



# Cisco SPS Overview

---

The Cisco SIP proxy server (Cisco SPS) is a call-control software package that enables service providers and others to build scalable, reliable packet voice networks and provides the capabilities required for call-session management in a VoIP network. It also provides a full array of call-routing capabilities for maximizing network performance in both small and large packet voice networks.

Cisco SPS has the capabilities of an edge proxy server, performing such functions as authentication, accounting, registration, network-access control, and security. It also has the capabilities of an infrastructure proxy server, performing such functions as next-hop routing based on received or translated destination URLs.

Cisco SPS is based on Session Initiation Protocol (SIP), a text-based protocol for setting up, modifying, and tearing down both unicast and multicast multimedia conferences.

This chapter provides an overview of Cisco SPS. It contains the following sections:

- [Prerequisites, page 1-1](#)
- [Restrictions, page 1-2](#)
- [SIP Basics, page 1-3](#)
- [Cisco SPS Basics, page 1-6](#)
- [Features, page 1-10](#)
- [Additional References, page 1-22](#)

## Prerequisites

### Hardware and Software

Hardware and software requirements are described in the *Cisco SIP Proxy Server Installation Guide*, available on your Cisco SPS CD or at

[http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod\\_technical\\_documentation.html](http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html).

## System Permissions, Strategy, and Functionality

### Permissions and Knowledge

You must have the following:

- Administrator privileges
- Familiarity with UNIX commands and shells

### Architecture Strategy

You must have determined which one of the following proxy-server architectures to implement:

- One standalone server
- A farm of two servers

### Configuration Strategy

You must have determined which one of the following strategies to use to configure SIP directives:

- GUI-based provisioning system, also called the provisioning GUI
- Manual configuration of the SIP directives (sipd.conf) file



#### Note

As a general rule, use GUI rather than manual methods to configure and operate your system.

- GUI methods are described in [Chapter 2, “Configuring Cisco SPS”](#) and [Chapter 3, “Operating and Maintaining Cisco SPS.”](#)
- Manual methods are described in [Appendix B, “Manual Configuration”](#) and [Appendix C, “Manual Operation and Maintenance.”](#)

## Restrictions

### Hardware and Software

Members of a proxy-server farm may run on different hardware. However, they must run the same operating system (mixed Linux/Solaris farms are not supported) and the same version of the Cisco SPS software.

### Time Synchronization

Members of a proxy-server farm must be time-synchronized to a common clock.

### Network Management

Do not run Network Information System (NIS)—a network-lookup service for managing a network of computers—on Cisco SPS systems. Doing so can cause long delays when you add a new user or user group. If you do run NIS, be sure to note any instructions for stopping it before doing configuration tasks.

### Permissions

You must run the proxy servers with either csp (default account ID) or root permission. Log ownership prevents csp from running a log after root has run the log.

# SIP Basics

SIP is the Internet Engineering Task Force (IETF) standard for multimedia conferencing over IP. A lightweight, generic, ASCII-based signaling protocol that provides session control, it is complementary to Simple Gateway Control Protocol (SGCP) and Media Gateway Control Protocol (MGCP), both of which provide device control. It offers significant performance advantages over H.323; it also more easily enables applications, provides a simple method to map Signaling System 7 (SS7) to Internet networks, provides loop detection, and supports conferencing.

Like other VoIP protocols, SIP provides signaling and session management within a packet-telephony network. Signaling allows call information to be carried across network boundaries. Session management controls the attributes of an end-to-end call.

SIP does the following, in the stated order:

1. Determines the location of the target endpoint.

SIP supports address resolution, name mapping, and call redirection.

2. Determines the media capabilities of the target endpoint.

Via Session Description Protocol (SDP), SIP determines the highest level of common services between the endpoints. Conferences are established through use of only those media capabilities that can be supported by all endpoints.

**Note**

A conference is an established session (or call) between two or more endpoints. Conferences consist of two or more users and can be established through multicast or multiple unicast sessions. In this document, a conference and a call are synonymous.

3. Determines the availability of the target endpoint.

If a call cannot be completed because the target endpoint is unavailable, SIP determines whether the called party is already on the phone or did not answer in the allotted number of rings. It then returns a message indicating why the target endpoint is unavailable.

4. Establishes a session between the originating and target endpoint.

If a call can be completed, SIP establishes a session between the endpoints. SIP also supports midcall changes such as addition of another endpoint to the conference or changing of a media characteristic or codec.

5. Transfers calls from one endpoint to another and terminates calls.

During call transfer, SIP simply establishes a session between the transferee and a new endpoint (specified by the transferring party) and terminates the session between the transferee and the transferring party. At the end of a call, SIP terminates sessions between all parties.

## SIP Components

SIP is a peer-to-peer protocol. The peers in a session are called user agents (UAs). A user agent can function in either of two roles: user-agent client or user-agent server. Definitions in the following sections are adapted from RFC 3261.

## User-Agent Client

A user-agent client (UAC) is a peer that initiates a SIP call request. More formally, it is a logical entity that creates a new request and then uses the client-transaction-state machinery to send it. The role of UAC lasts only for the duration of that transaction. In other words, if a piece of software initiates a request, it acts as a UAC for the duration of that transaction. If it receives a request later, it assumes the role of a user-agent server for the processing of that transaction.

Clients are of the following types:

- Phones that can act as either UAS or UAC. Softphones (PCs with installed phone capabilities) and Cisco SIP IP phones can initiate and respond to requests.
- Gateways that translate transmission format, communications procedures, and codecs between SIP conferencing endpoints and other terminal types. Other call-control functions include call setup and clearing on both the LAN and the switched-circuit network side.

The UAC core is the set of processing functions that are required of a UAC and that reside above the transaction and transport layers.

## User-Agent Server

A user-agent server (UAS) is a peer that receives and responds to call requests. More formally, it is a logical entity that generates a response to a SIP request; the response is to accept, reject, or redirect the request. This role lasts only for the duration of that transaction. In other words, if a piece of software responds to a request, it acts as a UAS for the duration of that transaction. If it generates a request later, it assumes the role of a UAC for the processing of that transaction.

A UAS can interact with other applications such as Lightweight Directory Access Protocol (LDAP) servers, databases, and Extensible Markup Language (XML) applications that provide back-end services such a directory, authentication, and billing.

