



CHAPTER 39

Configuring Cisco IOS IP SLAs Operations

This chapter describes how to use Cisco IOS IP Service Level Agreements (SLAs) and the IETF Two-Way Active Measurement Protocol (TWAMP) on the Cisco ME 3400 Ethernet Access switch. Cisco IP SLAs is a part of Cisco IOS software that allows Cisco customers to analyze IP service levels for IP applications and services by using active traffic monitoring—the generation of traffic in a continuous, reliable, and predictable manner—for measuring network performance. With Cisco IOS IP SLAs, service provider customers can measure and provide service level agreements, and enterprise customers can verify service levels, verify outsourced service level agreements, and understand network performance. Cisco IOS IP SLAs can perform network assessments, verify quality of service (QoS), ease the deployment of new services, and assist with network troubleshooting.



Note

Full Cisco IP SLAs functionality is supported only on switches that are running the metro IP access or metro access image. Switches running the metro base image, such as the ME 2400, support only IP SLAs responder functionality and must be configured with another device that supports full IP SLAs functionality.

TWAMP defines a standard for measuring round-trip network performance between any two devices that support the protocol. Beginning with Cisco IOS Release 12.2(52)SE, you can implement TWAMP on the Cisco ME 3400 Ethernet Access switch.

For more information about IP SLAs, see the *Cisco IOS IP SLAs Configuration Guide, Release 12.4T* at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

For command syntax information, see the command reference at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_book.html

This chapter consists of these sections:

- [Understanding Cisco IOS IP SLAs, page 39-2](#)
- [Configuring IP SLAs Operations, page 39-6](#)
- [Monitoring IP SLAs Operations, page 39-13](#)
- [Understanding TWAMP, page 39-14](#)
- [Configuring TWAMP, page 39-15](#)

Understanding Cisco IOS IP SLAs

Cisco IOS IP SLAs sends data across the network to measure performance between multiple network locations or across multiple network paths. It simulates network data and IP services and collects network performance information in real time. Cisco IOS IP SLAs generates and analyzes traffic either between Cisco IOS devices or from a Cisco IOS device to a remote IP device such as a network application server. Measurements provided by the various Cisco IOS IP SLAs operations can be used for troubleshooting, for problem analysis, and for designing network topologies.

Depending on the specific Cisco IOS IP SLAs operation, various network performance statistics are monitored within the Cisco device and stored in both command-line interface (CLI) and Simple Network Management Protocol (SNMP) MIBs. IP SLAs packets have configurable IP and application layer options such as source and destination IP address, User Datagram Protocol (UDP)/TCP port numbers, a type of service (ToS) byte (including Differentiated Services Code Point [DSCP] and IP Prefix bits), Virtual Private Network (VPN) routing/forwarding instance (VRF), and URL web address.

Because Cisco IP SLAs is Layer 2 transport independent, you can configure end-to-end operations over disparate networks to best reflect the metrics that an end user is likely to experience. IP SLAs collects a unique subset of these performance metrics:

- Delay (both round-trip and one-way)
- Jitter (directional)
- Packet loss (directional)
- Packet sequencing (packet ordering)
- Path (per hop)
- Connectivity (directional)
- Server or website download time

Because Cisco IOS IP SLAs is SNMP-accessible, it can also be used by performance-monitoring applications like CiscoWorks Internetwork Performance Monitor (IPM) and other third-party Cisco partner performance management products. You can find more details about network management products that use Cisco IOS IP SLAs at this URL:

<http://www.cisco.com/go/ipsla>

Using IP SLAs can provide these benefits:

- Service-level agreement monitoring, measurement, and verification.
- Network performance monitoring
 - Measures the jitter, latency, or packet loss in the network.
 - Provides continuous, reliable, and predictable measurements.
- IP service network health assessment to verify that the existing QoS is sufficient for new IP services.
- Edge-to-edge network availability monitoring for proactive verification and connectivity testing of network resources (for example, shows the network availability of an NFS server used to store business critical data from a remote site).
- Troubleshooting of network operation by providing consistent, reliable measurement that immediately identifies problems and saves troubleshooting time.
- Multiprotocol Label Switching (MPLS) performance monitoring and network verification (if the switch supports MPLS)


