



Overview of IP Service Level Agreements

This chapter provides an overview of Cisco NX-OS IP Service Level Agreements (SLAs), and includes the following sections:

- [Information About Cisco NX-OS IP SLAs, page 1](#)
- [Guidelines and Restrictions for Implementing IP Service Level Agreements, page 3](#)
- [Network Performance Measurement Using Cisco NX-OS IP SLAs, page 4](#)
- [Cisco NX-OS IP SLA Operation Types, page 5](#)
- [Cisco NX-OS IP SLA Responder and IP SLA Control Protocol, page 5](#)
- [Cisco NX-OS IP SLA Operation Scheduling, page 5](#)
- [Cisco NX-OS IP SLA Operation Threshold Monitoring, page 6](#)
- [MPLS VPN Awareness, page 6](#)
- [History Statistics, page 6](#)

Information About Cisco NX-OS IP SLAs

Many companies 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 (SLA), 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, Multiprotocol Label Switching (MPLS), and Virtual Private Networks (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. IP Service Level Agreements (SLAs) allow you to manage IP service levels for IP applications and services.

The Cisco NX-OS IP SLAs use active traffic monitoring--the generation of traffic in a continuous, reliable, and predictable manner--for measuring network performance. Cisco NX-OS IP SLAs send 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 the response time, one-way latency, jitter (interpacket delay variance), packet loss, voice quality scoring, network resource availability, application performance, and server response time. Cisco NX-OS IP SLAs performs active monitoring by generating and analyzing traffic to measure

performance either between Cisco NX-OS devices or from a Cisco NX-OS device to a remote IP device such as a network application server. Measurement statistics provided by the various Cisco NX-OS IP SLAs operations can be used for troubleshooting, problem analysis, and designing network topologies.




Note IPSLA do not support rollback. The rollback is related to IPSLA configuration via CLI.

Cisco NX-OS 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 Cisco NX-OS IP SLAs easier and cheaper to implement than the physical probes often required with traditional service level agreements.
- Application-aware monitoring—Cisco NX-OS 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—Cisco NX-OS IP SLAs support exists in Cisco networking devices that range from low-end to high-end switches. This wide range of deployment gives Cisco NX-OS IP SLAs more flexibility over traditional service level agreements.

The following figure shows how Cisco NX-OS IP SLAs have 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 Cisco NX-OS IP SLAs



Using Cisco NX-OS IP SLAs, you can measure, provide, and verify service level agreements. You can also analyze and troubleshoot network performance for IP services and applications. Depending on the specific Cisco NX-OS 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.

Because Cisco NX-OS IP SLAs are accessible using SNMP, it also can be used by performance monitoring applications such as CiscoWorks Internetwork Performance Monitor (IPM) and other third-party, Cisco partner performance management products.

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

Guidelines and Restrictions for Implementing IP Service Level Agreements

- IPv6 is available only from Cisco NX-OS Release 8.0 onwards.
- The maximum number of IP SLA-configurable operations that is supported by Cisco NX-OS software is 500.
- The current validated scale numbers for scheduling operations are as follows:
 - The number of UDP echo operations is 400 operations with default frequency.
 - The number of UDP jitter operations is 500 operations with default frequency.
 - The number of ICMP IPv4 or IPv6 echo operations is 500 operations with default frequency.
 - The number of TCP connect operations is 100 operations with default frequency.
 - We do not recommend scheduling more than 10 operations per second at the same start time because this might affect the performance of the network. We recommend the use of group scheduling configuration.



Note

Setting the frequency to less than 60 seconds will increase the number of packets that will be sent. But this might negatively impact the performance of IP SLA operations when the scheduled operations have the same start time.

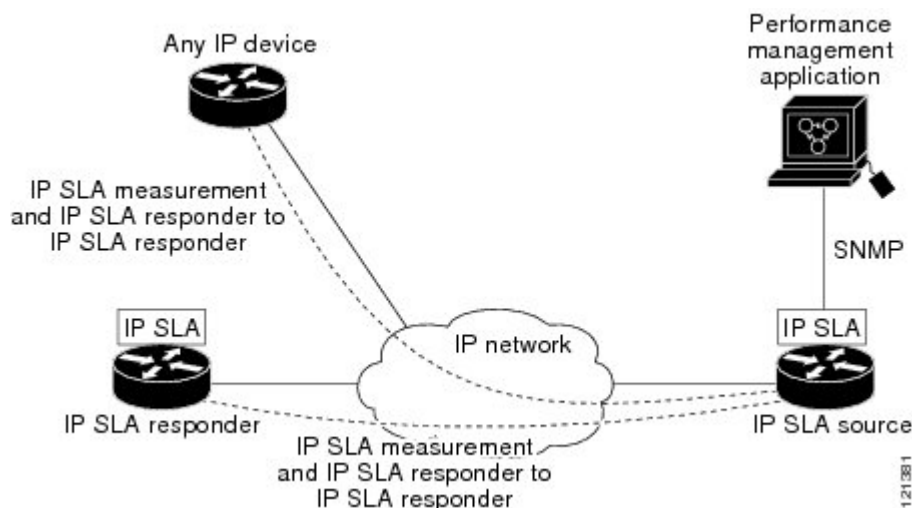
- IP SLA is not HA capable.
- Consider the following guidelines before configuring the **frequency**, **timeout**, and **threshold** commands:
 - For the UDP and ICMP jitter operations, we recommend the following guidelines:
 - $\text{frequency} > \text{timeout} + 2 \text{ seconds} + \text{num_packets} * \text{packet_interval}$
 - $\text{timeout} > \text{rtt_threshold}$
 - $\text{num_packet} > \text{loss_threshold}$
 - For all other IP SLAs operations, we recommend the $\text{frequency} > \text{timeout} > \text{rtt_threshold}$ guideline.