Servers are of the following types:

- Proxy servers
- Redirect servers
- Registrar servers
- Location servers

### Proxy Servers

A proxy server initiates requests on behalf of and receives requests from a client. More formally, a proxy server is an intermediate entity that acts as both a server and a client for the purpose of making requests on behalf of other clients. A proxy server primarily handles routing, ensuring that a request is sent to another entity closer to the targeted user. Proxy servers are also useful for enforcing policy (for example, ensuring that a user is allowed to make a call). A proxy server interprets, and, if necessary, rewrites specific parts of a request message before forwarding it.

Proxy servers provide functions such as authentication, authorization, network access control, routing, reliable request retransmission, and security. They are often colocated with redirect or registrar servers.

### Redirect Servers

A redirect server is a UAS that generates 3xx responses to requests that it receives, directing the client to contact an alternate set of URIs. It receives requests, strips out the address in the request, checks its address tables for any other addresses that might be mapped to the one in the request, and returns the

results of the address mapping to the client for the next one or more hops that a message should take. The client then contacts the next-hop server or UAS directly. Redirect servers are often colocated with proxy or registrar servers.

### Registrar Servers

A registrar server is a server that accepts REGISTER requests from UACs for registration of their current location. It places the information received in those requests into the location service for the domain that it handles. Registrar servers are often colocated with proxy or redirect servers.

### Location Services

A location service is used by a SIP redirect or proxy server to obtain information about a called party's possible locations. It contains a list of bindings of address-of-record keys to zero or more contact addresses. Bindings can be created and removed in many ways; the SIP specification defines a REGISTER method for updating bindings. Location services are often colocated with redirect servers.

## SIP Network Architecture

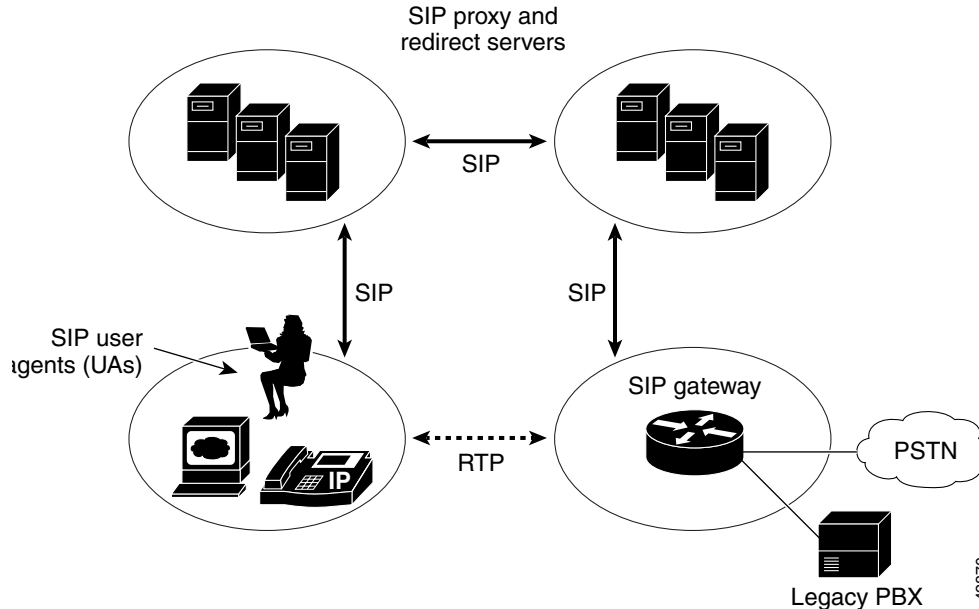
A typical SIP endpoint can function as either a UAC or a UAS during a transaction, its role depending on whether it initiates or receives and responds to a request.

A UAC can directly contact a UAS if it knows the location of the UAS and does not want any special services from the network. However, a UAC typically initiates a call through a proxy server and relies on the proxy server to locate the desired UAS and obtain any special services from the network. The SIP messaging path from UAC to UAS can involve multiple proxy servers, and in such scenarios Cisco SPS interfaces at a peer level with other proxy servers.

SIP requests can be sent with any reliable or unreliable protocol. Cisco SPS supports the use of User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Transport Layer Security (TLS) for sending and receiving SIP requests and responses.

[Figure 1-1](#) shows the architecture of a SIP network.

Figure 1-1 SIP Network Architecture



Users in a SIP network are identified by unique SIP addresses. A SIP address is similar to an e-mail address and is in the form `sip:userID@gateway.com`. The user ID can be either a username or an E.164 address.

Users register with a registrar server using their assigned SIP addresses. The registrar server provides this information to the location server upon request.

A user initiates a call by sending a SIP request to a SIP server (either a proxy server or a redirect server). The request includes the address of the caller (in the From header field) and the address of the called party (in the To header field).

Over time, a SIP end user—that is, a subscriber—might move between end systems. The location of the subscriber can be dynamically registered with the SIP server. The location server can use one or more protocols (including finger, rwhois, and LDAP) to locate the subscriber. Because the subscriber can be logged in at more than one station, the server might return more than one address. If the request comes through a SIP proxy server, the proxy server tries each of the returned addresses until it locates the subscriber. If the request comes through a SIP redirect server, the redirect server forwards all of the addresses to the caller listed in the Contact header field of the invitation response.

## Cisco SPS Basics

Cisco SPS sits in the core of a SIP network and routes calls among other proxy servers, voice gateways, IP endpoints (such as IP phones), and application servers. It enables user registration, authentication, and call-routing decisions to be made within the network, and identifies the next hop in the path to the called party. During call setup and teardown, Cisco SPS can generate accounting messages and pass them to a RADIUS server to form call-detail records (CDRs).