This section includes this information about IP SLAs functionality:

- [Using Cisco IOS IP SLAs to Measure Network Performance](#), page 39-3
- [IP SLAs Responder and IP SLAs Control Protocol](#), page 39-4
- [Response Time Computation for IP SLAs](#), page 39-4
- [IP SLAs Operation Scheduling](#), page 39-5
- [IP SLAs Operation Threshold Monitoring](#), page 39-5

Using Cisco IOS IP SLAs to Measure Network Performance

You can use IP SLAs to monitor the performance between any area in the network—core, distribution, and edge—without deploying a physical probe. It uses generated traffic to measure network performance between two networking devices. [Figure 39-1](#) shows how IP SLAs begins 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 39-1 Cisco IOS IP SLAs Operation



To implement IP SLAs network performance measurement, you need to perform these tasks:

1. Enable the IP SLAs responder, if required.
2. Configure the required IP SLAs operation type.
3. Configure any options available for the specified operation type.
4. Configure threshold conditions, if required.
5. Schedule the operation to run, then let the operation run for a period of time to gather statistics.
6. Display and interpret the results of the operation using the Cisco IOS CLI or a network management system (NMS) system with SNMP.

For more information about IP SLAs operations, see the operation-specific chapters in the *Cisco IOS IP SLAs Configuration Guide* at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

**Note**

The switch does not support IP SLAs Voice over IP (VoIP) service levels using the gatekeeper registration delay operations measurements. Before configuring any IP SLAs application, you can use the **show ip sla application** privileged EXEC command to verify that the operation type is supported on your software image.

IP SLAs Responder and IP SLAs Control Protocol

The IP SLAs responder is a component embedded in the destination Cisco device that allows the system to anticipate and respond to IP SLAs request packets. The responder provides accurate measurements without the need for dedicated probes. The responder uses the Cisco IOS IP SLAs Control Protocol to provide a mechanism through which it can be notified on which port it should listen and respond. Only a Cisco IOS device can be a source for a destination IP SLAs Responder.

**Note**

The IP SLAs responder can be a Cisco IOS Layer 2, responder-configurable switch, such as a Catalyst 2960 or Cisco ME 2400 switch or a Cisco ME 3400 switch running the metro base image. The responder does not need to support full IP SLAs functionality.

Figure 39-1 shows where the Cisco IOS IP SLAs responder fits in the IP network. The responder listens on a specific port for control protocol messages sent by an IP SLAs operation. Upon receipt of the control message, it enables the specified UDP or TCP port for the specified duration. During this time, the responder accepts the requests and responds to them. It disables the port after it responds to the IP SLAs packet, or when the specified time expires. MD5 authentication for control messages is available for added security.

You do not need to enable the responder on the destination device for all IP SLAs operations. For example, a responder is not required for services that are already provided by the destination router (such as Telnet or HTTP). You cannot configure the IP SLAs responder on non-Cisco devices and Cisco IOS IP SLAs can send operational packets only to services native to those devices.


Response Time Computation for IP SLAs

Switches and routers can take tens of milliseconds to process incoming packets due to other high priority processes. This delay affects the response times because the test-packet reply might be in a queue while waiting to be processed. In this situation, the response times would not accurately represent true network delays. IP SLAs minimizes these processing delays on the source device as well as on the target device (if the responder is being used) to determine true round-trip times. IP SLAs test packets use time stamping to minimize the processing delays.

When the IP SLAs responder is enabled, it allows the target device to take time stamps when the packet arrives on the interface at interrupt level and again just as it is leaving, eliminating the processing time. This time stamping is made with a granularity of sub-milliseconds (ms).

Figure 39-2 demonstrates how the responder works. Four time stamps are taken to make the calculation for round-trip time. At the target router, with the responder functionality enabled, time stamp 2 (TS2) is subtracted from time stamp 3 (TS3) to produce the time spent processing the test packet as represented by delta. This delta value is then subtracted from the overall round-trip time. Notice that the same principle is applied by IP SLAs on the source router where the incoming time stamp 4 (TS4) is also taken at the interrupt level to allow for greater accuracy.

