices (such as routers or switches) that acts as the target for IP SLA test packets. When you run an IP SLA test, the responder is enabled on the destination device so it can recognize these specific test packets and respond to them.

The responder helps provide very accurate measurements of network performance, such as delay or packet loss, because it processes the test packets quickly and can add precise time-stamp information. Since the responder is built into Cisco devices, you don't need to deploy or buy any extra hardware (like dedicated probes).

The responder uses the Cisco IOS IP SLA Control Protocol, which tells the device exactly which port (a network communication endpoint) to listen on and for how long. This ensures the responder is only active when needed and on the correct port, improving both security and accuracy. The responder can be enabled on Cisco devices operating at Layer 2 (the data link layer), and it doesn't need to support all IP SLA features—just the ability to recognize and reply to test packets.

The following figure shows where the Cisco IOS IP SLA responder fits in the IP network. The Cisco IOS IP SLA responder operates by listening on a designated port for control protocol messages that are sent by an IP SLA operation. When the responder receives a control message, it temporarily enables the specified UDP or TCP port for a set period of time. During this active window, the responder accepts incoming requests and replies to them, facilitating precise network measurements. Once it has responded to the IP SLA packet or the configured duration ends, the responder disables the port to maintain security and efficiency. For enhanced security, MD5 authentication can be used for the control messages, ensuring only authorized operations are processed.

Figure 2: Cisco IOS IP SLAs Operation



It is not always necessary to enable the responder on the destination device for every IP SLA operation. If the IP SLA test is targeting services that are already running on the destination device, such as Telnet or HTTP, the responder feature is not needed. In these scenarios, the IP SLA operation can interact directly with the existing service, which simplifies the configuration process and eliminates the need for extra setup steps on the destination device. This makes it easier and faster to deploy IP SLA monitoring for commonly used network services.

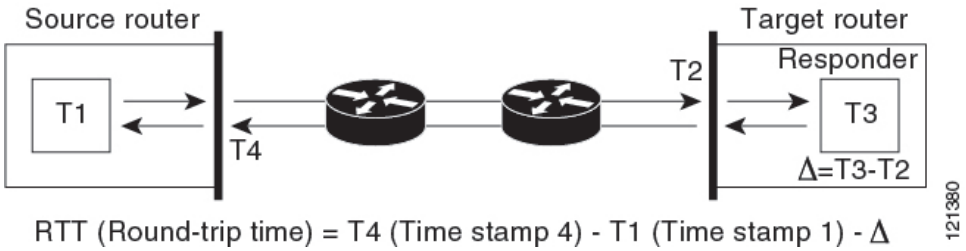
Response time computation for IP SLAs

When measuring network performance, it's important to get accurate response times. However, network devices like switches, controllers, and routers sometimes have to handle many tasks at once. This can cause short processing delays—so when a test packet arrives, it might wait in a queue before being processed or replied to. If these processing delays are included in the response time measurement, the results may not reflect the actual time it takes for data to travel across the network.

To overcome this, IP SLAs use precise time stamps to measure the exact moment a test packet enters and leaves a device. When the IP SLA responder feature is enabled, the device marks the time as soon as the packet arrives (at the interrupt level—before it's queued for processing) and again when it leaves. By subtracting out the time the packet spends being processed inside the device, IP SLAs provide a much more accurate measurement of the real network delay, not just how fast the device processes packets. This time stamping is done with very fine (sub-millisecond) accuracy, ensuring that even small delays are measured precisely.

The following figure demonstrates how the responder works. When the IP SLA responder is enabled, the round-trip time calculation becomes highly accurate by using four time stamps. As the test packet travels, the target router records the exact time it receives the packet (TS2) and the time it sends the response back (TS3). The difference between these two times, called delta, represents the processing time on the target device. This processing time is then subtracted from the total round-trip time to eliminate any delays caused by the device itself. Similarly, on the source router, the final arrival time of the response (TS4) is also captured at the interrupt level for maximum accuracy. By using this method, IP SLAs ensure that the measured round-trip time closely reflects only the actual network delay, not any internal device processing time.

Figure 3: Cisco IOS IP SLA Responder Time Stamping



Another important benefit of taking two time stamps at the target device is that it enables monitoring of advanced performance metrics like one-way delay, jitter, and directional packet loss. This is valuable because network traffic can often behave differently in each direction, so having detailed, directional statistics provides a more accurate picture of network health. To measure one-way delay accurately, both the source and target routers must have their clocks synchronized, typically using the Network Time Protocol (NTP). However, one-way jitter can still be measured even if the clocks are not synchronized, which allows administrators to assess variations in packet transit times without needing exact time alignment between devices. This flexibility makes it easier to analyze and troubleshoot network performance under real-world conditions.