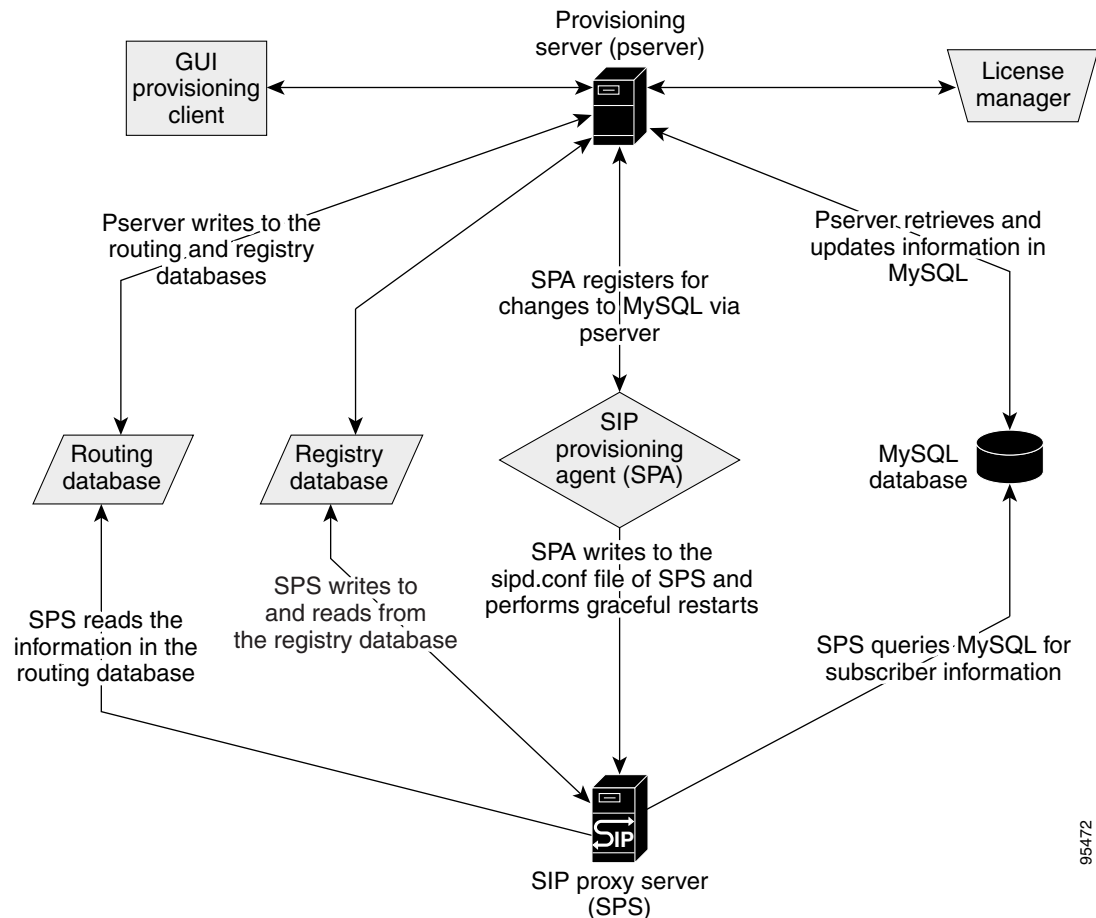
## Components

Cisco SPS components (see [Figure 1-2](#)) include the following:

- Proxy server
- Provisioning server
- MySQL database
- Provisioning GUI client of the Cisco SPS GUI-based provisioning system
- SIP provisioning agent
- Registry database
- Routing database
- License manager

Each of these elements is described after [Figure 1-2](#).

**Figure 1-2 Cisco SPS Components**



95472

**Proxy Server (sipd)**

The proxy server provides the primary capabilities required for call-session management in a VoIP network. It processes SIP requests and responses. It can be configured to function as a transaction stateful or stateless server and to provide additional server modes and features.

Cisco SPS supports the use of a proxy-server farm containing up to two proxy servers.

**Provisioning Server (pserver)**

The provisioning server is used by the Cisco SPS GUI-based provisioning system. A license manager is automatically installed when the provisioning server is installed.

**MySQL Database**

MySQL is a popular open-source database whose architecture makes it extremely fast and easy to customize. This database stores and accesses provisioning-system and subscriber-feature data. Subscriber features (call forwarding and local authentication, but not RADIUS) are automatically included.

Cisco SPS supports up to two replicated and synchronized MySQL databases.

**Provisioning GUI Client of the Cisco SPS GUI-Based Provisioning System**

The provisioning GUI client—also called the GUI-based provisioning system—retrieves and displays current information in the MySQL database by means of the provisioning server. It also retrieves and displays information outside of the MySQL database, such as registrations and routes.

If you modify this information, the GUI sends the updates to the provisioning server. If you modify information that is related to the license, the license manager does additional processing on this information, after which the provisioning server updates the database with the information. Meanwhile, each proxy-server provisioning client (there can be more than one on a farm) has registered to receive specific changes to the database. If a client is notified of a change, it requests that the provisioning server send the changed information and then updates its configuration file and its routing and registry data. The proxy server can use the new registry and routing information immediately, and uses the new configuration file if it is gracefully restarted.

The GUI is currently designed to change only certain things, such as sipd.conf, subscriber data, routes, and registries. It does not change other things such as spa.conf and ps.conf.

You can install the client independently of the provisioning server.

**SIP Provisioning Agent (spa)**

The SIP provisioning agent resides on a farm member and handles requests that the provisioning server gets from the GUI. It receives requests from the provisioning server, accesses and updates (as needed) the SIP directives (sipd.conf) file, and provides feedback, by way of the provisioning server, to the GUI.

The IP address in the spa (spa.conf) file is needed so that the spa process can determine whether the data is intended for this particular process. The spa process makes the determination by comparing the data with the IP address in spa.conf; if they do not match, it decides that the data is not intended for this process.

Changes to the MySQL database that the SIP provisioning agent does not register for include subscriber information such as call-forwarding destinations and authentication passwords. These changes pass from the database directly to sipd.

### Registry Database

The registry database contains location information for registered endpoints. The database is stored in memory-mapped files within shared memory so that information persists between restarts. Registry information is exchanged with all members of a proxy-server farm. The database contains two types of information: dynamic information that is received from endpoints when they register their contact information and static information that is configured by means of the provisioning GUI client.

Information pertaining to a single registered user (SIP endpoint) is called a registration. The registry database is thus a collection of registrations.

Registration is made or renewed either dynamically or statically:

- **Dynamic registration**—An endpoint registers itself by sending a REGISTER request to a Cisco SPS, which then adds a registration to the registry. By default, SIP endpoints register once every hour, and each previous registration expires after an hour. Dynamic registration is the norm.
- **Static registration**—You, as system administrator, explicitly create a registration for a SIP endpoint that is incapable of registering for itself (such as a SIP gateway). Static registrations are generally permanent. They are not the norm, because such endpoints are generally represented by a route instead of a registration.

### Routing Database

The routing database contains static route information that the proxy server uses to forward requests toward endpoints that are either not registered with the local registrar server, reside within a different domain, or exist in the PSTN. Static routes are configured based on next-hop IP addresses or next-hop domains. Routing information is configured by means of the provisioning GUI client. As with the registry database, the routing database is stored in memory-mapped files within shared memory so that information persists between restarts.