Figure 39-2 Cisco IOS IP SLAs Responder Time Stamping



An additional benefit of the two time stamps at the target device is the ability to track one-way delay, jitter, and directional packet loss. Because much network behavior is asynchronous, it is critical to have these statistics. However, to capture one-way delay measurements, you must configure both the source router and target router with Network Time Protocol (NTP) so that the source and target are synchronized to the same clock source. One-way jitter measurements do not require clock synchronization.

IP SLAs Operation Scheduling

When you configure an IP SLAs operation, you must schedule the operation to begin capturing statistics and collecting error information. You can schedule an operation to start immediately or to start at a certain month, day, and hour. You can use the *pending* option to set the operation to start at a later time. The pending option is an internal state of the operation that is visible through SNMP. The pending state is also used when an operation is a reaction (threshold) operation waiting to be triggered. You can schedule a single IP SLAs operation or a group of operations at one time.

You can schedule several IP SLAs operations by using a single command through the Cisco IOS CLI or the CISCO RTTMON-MIB. Scheduling the operations to run at evenly distributed times allows you to control the amount of IP SLAs monitoring traffic. This distribution of IP SLAs operations helps minimize the CPU utilization and thus improves network scalability.

For more details about the IP SLAs multioperations scheduling functionality, see the “IP SLAs—Multiple Operation Scheduling” chapter of the *Cisco IOS IP SLAs Configuration Guide* at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

IP SLAs Operation Threshold Monitoring

To support successful service level agreement monitoring, you must have mechanisms that notify you immediately of any possible violation. IP SLAs can send SNMP traps that are triggered by events such as these:

- Connection loss
- Timeout

- Round-trip time threshold
- Average jitter threshold
- One-way packet loss
- One-way jitter
- One-way mean opinion score (MOS)
- One-way latency

An IP SLAs threshold violation can also trigger another IP SLAs operation for further analysis. For example, the frequency could be increased or an ICMP path echo or ICMP path jitter operation could be initiated for troubleshooting.

Determining the type of threshold and the level to set can be complex, and depends on the type of IP service being used in the network. For more details on using thresholds with Cisco IOS IP SLAs operations, see the “IP SLAs—Proactive Threshold Monitoring” chapter of the Cisco IOS IP SLAs Configuration Guide at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

Configuring IP SLAs Operations

This section does not include configuration information for all available operations as the configuration information details are included in the *Cisco IOS IP SLAs Configuration Guide*. It does include several operations as examples, including configuring the responder, configuring UDP jitter operation, which requires a responder, and configuring ICMP echo operation, which does not require a responder. For details about configuring other operations, see the *Cisco IOS IP SLAs Configuration Guide* at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

This section includes this information:

- [Default Configuration, page 39-6](#)
- [Configuration Guidelines, page 39-6](#)
- [Configuring the IP SLAs Responder, page 39-8](#)
- [Analyzing IP Service Levels by Using the UDP Jitter Operation, page 39-8](#)
- [Analyzing IP Service Levels by Using the ICMP Echo Operation, page 39-11](#)

Default Configuration

No IP SLAs operations are configured.

Configuration Guidelines

For information on the IP SLAs commands, see the *Cisco IOS IP SLAs Command Reference, Release 12.4T* command reference at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_book.html

For detailed descriptions and configuration procedures, see the *Cisco IOS IP SLAs Configuration Guide, Release 12.4T* at this URL:

http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

Note that not all of the IP SLAs commands or operations described in this guide are supported on the switch. The switch supports IP service level analysis by using UDP jitter, UDP echo, HTTP, TCP connect, ICMP echo, ICMP path echo, ICMP path jitter, FTP, DNS, and DHCP, as well as multiple operation scheduling and proactive threshold monitoring. It does not support IP SLAs VoIP service levels using the gatekeeper registration delay operations measurements.

