



IP SLAs Overview

This module describes IP Service Level Agreements (SLAs). IP SLAs allows Cisco customers to analyze IP service levels for IP applications and services, to increase productivity, to lower operational costs, and to reduce the frequency of network outages. IP SLAs uses active traffic monitoring--the generation of traffic in a continuous, reliable, and predictable manner--for measuring network performance. Using 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. IP SLAs can perform network assessments, verify quality of service (QoS), ease the deployment of new services, and assist administrators with network troubleshooting. IP SLAs can be accessed using the Cisco software commands or Simple Network Management Protocol (SNMP) through the Cisco Round-Trip Time Monitor (RTTMON) and syslog Management Information Bases (MIBs).

- [Information About IP SLAs, on page 1](#)

Information About IP SLAs

IP SLAs Technology Overview

Cisco IP SLAs uses active traffic monitoring--the generation of traffic in a continuous, reliable, and predictable manner--for measuring network performance. 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. The information collected includes data about response time, one-way latency, jitter (interpacket delay variance), packet loss, voice quality scoring, network resource availability, application performance, and server response time. IP SLAs performs active monitoring by generating and analyzing traffic to measure performance either between Cisco devices or from a Cisco device to a remote IP device such as a network application server. Measurement statistics provided by the various IP SLAs operations can be used for troubleshooting, for problem analysis, and for designing network topologies.

Using 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 for new or existing IP services and applications. IP SLAs uses unique service level assurance metrics and methodology to provide highly accurate, precise service level assurance measurements.

Depending on the specific IP SLAs operation, statistics of delay, packet loss, jitter, packet sequence, connectivity, path, server response time, and download time can be monitored within the Cisco device and stored in both CLI and SNMP MIBs. The packets have configurable IP and application layer options such as a 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), a Virtual Private Network (VPN) routing/forwarding instance (VRF), and a URL web address.

Being Layer-2 transport independent, IP SLAs can be configured end-to-end over disparate networks to best reflect the metrics that an end-user is likely to experience. Performance metrics collected by IP SLAs operations include the following:

- 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
- Voice quality scores

Because IP SLAs is accessible using SNMP, it also can be used by performance monitoring applications like CiscoWorks Internetwork Performance Monitor (IPM) and other third-party Cisco partner performance management products. For details about network management products that use IP SLAs, see <http://www.cisco.com/go/ipsla>.

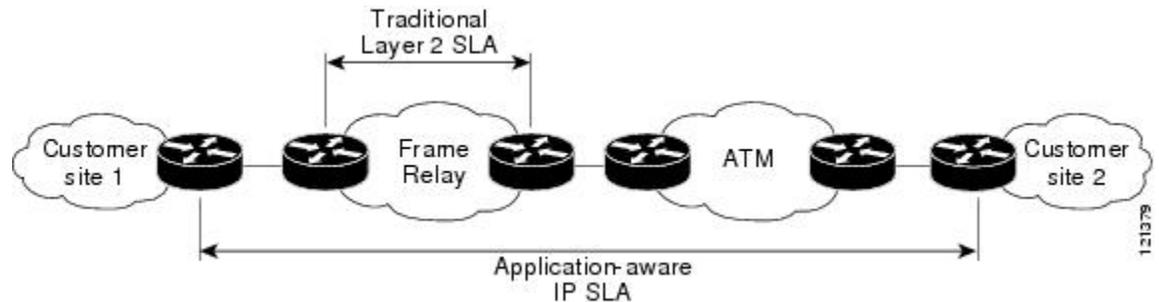
SNMP notifications based on the data gathered by an IP SLAs operation allow the router to receive alerts when performance drops below a specified level and when problems are corrected. IP SLAs uses the Cisco RTTMON MIB for interaction between external Network Management System (NMS) applications and the IP SLAs operations running on the Cisco devices. For a complete description of the object variables referenced by the IP SLAs feature, refer to the text of the CISCO-RTTMON-MIB.my file, available from the Cisco MIB website.