### License Manager

The license manager maintains the license key required for activation of the proxy server. You need a valid license to install, start, and restart SPS. Two types of licenses are available: evaluation and permanent. An evaluation license expires after a set period of time; after the expiration date, the proxy server can only be gracefully restarted.

## Server Modes

Cisco SPS can function as a proxy server, a redirect server, or a registrar server. The first two can be transaction-stateful or transaction-stateless.

- A transaction-stateful server remembers incoming and outgoing requests, provides reliable retransmission of proxied requests, and returns the best final responses.
- A transaction-stateless server forgets all information once a request or response has been processed. It merely forwards requests and responses.

A transaction includes the following:

- Received request
- Request or requests (if forked) forwarded downstream
- Responses received from downstream hosts
- Best response returned upstream

**Proxy Server**

A proxy server is an intermediate device that receives SIP requests from a client and then initiates requests on the client's behalf. It can be transaction-stateful or transaction-stateless. If transaction-stateful, it also does the following:

- Creates a transaction control block (TCB)
- Remembers incoming and outgoing requests
- Provides reliable retransmissions for unreliable transports
- Returns the best final response or responses upstream

Cisco SPS functions by default as both a proxy server and a registrar server for all calls, although you can change this default functionality.

**Redirect Server**

A redirect server does the following:

- Accepts SIP requests
- Maps the address in the request-URI to zero or more new addresses
- Returns these addresses as contacts in a SIP 3xx response to the UAC

A redirect server can be transaction-stateful or transaction-stateless. If transaction-stateless, it does not create a TCB on receiving an INVITE request.

**Registrar Server**

A Cisco SPS functioning in the role of registrar server does the following:

- Processes requests from UACs for registration of their current location
- Maintains registration information, which it can share with other registrar servers in its server farm
- Provides location services to the proxy server

Cisco SPS functions by default as both a proxy server and a registrar server for all calls, although this default functionality can be changed.

You must configure each registrar server to function also as either a proxy server or a redirect server. Standalone registrar servers are not supported.

## User IDs and Passwords

As a general rule, use the default user IDs and passwords listed in [Table 1-1](#).

**Table 1-1 System Account IDs and Passwords**

System	Default Account ID	Default Password
Cisco SPS	csps	cspsuser
MySQL	guest	nobody

## Features

Cisco SPS supports the following major features and capabilities, listed alphabetically:

- [Access and Error Logging, page 1-11](#)

- [Accounting, page 1-12](#)
- [Address Translation, Next-Hop Routing, and IP Resolution, page 1-12](#)
- [Authentication and Authorization, page 1-15](#)
- [DNS Support, page 1-16](#)
- [IP Security, page 1-17](#)
- [Proxy-Server Farms, page 1-17](#)
- [Registrar Server for Multiple Domains, page 1-17](#)
- [Registry and Route Configurations, page 1-19](#)
- [Spiralled and Looped Request Detection, page 1-20](#)
- [Subscribers, page 1-21](#)
- [TLS Support, page 1-21](#)

**Note**

For information on configuring these features and capabilities, see [Chapter 2, “Configuring Cisco SPS.”](#)

In addition, Cisco SPS provides the following additional capabilities, some of which were described earlier:

- Functions as a transaction-stateful or transaction-stateless proxy server, transaction-stateful or transaction-stateless redirect server, and registrar server
- Handles call forwarding
- Has a GUI-based provisioning system for configuring the server and accessing the embedded routing and registry databases and the embedded MySQL subscriber database
- Forks requests and distinguishes spiralled requests from looped requests
- Supports SIP over UDP or TCP
- Interoperates with Cisco SIP gateways, SIP IP phones, and unified messaging
- Supports parameters relevant to Network Address Translation (NAT) traversal
- Has a domain-specific registration, authentication, accounting, and subscriber database
- Handles preauthorization queries to a resource-policy-management system
- Offers an SNMP interface by means of CIAgent with basic platform MIBs for server status and start/stop/restart
- Provides database robustness:
  - Registration and route databases are memory-mapped files that facilitate both quick access and persistence should servers be stopped and restarted.
  - The subscriber database is stored in MySQL. If a proxy-server farm contains multiple MySQL databases, subscriber data is synchronized between them and persists between server stops and restarts.

## Access and Error Logging

Cisco SPS uses both standard Apache and SIP-specific logging functionality. Standard Apache logging functionality is configured on Cisco SPS as a whole. SIP-specific logging functionality is configured on a per-module basis.

For information on log locations, interpreting log files, and customizing the type and amount of information that they include, see [Chapter 3, “Operating and Maintaining Cisco SPS.”](#)

## Accounting

If you enable accounting services and configure the interface to a RADIUS server, Cisco SPS sends accounting records to the RADIUS server.

Cisco SPS uses basic start-stop records with a combination of standard RADIUS attributes and Cisco vendor-specific attributes (VSAs). Additionally, you can configure Cisco SPS to add any desired SIP headers as VSAs in accounting requests.

Cisco SPS receives the BYE message only if record-route is enabled. In an unsuccessful call, it writes a STOP message when the best non-200 (4/5/6XX) response is received for an INVITE message.



### Tip

---

For information about Cisco SPS and the RADIUS server, refer to the RADIUS interface specifications at <http://www.cisco.com/univercd/cc/td/doc/product/voice/sipproxy/index.htm> or <http://www.cisco.com/en/US/products/sw/voicesw/ps2157/index.html>.

---

## Address Translation, Next-Hop Routing, and IP Resolution

Cisco SPS performs the following steps to deliver messages from endpoint to endpoint:

1. Address translation—Translates an incoming request-URI into an outgoing request-URI.
2. Next-hop routing—Obtains a set of fully qualified domain names (FQDNs) or IP addresses with transport type and port numbers for each of the SIP entities found in the translation step. This step involves the following features and more:
  - SRV lookup for static route—SRV lookup on the Static\_Route\_NextHop field if the Static\_Route\_NextHopPort field in a static route is not specified or is zero.
  - Registration, Admission, and Status Protocol (RAS) LRQ message transmission to gatekeeper—Sending the LRQ message to the H.323 gatekeeper to obtain the next-hop gateway transport address.
3. IP resolution—Converts each next hop found in the next-hop route lookup step into an IP address.



### Note

---

If a route header is present in the SIP request message, Cisco SPS bypasses translation and next-hop routing steps. You must enable record-route on the server for subsequent requests to contain a route header.