Before configuring any IP SLAs application, you can use the **show ip sla application** privileged EXEC command to verify that the operation type is supported on your software image. This is an example of the output from the command:

```
Switch# show ip sla application
      IP SLAs
Version: 2.2.0 Round Trip Time MIB, Infrastructure Engine-II
Time of last change in whole IP SLAs: 22:17:39.117 UTC Fri Jun
Estimated system max number of entries: 15801

Estimated number of configurable operations: 15801
Number of Entries configured      : 0
Number of active Entries          : 0
Number of pending Entries         : 0
Number of inactive Entries        : 0

      Supported Operation Types
Type of Operation to Perform: 802.1agEcho
Type of Operation to Perform: 802.1agJitter
Type of Operation to Perform: dhcp
Type of Operation to Perform: dns
Type of Operation to Perform: echo
Type of Operation to Perform: ftp
Type of Operation to Perform: http
Type of Operation to Perform: jitter
Type of Operation to Perform: pathEcho
Type of Operation to Perform: pathJitter
Type of Operation to Perform: tcpConnect
Type of Operation to Perform: udpEcho

IP SLAs low memory water mark: 21741224
```

Configuring the IP SLAs Responder

The IP SLAs responder is available only on Cisco IOS software-based devices, including some Layer 2 switches that do not support full IP SLAs functionality, such as the Catalyst 2960 or the Cisco ME 2400 switch or a Cisco ME 3400 switch running the metro base image. Beginning in privileged EXEC mode, follow these steps to configure the IP SLAs responder on the target device (the operational target):

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>configure terminal</code> | Enter global configuration mode. |
| Step 2 | <code>ip sla responder {tcp-connect udp-echo} ipaddress ip-address port port-number</code> | Configure the switch as an IP SLAs responder. The keywords have these meanings: <ul style="list-style-type: none"> • tcp-connect—Enable the responder for TCP connect operations. • udp-echo—Enable the responder for User Datagram Protocol (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 SLAs operation. |
| Step 3 | <code>end</code> | Return to privileged EXEC mode. |
| Step 4 | <code>show ip sla responder</code> | Verify the IP SLAs responder configuration on the device. |
| Step 5 | <code>copy running-config startup-config</code> | (Optional) Save your entries in the configuration file. |

To disable the IP SLAs responder, enter the **no ip sla responder** global configuration command. This example shows how to configure the device as a responder for the UDP jitter IP SLAs operation in the next procedure:

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

Analyzing IP Service Levels by Using the UDP Jitter Operation

Jitter means interpacket delay variance. When multiple packets are sent consecutively 10 ms apart from source to destination, if the network is behaving correctly, the destination should receive them 10 ms apart. But if there are delays in the network (like queuing, arriving through alternate routes, and so on) the arrival delay between packets might be more than or less than 10 ms with a positive jitter value meaning that the packets arrived more than 10 ms apart. If the packets arrive 12 ms apart, positive jitter is 2 ms; if the packets arrive 8 ms apart, negative jitter is 2 ms. For delay-sensitive networks, positive jitter values are undesirable, and a jitter value of 0 is ideal.

In addition to monitoring jitter, the IP SLAs UDP jitter operation can be used as a multipurpose data gathering operation. The packets IP SLAs generates carry packet sending and receiving sequence information and sending and receiving time stamps from the source and the operational target. Based on these, UDP jitter operations measure this data:

- Per-direction jitter (source to destination and destination to source)
- Per-direction packet-loss

- Per-direction delay (one-way delay)
- Round-trip delay (average round-trip time)

Because the paths for the sending and receiving of data can be different (asymmetric), you can use the per-direction data to more readily identify where congestion or other problems are occurring in the network.

The UDP jitter operation generates synthetic (simulated) UDP traffic and sends a number of UDP packets, each of a specified size, sent a specified number of milliseconds apart, from a source router to a target router, at a given frequency. By default, ten packet-frames, each with a payload size of 10 bytes are generated every 10 ms, and the operation is repeated every 60 seconds. You can configure each of these parameters to best simulate the IP service you want to provide.