Service Level Agreements

Internet commerce has grown significantly in the past few years as the technology has advanced to provide faster, more reliable access to the Internet. Many companies now need online access and conduct most of their business online and any loss of service can affect the profitability of the company. Internet service providers (ISPs) and even internal IT departments now offer a defined level of service--a service level agreement--to provide their customers with a degree of predictability.

The latest performance requirements for business-critical applications, voice over IP (VoIP) networks, audio and visual conferencing, and VPNs are creating internal pressures on converged IP networks to become optimized for performance levels. Network administrators are increasingly required to support service level agreements that support application solutions. The figure below shows how IP SLAs has taken the traditional concept of Layer 2 service level agreements and applied a broader scope to support end-to-end performance measurement, including support of applications.

Figure 1: Scope of Traditional Service Level Agreement Versus IP SLAs



IP SLAs provides the following improvements over a traditional service level agreement:

- End-to-end measurements--The ability to measure performance from one end of the network to the other allows a broader reach and more accurate representation of the end-user experience.
- Sophistication--Statistics such as delay, jitter, packet sequence, Layer 3 connectivity, and path and download time that are broken down into bidirectional and round-trip numbers provide more data than just the bandwidth of a Layer 2 link.
- Ease of deployment--Leveraging the existing Cisco devices in a large network makes IP SLAs easier and cheaper to implement than the physical probes often required with traditional service level agreements.
- Application-aware monitoring--IP SLAs can simulate and measure performance statistics generated by applications running over Layer 3 through Layer 7. Traditional service level agreements can only measure Layer 2 performance.
- Pervasiveness--IP SLAs support exists in Cisco networking devices ranging from low-end to high-end devices and switches. This wide range of deployment gives IP SLAs more flexibility over traditional service level agreements.

When you know the performance expectations for different levels of traffic from the core of your network to the edge of your network, you can confidently build an end-to-end application-aware service level agreement.

Benefits of IP SLAs

- IP SLAs monitoring
 - Provides 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
 - Verifies that the existing QoS is sufficient for new IP services.
- Edge-to-edge network availability monitoring

- Provides proactive verification and connectivity testing of network resources (for example, indicates the network availability of a Network File System (NFS) server used to store business critical data from a remote site).
- Troubleshooting of network operation
 - Provides consistent, reliable measurement that immediately identifies problems and saves troubleshooting time.
- Voice over IP (VoIP) performance monitoring
- Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) performance monitoring and network verification

Restriction for IP SLAs

- The maximum supported scale number of CFM and IP SLA over the port channel is only .

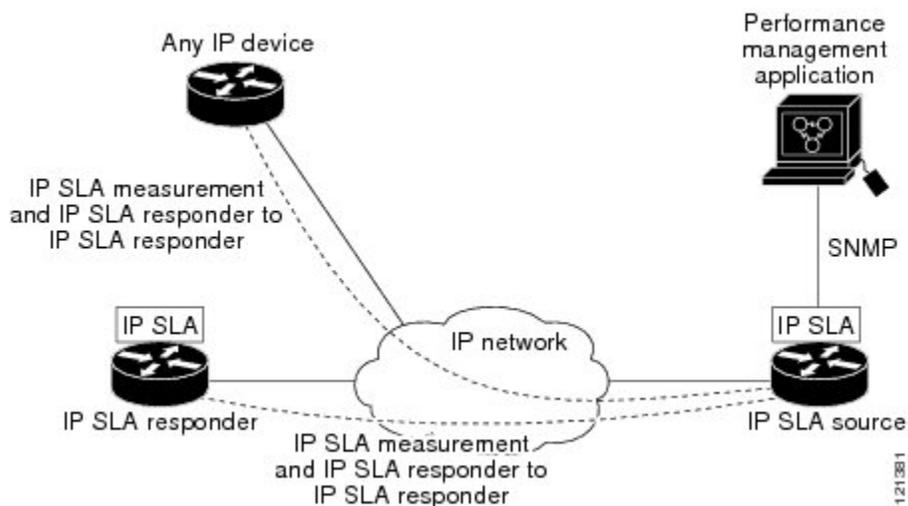
Network Performance Measurement Using IP SLAs

Using IP SLAs, a network engineer can monitor the performance between any area in the network: core, distribution, and edge. Monitoring can be done anytime, anywhere, without deploying a physical probe.