---

## Address Translation

During address translation, Cisco SPS processes the request-URI of an incoming request and returns a list of contacts, each providing a URL for use in the outgoing request.

If you enable number expansion, Cisco SPS applies the global set of expansion rules to the user portion of the relevant URLs for which the host portion is the Cisco SPS. For REGISTER messages, this applies to the To, From, Contact, and optionally the Authorization headers. For INVITE messages, this applies to the Request URI, From, and optionally the Proxy-Authorization headers.

**Note**

Headers are not rewritten. The expanded versions are used internally for authentication, accounting, translation, and routing purposes.

Cisco SPS translation modules, in the order in which Cisco SPS calls them, are as follows:

- Call Forward Unconditional (mod\_sip\_call\_forward)
- Registry (mod\_sip\_registry)
- ENUM (mod\_sip\_enum)
- GKTMP (mod\_sip\_gktmp)

The first module to return one or more contacts completes the translation step, and the remaining modules are not called. For example, if the Registry module returns a contact, then neither the ENUM nor the GKTMP module is called. If none of the translation modules returns a contact, the core proxy module (mod\_sip) returns a contact based on the incoming Request-URI and that Request-URI is used in the next-hop routing step.

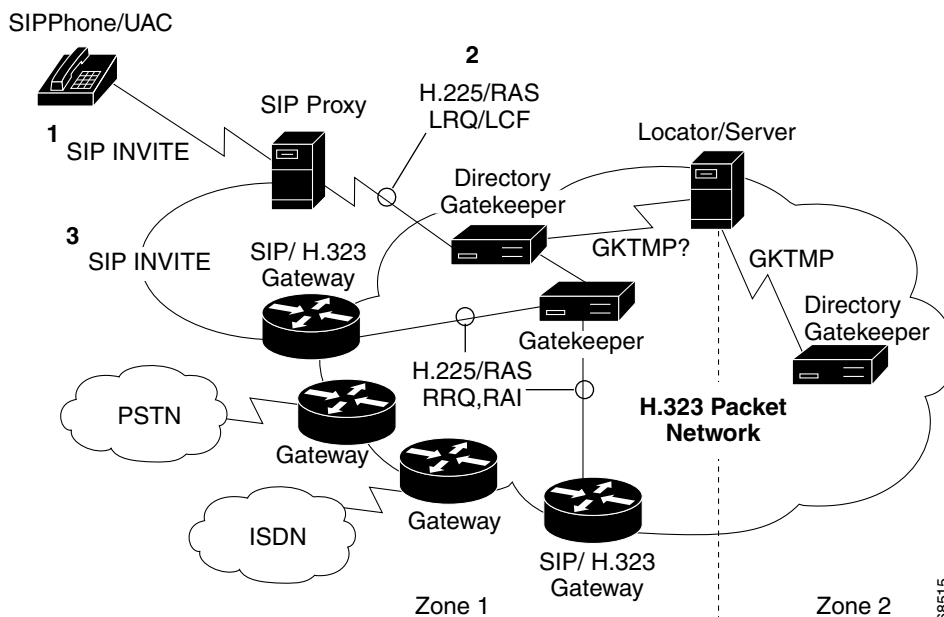
## Next-Hop Routing

The next step in SIP request processing is to determine the next-hop route for each contact. Next-hop routing takes each translated request-URI (contact) and locates a set of next-hop SIP entities capable of routing a message to the new request-URI. This step involves two advanced features:

- SRV Lookup for Static Route—If the Static\_Route\_NextHopPort field of a static route is not specified or is zero, Cisco SPS performs an SRV lookup on the Static\_Route\_NextHop field.
- H.323 RAS module—The RAS module allows communications between a SIP proxy server and an H.323 gatekeeper. Cisco SPS can send an ASN.1-encoded RAS LRQ message to an H.323 gatekeeper that responds with a RAS LCF message.

[Figure 1-3](#) illustrates how the proxy server sends an LRQ H.225 RAS message to a gatekeeper that responds with an LCF message that can contain a gateway transport address.

**Figure 1-3 Cisco SPS and H.323 Gatekeeper Message Processing**



The RAS module supports the following:

- H.323 request in progress (RIP) message—This message contains a time-delay value. Upon receipt of this message, Cisco SPS resets its timer and waits for the final response from the gatekeeper.
- Directory gatekeeper—This gatekeeper forwards requests to the other gatekeepers and returns the highest-priority and lowest-cost gateway information to the requesting endpoint. The value of the time-to-live (TTL) field in the LRQ message must be greater than zero. The default is 6 hops.
- Sequential LRQ—Cisco SPS can send LRQ requests sequentially based on the priority specified for the gatekeeper clusters in the sipd.conf file. Configurable parameters include a timeout value for each individual request, a total LRQ request window, and an indication as to whether the module should wait for the best LCF within the LRQ time window or return the first valid LCF response to the proxy server.
- Blast LRQ—Cisco SPS can send LRQ requests to all gatekeeper clusters before listening to the responses from the gatekeepers. It can also be configured to either take the first valid LCF response or wait for the best response within the LRQ time window.
- Multi-alternate endpoint—If the selected LCF from a gatekeeper contains alternate endpoints, Cisco SPS can parse these endpoints and store them in the route table. The routes in the alternate endpoints have lower priority than the route for the primary endpoint.
- Redundant gatekeeper—To establish redundancy, you must maintain a list of prioritized gatekeeper clusters in the Cisco SPS configuration. The following trial process occurs:
  1. The proxy server tries one of the gatekeepers randomly in the cluster that has the highest priority.
  2. If the trial goes beyond the timeout value specified in the RAS`TimeoutInterval` directive, the proxy server tries the next gatekeeper in the same cluster in a round-robin manner.
  3. If the proxy server receives an LRJ message, it tries a gatekeeper in the next-highest-priority cluster and so forth.

- Tech prefix—Cisco SPS prepends the technology prefix (one, two, or three digits followed by #) to the expanded dialed number in the LRQ request. The prefix can also be included in the Request-URI of the INVITE message forwarded to the SIP/H.323 gateway. Multiple technology prefixes can be configured in the sipd.conf file based on the expanded dialed number.
- Pavo extension—Cisco SPS can include Pavo extensions (CallIdentifier, RedirectIEInfo, CallingOctet3a) in the LRQ from the CC-Diversion header of the incoming INVITE message.