To provide accurate one-way delay (latency) measurements, time synchronization, such as that provided by NTP, is required between the source and the target device. Time synchronization is not required for the one-way jitter and packet loss measurements. If the time is not synchronized between the source and target devices, one-way jitter and packet loss data is returned, but values of 0 are returned for the one-way delay measurements provided by the UDP jitter operation

**Note**

Before you configure a UDP jitter operation on the source device, you must enable the IP SLAs responder on the target device (the operational target).

Beginning in privileged EXEC mode, follow these steps to configure UDP jitter operation on the source device:

| | Command | Purpose |
|---------------|--------------------------------|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | ip sla operation-number | Create an IP SLAs operation, and enter IP SLAs configuration mode. |

| | Command | Purpose |
|--------|---|---|
| Step 3 | udp-jitter { <i>destination-ip-address</i> <i>destination-hostname</i> } <i>destination-port</i> [source-ip { <i>ip-address</i> <i>hostname</i> }] [source-port <i>port-number</i>] [control { enable disable }] [num-packets <i>number-of-packets</i>] [interval <i>interpacket-interval</i>] | Configure the IP SLAs operation as a UDP jitter operation, and enter UDP jitter configuration mode. <ul style="list-style-type: none"> • <i>destination-ip-address</i> <i>destination-hostname</i>—Specify the destination IP address or hostname. • <i>destination-port</i>—Specify the destination port number in the range from 1 to 65535. • (Optional) source-ip {<i>ip-address</i> <i>hostname</i>}—Specify the source IP address or hostname. When a source IP address or hostname is not specified, IP SLAs chooses the IP address nearest to the destination • (Optional) source-port <i>port-number</i>—Specify the source port number in the range from 1 to 65535. When a port number is not specified, IP SLAs chooses an available port. • (Optional) control—Enable or disable sending of IP SLAs control messages to the IP SLAs responder. By default, IP SLAs control messages are sent to the destination device to establish a connection with the IP SLAs responder • (Optional) num-packets <i>number-of-packets</i>—Enter the number of packets to be generated. The range is 1 to 6000; the default is 10. • (Optional) interval <i>inter-packet-interval</i>—Enter the interval between sending packets in milliseconds. The range is 1 to 6000; the default value is 20 ms. |
| Step 4 | frequency <i>seconds</i> | (Optional) Set the rate at which a specified IP SLAs operation repeats. The range is from 1 to 604800 seconds; the default is 60 seconds. |
| Step 5 | exit | Exit UDP jitter configuration mode, and return to global configuration mode. |
| Step 6 | ip sla monitor schedule <i>operation-number</i> [life { forever <i>seconds</i> }] [start-time { <i>hh:mm[:ss]</i> <i>month day</i> <i>day month</i> } pending now after <i>hh:mm:ss</i>] [ageout <i>seconds</i>] [recurring] | Configure the scheduling parameters for an individual IP SLAs operation. <ul style="list-style-type: none"> • <i>operation-number</i>—Enter the RTR entry number. • (Optional) life—Set the operation to run indefinitely (forever) or for a specific number of <i>seconds</i>. The range is from 0 to 2147483647. The default is 3600 seconds (1 hour). • (Optional) start-time—Enter the time for the operation to begin collecting information: <ul style="list-style-type: none"> – To start at a specific time, enter the hour, minute, second (in 24-hour notation), and day of the month. If no month is entered, the default is the current month. – Enter pending to select no information collection until a start time is selected. – Enter now to start the operation immediately. – Enter after <i>hh:mm:ss</i> to show that the operation should start after the entered time has elapsed. • (Optional) ageout <i>seconds</i>—Enter the number of seconds to keep the operation in memory when it is not actively collecting information. The range is 0 to 2073600 seconds, the default is 0 seconds (never ages out). • (Optional) recurring—Set the operation to automatically run every day. |