Monitoring and storing network performance statistics

When you set up IP SLA operations on a Cisco device, the device continuously monitors various network performance statistics—such as delay, jitter, and packet loss—based on the type of test you configure. The statistics are saved directly on the Cisco device that is performing the IP SLA operation.

There are two options available for monitoring and storing network performance statistics. They are:

- You can use Cisco IOS commands to view the results directly from the device's console or terminal.
- The data is also available via SNMP, a standard network management protocol. This means you can use network monitoring tools to automatically collect, analyze, and display the performance statistics from multiple devices in a centralized system.

Customization of IP SLA packets

IP SLA packets can be customized with different IP and application layer options means that when you set up an IP SLA operation (a network test), you can adjust several parameters of the test packets. This customization helps you make the test simulate real network conditions or target specific parts of your network.

The options you can customize include:

- Source and Destination IP Addresses

  You can choose which device (IP address) sends the test packets and which device receives them.

- UDP or TCP Port Numbers

  You can specify which application port to use (for example, port 80 for HTTP or port 5060 for VoIP/SIP) so you can test the path for specific applications or services.

- Type of Service (ToS) Byte Settings

  This includes settings like DSCP and IP Precedence, which are used for Quality of Service (QoS). This lets you see how high-priority or low-priority traffic performs across your network.

- VRF Instances

  If your network uses VPNs or multiple routing tables, you can specify which VPN or VRF context to use for the test, ensuring you're measuring performance for the correct network segment.

- URL Web Addresses

  You can target a specific website or web application to test if it's reachable and how it performs from your network's perspective.

Network administrators often need to monitor different traffic types and paths. With these customization options, they can:

- Test exactly the traffic their users or applications use.
- Simulate real-life network scenarios for accurate measurement.
- Verify the performance of specific business-critical services or applications.
- Ensure the network meets required service levels for different traffic types or customers.

Benefits of IP SLA

Using IP SLAs offers several important benefits for network management and monitoring:

- Service-Level Agreement (SLA) monitoring and verification

IP SLAs help you track and verify whether your network is meeting the performance standards promised in SLAs with service providers or internal customers. This means you can prove the network is delivering the agreed level of service.

- Comprehensive network performance monitoring

  By measuring critical metrics like jitter, latency, and packet loss, IP SLAs provide a detailed and ongoing view of network health. These reliable and predictable measurements make it easier to assess network performance over time.

- Quality of Service (QoS) assessment

  IP SLAs allow you to check if your current network configuration can support new IP-based services (like voice or video), and whether your QoS settings are effective. This is important before rolling out new applications that require certain performance levels.

- End-to-End Network Availability

  You can use IP SLAs to test network connections from one end to the other (edge to edge), including remote sites. This is useful for confirming that important resources, like servers storing business data, are reachable at all times.

- Efficient Troubleshooting

  IP SLAs provide consistent and accurate performance data, helping you quickly pinpoint and fix network issues. This saves time and reduces the impact of problems on users and business operations.

- Support for MPLS Networks

  On devices that use MPLS, IP SLAs can also measure and verify the performance of MPLS paths, ensuring that these advanced network segments meet business requirements.

Configure the IP SLA responder

The IP SLA responder is available only on Cisco IOS software-based devices, including some Layer 2 devices that do not support full IP SLA functionality.

Perform this task to configure the IP SLA responder on the target device (the operational target)

Procedure

Step 1 **enable**

Example:

```
Device> enable
```

Enables privileged EXEC mode.

Enter your password, if prompted.

Step 2 **configure terminal**

Example:

```
Device# configure terminal
```

Enters global configuration mode.

Step 3 **ip sla responder** {**tcp-connect** | **udp-echo**} **ipaddress** *ip-address* **port** *port-number*

Example:

```
Device(config)# ip sla responder udp-echo 172.29.139.134 5000
```

Configures the device as an IP SLA responder.

- **tcp-connect**: Enables the responder for TCP connect operations.
- **udp-echo**: Enables the responder for User DatagramProtocol (UDP) echo or jitter operations.
- **ipaddress** *ip-address*: Enter the destination IP address.
- **port** *port-number*: Enter the destination port number.

  Note
  The IP address and port number must match those configured on the source device for the IP SLA operation.

Step 4 **end**

Example:

```
Device(config)# end
```

Exits to privileged EXEC mode.