During location of the next-hop SIP entities, the following occurs:

1. If the host portion of the new request-URI is the address (FQDN or IP address) of the server itself, next-hop routing is performed by means of the user portion of the request-URI. E.164 routing makes use of this method.
2. If the Static\_Route\_NextHopPort field of a static route is not specified or is 0, Cisco SPS tries to do an SRV lookup on the Static\_Route\_NextHop field.
  - If the lookup succeeds, it uses the algorithm outlined in RFC 2782 to select one destination.
  - If the lookup fails, it tries alternate destinations and the proxy server does a simple DNS A lookup on the Static\_Route\_NextHop field of the static route. This field should contain a name for the A lookup.
3. If no route is found in the E.164 routing, Cisco SPS sends an RAS LRQ message to the H.323 gatekeeper cluster that has the highest-priority gateway. According to configurations and availability of the gatekeepers being contacted, one of the following messages is returned before the time limit in the configured timeout window expires:
  - LCF (location confirm)
  - LRJ (location reject)
  - RIP (response in progress)The RIP message is activated when the gatekeeper cluster connects to a remote gatekeeper (by means of UDP).
4. If the host portion of the new request-URI is not the address (FQDN or IP address) of the server itself, Cisco SPS performs domain routing using the host portion of the Request-URI.

## IP Resolution

IP resolution is the conversion of each hop found by means of next-hop routing into an IP address. Standard IP resolution (through gethostbyname) is performed by either DNS, NIS, NIS+, or host file, depending on the IP resolution subsystem configured on the system where Cisco SPS is located.

## Authentication and Authorization

The Cisco SIP proxy server can provide authentication and authorization.

Two locations are supported: Authentication can occur at a RADIUS server or at the proxy server.

Two types of authentication are supported: HTTP digest authentication and HTTP basic authentication, both as described in RFC 2617. Either type can occur at either location.

A RADIUS server operates in accordance with the RADIUS protocol—an IETF protocol based on UDP. RADIUS supports the exchange of a set of attribute/value pairs between client and server. For example, a Cisco SPS acting as a RADIUS client exchanges attribute/value pairs with a RADIUS server to provide authentication.

During authentication, the UAC password is stored as follows:

- For RADIUS-supported authentication, it is stored at the RADIUS server.
- For proxy-supported authentication, it is stored in a subscriber table in a MySQL database.

## Authentication

The default authentication scheme is HTTP digest authentication performed at the Cisco SPS.

When digest authentication and basic authentication are performed at the proxy server, the username, as found in the authorization header or the proxy-authorization header, is the key to query the MySQL database.

If authentication takes place at the RADIUS server, Cisco SPS passes the username as one of the attribute/value pairs to the RADIUS server, where it can be used to key the user search before authentication. Additionally, you can configure Cisco SPS to add any desired SIP headers as VSAs in the authentication request to the RADIUS server.

Cisco SPS can expand the UserName before MySQL lookup or before passing it to the RADIUS server. This enables phone numbers to be expanded to full E.164 numbers before being processed.

If the virtual proxy host feature is enabled, the Username@domain (username found in the Authorization/proxy-Authorization header; domain name found in the From header) is the key used to query the MySQL database, just as when authentication is performed on a RADIUS server and the RadiusUserNameAttrAddDomain directive is enabled. Cisco SPS passes the Username@domain directive as one of the attribute/value pairs to the RADIUS server, where it can be used as the key for user searches.



### Note

---

Cisco SPS does not support native Apache-based virtual hosts. Native Apache-based virtual host refers to providing the illusion of more than one server on one system. For example, companies sharing a web server can have their own domains (www.company1.com and www.company2.com) and access to the web server.

---

Cisco SPS provides access-control lists and user authentication that you can combine to provide more complex access control. Some devices such as SIP gateways do not authenticate themselves. You can specify them in the access-control list, which instructs SPS to accept REGISTER and INVITE requests from them without the need for additional authentication.

You can configure Cisco SPS to challenge REGISTER and INVITE requests from other devices for user authentication.

## Authorization

An authenticated user is authorized. Currently, you cannot authorize authenticated users according to specific user capabilities for service-provider voice applications.

## DNS Support

Cisco SPS implements (in accordance with RFC 3263) DNS procedures for SIP clients and servers, including the use of naming-authority pointer (NAPTR), server (SRV), and ready-to-receive (RR) lookups by SIP entities. It also uses SIP-specific extensions for handling failures during lookups. RFC 3263 obsoletes RFC 2543 guidelines for DNS procedures and provides DNS configuration guidelines to SIP server administrators to help create secure and correct SIP services.

Cisco SPS uses DNS procedures for the following reasons:

- To support scalability and high availability. Typically, customers deploy SIP services in a farm of homogeneously configured proxy servers. DNS enables you to configure these farm members with prioritization and weights (thus supplying a crude level of capacity-based load balancing). DNS also provides failover capability in both upstream and downstream directions.
- To discover SIP servers in external domains. Specifically, SPS needs to determine the IP address, port, and transport protocol for the server. SPS can support TCP, UDP, and TLS for SIP signaling. SPS needs to be able to automatically determine which transport protocols are available with next-hop servers and in what order of preference.

Cisco SPS configuration for advanced DNS support is tightly coupled with the type of DNS support and configuration available. It is highly recommended that you configure the domain for which you are installing the proxy server using the guidelines mentioned in this document.

## IP Security

IP security (IPSec) provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts at the network layer, protecting and authenticating IP packets moving between participating IPSec devices (peers) such as Cisco SPS and a UAC, SIP gateway, or another Cisco SPS. With IPSec, data can move across a public network without fear of observation, modification, or spoofing.

All IPSec combinations can successfully secure traffic to and from Cisco SPS on the specified Solaris and Red Hat Linux platforms.

## Proxy-Server Farms

Cisco SPS is designed for deployment in proxy-server farms of up to two servers that act as one virtual server. Multiple proxy servers share database information and provide both redundancy and high availability.

A farm is represented by a domain or shared server name, where DNS determines the appropriate IP address. You enter only one host name or IP address; the information is immediately made available to both farm members.

Registry information and routing information are synchronized across farm members dynamically. You enter a registration or route on one farm member; the new information is immediately made available to both farm members.

A farm supports up to 20,000 static registries and routes.

## Registrar Server for Multiple Domains

Cisco SPS can act as a registrar server for multiple domains: a single domain (ProxyDomain) for fully expanded E.164 numbers plus and zero or more additional domains within which private name spaces can exist.