| | Command | Purpose |
|--------|---|--|
| Step 7 | end | Return to privileged EXEC mode. |
| Step 8 | show ip sla configuration [<i>operation-number</i>] | (Optional) Display configuration values, including all defaults for all IP SLAs operations or a specified operation. |
| Step 9 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To disable the IP SLAs operation, enter the `no ip sla operation-number` global configuration command. This example shows how to configure a UDP jitter IP SLAs operation:

```
Switch(config)# ip sla 10
Switch(config-ip-sla)# udp-jitter 172.29.139.134 5000
Switch(config-ip-sla-jitter)# frequency 30
Switch(config-ip-sla-jitter)# exit
Switch(config)# ip sla schedule 5 start-time now life forever
Switch(config)# end
Switch# show ip sla configuration 10
IP SLAs, Infrastructure Engine-II.

Entry number: 10
Owner:
Tag:
Type of operation to perform: udp-jitter
Target address/Source address: 1.1.1.1/0.0.0.0
Target port/Source port: 2/0
Request size (ARR data portion): 32
Operation timeout (milliseconds): 5000
Packet Interval (milliseconds)/Number of packets: 20/10
Type Of Service parameters: 0x0
Verify data: No
Vrf Name:
Control Packets: enabled
Schedule:
  Operation frequency (seconds): 30
  Next Scheduled Start Time: Pending trigger
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): 3600
  Entry Ageout (seconds): never
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): notInService
Threshold (milliseconds): 5000
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
Enhanced History:
```

Analyzing IP Service Levels by Using the ICMP Echo Operation

The ICMP echo operation measures end-to-end response time between a Cisco device and any devices using IP. Response time is computed by measuring the time taken between sending an ICMP echo request message to the destination and receiving an ICMP echo reply. Many customers use IP SLAs ICMP-based operations, in-house ping testing, or ping-based dedicated probes for response time measurements between the source IP SLAs device and the destination IP device. The IP SLAs ICMP echo operation conforms to the same specifications as ICMP ping testing, and the two methods result in the same response times.

**Note**

This operation does not require the IP SLAs responder to be enabled.

Beginning in privileged EXEC mode, follow these steps to configure an ICMP echo operation on the source device:

| | Command | Purpose |
|--------|--|---|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | ip sla operation-number | Create an IP SLAs operation and enter IP SLAs configuration mode. |
| Step 3 | icmp-echo { <i>destination-ip-address</i> <i>destination-hostname</i> } [source-ip { <i>ip-address</i> <i>hostname</i> } source-interface <i>interface-id</i>] | Configure the IP SLAs operation as an ICMP Echo operation and enter ICMP echo configuration mode. <ul style="list-style-type: none"> • <i>destination-ip-address</i> <i>destination-hostname</i>—Specify the destination IP address or hostname. • (Optional) source-ip { <i>ip-address</i> <i>hostname</i> }—Specify the source IP address or hostname. When a source IP address or hostname is not specified, IP SLAs chooses the IP address nearest to the destination • (Optional) source-interface <i>interface-id</i>—Specify the source interface for the operation. |
| Step 4 | frequency <i>seconds</i> | (Optional) Set the rate at which a specified IP SLAs operation repeats. The range is from 1 to 604800 seconds; the default is 60 seconds. |
| Step 5 | exit | Exit UDP jitter configuration mode, and return to global configuration mode. |
| Step 6 | ip sla schedule <i>operation-number</i> [life { forever <i>seconds</i> }] [start-time { <i>hh:mm</i> [: <i>ss</i>] [<i>month</i> <i>day</i> <i>day</i> <i>month</i>] } pending now after <i>hh:mm:ss</i>] [ageout <i>seconds</i>] [recurring] | Configure the scheduling parameters for an individual IP SLAs operation. <ul style="list-style-type: none"> • <i>operation-number</i>—Enter the RTR entry number. • (Optional) life—Set the operation to run indefinitely (forever) or for a specific number of <i>seconds</i>. The range is from 0 to 2147483647. The default is 3600 seconds (1 hour) • (Optional) start-time—Enter the time for the operation to begin collecting information: <ul style="list-style-type: none"> – To start at a specific time, enter the hour, minute, second (in 24-hour notation), and day of the month. If no month is entered, the default is the current month. – Enter pending to select no information collection until a start time is selected. – Enter now to start the operation immediately. – Enter after <i>hh:mm:ss</i> to indicate that the operation should start after the entered time has elapsed. • (Optional) ageout <i>seconds</i>—Enter the number of seconds to keep the operation in memory when it is not actively collecting information. The range is 0 to 2073600 seconds; the default is 0 seconds (never ages out). • (Optional) recurring—Set the operation to automatically run every day. |
| Step 7 | end | Return to privileged EXEC mode. |