Network Performance Measurement Using Cisco NX-OS IP SLAs

Using Cisco NX-OS IP SLAs, you 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.

Cisco NX-OS IP SLAs use generated traffic to measure network performance between two networking devices such as switches. The following figure shows how Cisco NX-OS IP SLAs start when the Cisco NX-OS IP SLAs device sends a generated packet to the destination device. After the destination device receives the packet, and depending on the type of Cisco NX-OS IP SLAs operation, the device responds with time-stamp information for the source to make the calculation on performance metrics. A Cisco NX-OS 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: Cisco NX-OS IP SLAs Operations



To implement a Cisco NX-OS IP SLAs network performance measurement, you must perform these tasks:

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

Cisco NX-OS IP SLA Operation Types

The various types of Cisco NX-OS IP SLA operations include the following:

- Domain Name System (DNS)
- FabricPath echo
- ICMP echo (IPv4)
- ICMP echo (IPv6)
- UDP jitter
- UDP echo
- Transmission Control Protocol (TCP) connect

Cisco NX-OS IP SLA Responder and IP SLA Control Protocol

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

The IP SLAs Responder listens on a specific port for control protocol messages sent by a Cisco NX-OS IP SLAs operation. Upon receipt of the control message, the responder enables 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 Cisco NX-OS IP SLAs packet or when the specified time expires.

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 switch (such as Telnet or HTTP) are chosen, the IP SLAs Responder does not need to be enabled. For non-Cisco devices, the IP SLAs Responder cannot be configured and Cisco NX-OS IP SLAs can send operational packets only to services native to those devices.

Cisco NX-OS IP SLA Operation Scheduling

After a Cisco NX-OS IP SLAs operation has been configured, you must schedule the operation to begin capturing statistics and collecting error information. When scheduling, an operation can start immediately or start at a certain month, day, and hour. There is a pending option to set the operation to start at a later time. The pending option is also an internal state of the operation 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 Cisco NX-OS IP SLAs operation or a group of operations at one time.

Multiperations scheduling allows you to schedule multiple Cisco NX-OS IP SLAs operations using a single command through the Cisco NX-OS CLI or the CISCO RTTMON-MIB. This feature allows you to control the amount of IP SLAs monitoring traffic by scheduling the operations to run at evenly distributed times. This distribution of IP SLAs operations allows you to minimize the CPU utilization and enhance the scalability of the network.

For more details about the IP SLAs multioperations scheduling functionality, see the IP SLAs Multioperation Scheduler section.

Cisco NX-OS IP SLA Operation Threshold Monitoring

To support successful service level agreement monitoring or to proactively measure network performance, threshold functionality is essential. Consistent reliable measurements immediately identify issues and can save troubleshooting time. To roll out a service level agreement, you must have mechanisms that notify you immediately of any possible violations. Cisco NX-OS IP SLAs can send SNMP traps that are triggered by events such as the following:

- 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

Alternately, a Cisco NX-OS IP SLAs threshold violation can trigger another Cisco NX-OS IP SLAs operation for further analysis.

For more details on using thresholds with Cisco NX-OS IP SLAs operations, see the Proactive Threshold Monitoring for IP SLAs Operations section.

MPLS VPN Awareness

The Cisco NX-OS IP SLAs MPLS VPN Awareness feature allows you to monitor IP service levels within Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). Using IP SLAs within MPLS VPNs allows service providers to plan, provision, and manage IP VPN services according to the service level agreement for a customer. IP SLAs operations can be configured for a specific VPN by specifying a VPN routing and forwarding (VRF) name.

History Statistics

Cisco NX-OS IP SLAs maintain the following three types of history statistics:

- Aggregated statistics--By default, IP SLAs maintain two hours of aggregated statistics for each operation. The value from each operation cycle is aggregated with the previously available data within a given hour. The Enhanced History feature in IP SLAs allows for the aggregation interval to be shorter than an hour.

- Operation snapshot history--IP SLAs maintain a snapshot of data for each operation instance that matches a configurable filter, such as all, over threshold, or failures. The entire set of data is available and no aggregation takes place.
- Distribution statistics--IP SLAs maintain a frequency distribution over configurable intervals. Each time IP SLAs starts an operation, a new history bucket is created until the number of history buckets that matches the specified size or the lifetime of the operation expires. By default, the history for an IP SLAs operation is not collected. If history is collected, each bucket contains one or more history entries from the operation. History buckets do not wrap.