For example, the single E.164 domain can be cisco.com; additional private domains can be a.com and b.com. In this situation, Cisco SPS accepts REGISTER messages for the following domains:

- <\*>@a.com
- <\*>@b.com
- <\*>@cisco.com

- <e.164-number>@cisco.com
- <e.164-number>@a.com
- <e.164-number>@b.com

Cisco SPS enters registrations for these domains in the registry database as follows:

- It treats an INVITE message that is received for <e.164-number>@\*.com as an INVITE message for <e.164-number>@cisco.com. All registrations for <e.164-number>@\*.com are represented by a single entry in the registry database.
- It treats an INVITE message that is received for user@a.com, user@b.com, user@cisco.com as an INVITE message for user@a.com, user@b.com, or user@cisco.com, respectively. Registrations for user@a.com, user@b.com, and user@cisco.com result in three separate entries in the registry database.

Note that creation and administration of multiple domains of overlapping name spaces requires coordinated effort on the part of the Cisco SPS administrator and the administrators of the supported domains.

As an example of multiple domains, consider CompanyA and CompanyB in area code 408. CompanyA uses four-digit extensions 2000 to 7999 and CompanyB uses four-digit extensions 2000 to 7999. No DID numbers are associated with these extensions. For users at CompanyA and CompanyB who need DID numbers, CompanyA gets 1-408-555-[8000-8999] and CompanyB gets 1-408-666-[8000-8999]. Cisco SPS enters registrations in this example as follows:

- 2xxx at CompanyA registers as 2xxx@companyA.com
- 2xxx at CompanyB registers as 2xxx@companyB.com
- 8xxx at CompanyA registers as +14085558xxx@companyA.com
- 8xxx at CompanyB registers as +14086668xxx@companyB.com




---

**Note** x is a wildcard and can be replaced by any digit. For example, 2xxx can be 2000, 2001, 2999, and so on.

---

Call processing occurs as follows:

1. 2001 calls 2000
  - 2001 at CompanyA calls 2000 at CompanyA as 2000@companyA.com [1]
  - 2001 at CompanyA calls 2000 at CompanyB as 2000@companyB.com [2]
  - 2001 at CompanyB calls 2000 at CompanyB as 2000@companyB.com [1]
  - 2001 at CompanyB calls 2000 at CompanyA as 2000@companyA.com [2]
2. 8001 calls 2000
  - 8001 at CompanyA calls 2000 at CompanyA as 2000@companyA.com [1]
  - 8001 at CompanyA calls 2000 at CompanyB as 2000@companyB.com [2]
  - 8001 at CompanyB calls 2000 at CompanyB as 2000@companyB.com [1]
  - 8001 at CompanyB calls 2000 at CompanyA as 2000@companyA.com [2]
3. 2001 calls 8000
  - 2001 at CompanyA calls 8000 at CompanyA as +14085558000@companyA.com[3]
  - 2001 at CompanyA calls 8000 at CompanyB as +14086668000@companyA.com[4]

- 2001 at CompanyB calls 8000 at CompanyB as +14086668000@companyB.com[3]
  - 2001 at CompanyB calls 8000 at CompanyA as +14085558000@companyB.com[4]
4. 8001 calls 8000
    - 8001 at CompanyA calls 8000 at CompanyA as +14085558000@companyA.com[3]
    - 8001 at CompanyA calls 8000 at CompanyB as +14086668000@companyA.com[4]
    - 8001 at CompanyB calls 8000 at CompanyB as +14086668000@companyB.com[3]
    - 8001 at CompanyB calls 8000 at CompanyA as +14085558000@companyB.com [4]
  5. 2000 and 8000 call 18183635839
    - 2000 at CompanyA calls 18183635839 as +18183635839@companyA.com [5]
    - 8000 at CompanyA calls 18183635839 as +18183635839@companyA.com [5]
    - 2000 at CompanyB calls 18183635839 as +18183635839@companyB.com [5]
    - 8000 at CompanyB calls 18183635839 as +18183635839@companyB.com [5]
  6. Authorized user John Doe (with neither CompanyA nor CompanyB) places calls as follows:
    - Calls 2000 at CompanyA as 2000@companyA.com [2]
    - Calls 2000 at CompanyB as 2000@companyB.com [2]
    - Calls 8000 at CompanyA as +14085558000@<proxy>.com [6]
    - Calls 8000 at CompanyB as +14086668000@<proxy>.com [6]

**Note**

- [1] The phones at both companies expand 2xxx to 2xxx@<proxy>.
- [2] This is done through full URL dialing.
- [3] The phones at companyA expand 8xxx to +14085558xxx@<proxy> and the phones at companyB expand 8xxx to +14086668xxx@<proxy>.
- [4] The phones at both companies expand [2-9]xxxxxx to +1408xxxxxx@<proxy>. All authenticated users, regardless of domain, have access to the same routes for forwarding calls for +1408xxxxxx.
- [5] The phones at both companies expand 1xxxxxxxx to +1xxxxxxxx@<proxy>. All authenticated users, regardless of domain, have access to the same routes for forwarding calls for +1xxxxxxxx.
- [6] Cisco SPS supports a default domain for E.164 numbers; such numbers are designated by the plus (+) sign.

## Registry and Route Configurations

### Basic Principles

A registry is a memory-mapped file maintained on each Cisco SPS server. It contains information for up to 20,000 registered users. (A user is a SIP endpoint that receives calls.) The file can be accessed quickly and persists across restarts. Registries are synchronized among farm members.

Information pertaining to a single registered user (SIP endpoint) is called a registration. The registry is thus a collection of registrations.

SPS does not determine capabilities based on the registration; it determines only location (that is, transport, IP address, and port to be used to reach the endpoint). Static registration information exists simply to interwork with SIP endpoints that are not capable of registering for themselves.

## Registry and Route Configuration Strategies

Whether you create registrations or routes is mostly a matter of preference. Most administrators tend to use static registration if the number of FXS ports is small and a route if the number is large. They also tend to use static registration if an FXS port is associated with a subscriber (for example, for call forwarding).

An example of a nonsubscriber registry is a Cisco IOS gateway connected to a PBX or a gateway that has analog phones connected via foreign exchange station (FXS) ports. If the gateway cannot dynamically register its endpoints with Cisco SPS, you must statically register them so that they can receive sessions from other endpoints. The FXS ports on the gateways are therefore subscribers with static registrations.