| | Command | Purpose |
|--------|---|---|
| Step 8 | show ip sla configuration [<i>operation-number</i>] | (Optional) Display configuration values including all defaults for all IP SLAs operations or a specified operation. |
| Step 9 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To disable the IP SLAs operation, enter the **no ip sla *operation-number*** global configuration command. This example shows how to configure an ICMP echo IP SLAs operation:

```
Switch(config)# ip sla 12
Switch(config-ip-sla)# icmp-echo 172.29.139.134
Switch(config-ip-sla-echo)# frequency 30
Switch(config-ip-sla-echo)# exit
Switch(config)# ip sla schedule 5 start-time now life forever
Switch(config)# end
Switch# show ip sla configuration 22
IP SLAs, Infrastructure Engine-1L.

Entry number: 12
Owner:
Tag:
Type of operation to perform: echo
Target address: 2.2.2.2
Source address: 0.0.0.0
Request size (ARR data portion): 28
Operation timeout (milliseconds): 5000
Type Of Service parameters: 0x0
Verify data: No
Vrf Name:
Schedule:
  Operation frequency (seconds): 60
  Next Scheduled Start Time: Pending trigger
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): 3600
  Entry Ageout (seconds): never
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): notInService
Threshold (milliseconds): 5000
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
History Statistics:
  Number of history Lives kept: 0
  Number of history Buckets kept: 15
  History Filter Type: None
Enhanced History:
```

Monitoring IP SLAs Operations

Use the User EXEC or Privileged EXEC commands in [Table 39-1](#) to display IP SLAs operations configuration and results.

Table 39-1 Monitoring IP SLAs Operations


| Command | Purpose |
|---|--|
| show ip sla application | Display global information about Cisco IOS IP SLAs. |
| show ip sla authentication | Display IP SLAs authentication information. |
| show ip sla configuration [entry-number] | Display configuration values including all defaults for all IP SLAs operations or a specific operation. |
| show ip sla enhanced-history {collection-statistics distribution statistics} [entry-number] | Display enhanced history statistics for collected history buckets or distribution statistics for all IP SLAs operations or a specific operation. |
| show ip sla ethernet-monitor configuration [entry-number] | Display IP SLAs automatic Ethernet configuration. |
| show ip sla group schedule [schedule-entry-number] | Display IP SLAs group scheduling configuration and details. |
| show ip sla history [entry-number full tabular] | Display history collected for all IP SLAs operations |
| show ip sla mpls-lsp-monitor {collection-statistics configuration ldp operational-state scan-queue summary [entry-number] neighbors} | Display MPLS label switched path (LSP) Health Monitor operations, |
| show ip sla reaction-configuration [entry-number] | Display the configured proactive threshold monitoring settings for all IP SLAs operations or a specific operation. |
| show ip sla reaction-trigger [entry-number] | Display the reaction trigger information for all IP SLAs operations or a specific operation. |
| show ip sla responder | Display information about the IP SLAs responder. |
| show ip sla statistics [entry-number aggregated details] | Display current or aggregated operational status and statistics. |

Understanding TWAMP

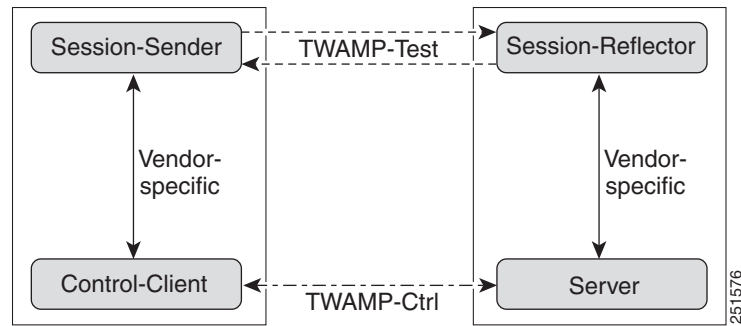
TWAMP consists of two related protocols. Use the TWAMP-Control protocol to start performance measurement sessions. Use TWAMP-Test to send and receive performance-measurement probes. You can deploy TWAMP in a simplified network architecture, with the control-client and the session-sender on one device and the server and the session-reflector on another device.