The IP SLAs Probe Enhancements feature is an application-aware synthetic operation agent that monitors network performance by measuring response time, network resource availability, application performance, jitter (interpacket delay variance), connect time, throughput, and packet loss. Performance can be measured between any Cisco device that supports this feature and any remote IP host (server), Cisco routing device, or mainframe host. Performance measurement statistics provided by this feature can be used for troubleshooting, for problem analysis, and for designing network topologies.

IP SLAs uses generated traffic to measure network performance between two networking devices. The figure below shows how IP SLAs starts when the IP SLAs device sends a generated packet to the destination device. After the destination device receives the packet, and depending on the type of IP SLAs operation, the device will respond 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 2: IP SLAs Operations



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

1. Enable the IP SLAs Responder, if appropriate.
2. Configure the required IP SLAs operation type.
3. Configure any options available for the specified IP SLAs 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 Cisco software commands or an NMS system with SNMP.

IP SLAs Responder and IP SLAs Control Protocol

The IP SLAs Responder is a component embedded in the destination Cisco routing device that allows the system to anticipate and respond to IP SLAs request packets. The IP SLAs Responder provides an enormous advantage with accurate measurements without the need for dedicated probes and additional statistics not available via standard ICMP-based measurements. The patented IP SLAs Control Protocol is used by the IP SLAs Responder providing a mechanism through which the responder can be notified on which port it should listen and respond. Only a Cisco device can be a source for a destination IP SLAs Responder.

The figure "IP SLAs Operations" in the "Network Performance Measurement Using IP SLAs" section shows where the IP SLAs Responder fits in relation to the IP network. The IP SLAs Responder listens on a specific port for control protocol messages sent by an IP SLAs operation. Upon receipt of the control message, the responder will enable the specified UDP or TCP port for the specified duration. During this time, the responder accepts the requests and responds to them. The responder disables the port after it responds to the IP SLAs packet, or when the specified time expires. For added security, MD5 authentication for control messages is available.

Enabling the IP SLAs Responder on the destination device is not required for all IP SLAs operations. For example, if services that are already provided by the destination device (such as Telnet or HTTP) are chosen,

the IP SLAs Responder need not be enabled. For non-Cisco devices, the IP SLAs Responder cannot be configured and IP SLAs can send operational packets only to services native to those devices.

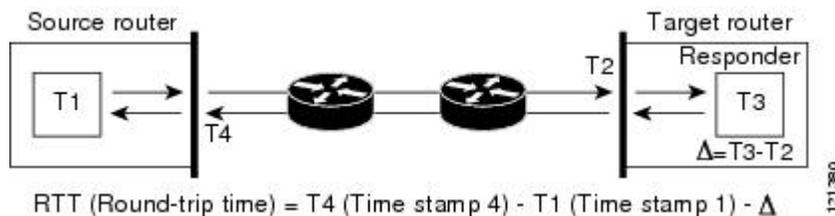
Response Time Computation for IP SLAs

Devices may take tens of milliseconds to process incoming packets, due to other high-priority processes. This delay affects the response times because the reply to test packets might be sitting on 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 IP SLAs Responder is being used), in order to determine true round-trip times. IP SLAs test packets use time stamping to minimize the processing delays.

When enabled, the IP SLAs Responder allows the target device to take two time stamps both when the packet arrives on the interface at interrupt level and again just as it is leaving, eliminating the processing time. At times of high network activity, an ICMP ping test often shows a long and inaccurate response time, while an IP SLAs test shows an accurate response time due to the time stamping on the responder.

The figure below demonstrates how the responder works. Four time stamps are taken to make the calculation for round-trip time. At the target device, 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 device where the incoming time stamp 4 (TS4) is also taken at the interrupt level to allow for greater accuracy.

Figure 3: 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 the configuration of both the source device and target device with Network Time Protocol (NTP) is required. Both the source and target need to be synchronized to the same clock source. One-way jitter measurements do not require clock synchronization.