The Cisco IOS software TWAMP implementation supports a basic configuration. [Figure 39-3](#) shows a sample deployment.

[Figure 39-4](#) shows the four logical entities that comprise TWAMP.

Figure 39-3 TWAMP Deployment

251575

Figure 39-4 TWAMP Architecture

Although each entity is separate, the protocol allows for logical merging of the roles on a single device.

Configuring TWAMP

- [Configuring the TWAMP Server, page 39-16](#)
- [Configuring the TWAMP Reflector, page 39-16](#)
- [Troubleshooting TWAMP, page 39-17](#)

Configuring the TWAMP Server

The TWAMP server and reflector functionality are configured on the same device.


Note

The switch does not support the TWAMP sender and client roles.

Beginning in privileged EXEC mode, follow these steps to configure the TWAMP server:

| | Command | Purpose |
|---------|--|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | ip sla server twamp | Configure the switch as a TWAMP server, and enter TWAMP configuration mode. |
| Step 3 | port <i>port-number</i> | (Optional) Specify the port to be used by the TWAMP server to listen for connection and control requests. The same port negotiates for the port to which performance probes are sent. The configured port should not be an IANA well-known port or any port used by other applications. The default is port 862. |
| Step 4 | timer inactivity <i>seconds</i> | (Optional) Set the maximum time, in seconds, the session can be inactive before the session ends. The range is 1–6000 seconds. The default is 900 seconds. |
| Step 5 | end | Return to privileged EXEC mode. |
| Step 6s | show ip sla standards | (Optional) Display the IP SLA standards supported on the switch. |
| Step 7 | show ip sla twamp connection requests | (Optional) Display the number and the source of TWAMP connections. |
| Step 8 | show ip sla twamp connection detail | (Optional) Display the connection ID, client IP address and port number, mode and status of TWAMP connections, and the number of test requests. |
| Step 9 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To disable the IP SLA TWAMP server, enter the **no ip sla server twamp** global configuration command. This example shows how to configure a switch as an IP SLA TWAMP server:

```
Switch(config)# ip sla server twamp
Switch(config-twamp-srvr)# port 9000
Switch(config-twamp-srvr)# timer inactivity 300
```

Configuring the TWAMP Reflector

The TWAMP server and reflector functionality are both configured on the same device.

Beginning in privileged EXEC mode, follow these steps to configure the TWAMP reflector:

| | Command | Purpose |
|--------|-------------------------------|--|
| Step 1 | configure terminal | Enter global configuration mode. |
| Step 2 | ip sla responder twamp | Configure the switch as a TWAMP responder, and enter TWAMP configuration mode. |

| | Command | Purpose |
|--------|---|--|
| Step 3 | timeout <i>seconds</i> | (Optional) Set the maximum time, in seconds, the session can be inactive before the session ends. The range is 1–604800 seconds. The default is 900 seconds. |
| Step 4 | end | Return to privileged EXEC mode. |
| Step 5 | show ip sla twamp session [source-ip <i>ip-address</i> source-port <i>port-number</i>] | (Optional) Display information about TWAMP test results for the specified client. |
| Step 6 | copy running-config startup-config | (Optional) Save your entries in the configuration file. |

To disable the IP SLA TWAMP reflector, enter the **no ip sla responder twamp** global configuration command. This example shows how to configure a switch as an IP SLA TWAMP reflector:

```
Switch(config)# ip sla responder twamp
Switch(config-twamp-srvr)# timeout 300
```

Troubleshooting TWAMP

Use these commands to troubleshoot TWAMP sessions:

```
debug ip sla error twamp [connection source-ip ip-address | control {reflector | server} | session source-ip ip-address]
```

```
debug ip sla trace twamp [connection source-ip ip-address | control {reflector | server} | session source-ip ip-address]
```