If a gateway has FXS ports for just a small number of extensions—say, for example, 5000 and 5001—you can create subscribers for 5000 and 5001, each with call forwarding to voice mail (5000@voicemail-server, 5001@voicemail-server), and then enter static registrations for 5000 and 5001 with the following contacts:

```
5000@<gateway-ip-address>:<gateway-port>;user=phone
5001@<gateway-ip-address>:<gateway-port>;user=phone
```

If a gateway has many (say, 10 or more) FXS ports, you might prefer to create a route rather than many static registrations. For example, if 100 FXS ports are mapped to users 5000–5099, you could create 100 subscribers with no static registrations, and then create a route such as the following:

```
destination-pattern: 50..
next-hop: <gateway-ip-address>
etc.
```

You could even create a route for a single endpoint such as the following:

```
destination-pattern: 5043
next-hop: <gateway-ip-address>
etc.
```

## Spiralled and Looped Request Detection

A spiralled request is a request that revisits Cisco SPS with a new request-URI and that SPS considers to be a new transaction. A looped request is a request that contains the proxy's Via header; the request-URI is the same as it was when it first visited the proxy server.

At the proxy server, this has the following effects:

- The same call can be logged more than once.
- Different features can be invoked.
- If the Record-Route field in the INVITE message is enabled in the proxy server, the record-route procedure is performed more than once.



### Note

The Call-ID, From, To, and CSeq.seqnum fields of a spiralled request are the same as those of the request that visited the proxy server the first time.

## Subscribers

A subscriber is a user (SIP endpoint) that has static subscriber information such as user ID, password, and call-forwarding settings. Subscribers generally register dynamically by means of a periodic REGISTER request from the subscriber's SIP endpoint (called a SIP user agent).

Subscriber records reside in the MySQL database subscriber table and are used for authorization and per-user call forwarding. Subscriber records are synchronized between replicated MySQL databases.

A subscriber can have multiple SIP agents and therefore multiple registrations. For example, a subscriber with a given user ID, password, call forwarding number, and so on might have an office phone, a cell phone, and a home phone, all of which register. When the subscriber receives a call, all three phones ring.

In general, you should not need to register a subscriber statically. In any case, you should never register the same contact for a subscriber both statically and dynamically. A mismatch could result between what the provisioning system considers to be registered and the registry information in shared memory that call processing uses.

## TLS Support

Transport Layer Security (TLS) is an IETF protocol that replaces Secure Socket Layer (SSL) encryption technology. It provides secure transactions such as transmission of credit-card numbers for e-commerce.

Cisco SPS provides TLS for SIP signaling according to RFC 3261 recommendations. It supports TLSv1 and the following two cipher suits:

- TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA
- TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA

The proxy server handles TLS using RSA key exchange. (RSA stands for Rivest, Shamir, and Adelman, inventors of the technique.) RSA is a public-key cryptographic system for encryption and authentication with Advanced Encryption Standard (AES)-128 or 3 Data Encryption Standard (DES) crypto methods in cipher-block-chaining (CBC) mode. It uses Hash-Based Message Authentication Code—Secure Hash Algorithm (HMAC-SHA) for message-authentication check. Cisco SPS does not explicitly change ciphers (a cipher is a cryptographic algorithm for encryption and decryption) during a session, but it allows session renegotiation (rehandshake) by clients. Cisco SPS supports session resumption; currently only 256 sessions can be cached.

Cisco SPS does not act as a certification authority, but requires that proper X.509 certificates and keys be installed. (For information on how to install TLS certificates and keys, refer to the *Cisco SIP Proxy Server Installation Guide*.) From the Cisco SPS point of view, it does not matter if the certificate is from a well-known authority or is self-signed by the user. The proxy server needs only to have its certificate, key, and the root certificates installed as part of configuration.

Cisco SPS does not distribute certificates during registration. Rather, you must manage and distribute certificates. You are free to act as certification authority and sign or distribute certificates to your clients, in which case you need to install your self-signed root certificate in the proxy-server configuration.

## TLS Client Behavior

The proxy server (as an SSL client) connects with a TLSv1 handshake. If there is a handshake failure, it falls back (for backward compatibility) to an SSLv23 handshake. (An SSL handshake differs from an SSL connection type, which is always TLSv1.) Once a handshake is successful, the proxy server verifies the authenticity of the downstream entity and tries to match the credentials presented in the certificate with the target of the URL for outgoing messages. In accordance with the SIP standard, the certificate

should prove the authenticity of the entity by matching certificate credentials—IP address, host name, FQDN, or DNS. The proxy server looks at the SubjAltName extension for these values and then, for backward compatibility, it looks at the CN (common name) field.

### TLS Server Behavior

On the server side, Cisco SPS accepts TLSv1 connections and can handshake in TLSv1 as well as SSLv23 mode. You can configure the proxy server to perform mutual TLS authentication, and hence challenge the client to prove its authenticity. It then extracts certificate credentials (IP, host name, FQDN, or DNS) and matches them with the Via header in the SIP message.

## Additional References

For additional information related to Cisco SPS, look at the material listed in the following sections.

### Related Documents

Related Topic	Document Title
Cisco SIP proxy server installation and configuration	<ul style="list-style-type: none"> <li>• <i>Cisco SIP Proxy Server Administrator Guide</i>, Version 2.1 at <a href="http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html">http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html</a></li> <li>• <i>Cisco SIP Proxy Server Installation Guide</i>, Version 2.1 at <a href="http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html">http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html</a></li> <li>• <i>Release Notes for the Cisco SIP Proxy Server Version 2.1</i> at <a href="http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html">http://www.cisco.com/en/US/products/sw/voicesw/ps2157/prod_technical_documentation.html</a></li> </ul>
Technologies referenced within Cisco SPS software and documentation	<ul style="list-style-type: none"> <li>• ADAPTIVE Communication Environment at <a href="http://www.cs.wustl.edu/~schmidt/ACE.html">http://www.cs.wustl.edu/~schmidt/ACE.html</a></li> <li>• Apache Software Foundation at <a href="http://www.apache.org">http://www.apache.org</a></li> <li>• Linux FreeS/WAN at <a href="http://www.freeswan.org/">http://www.freeswan.org/</a></li> <li>• Linux Online at <a href="http://www.linux.org">http://www.linux.org</a></li> <li>• MySQL at <a href="http://www.mysql.com">http://www.mysql.com</a></li> <li>• Red Hat at <a href="http://www.redhat.com">http://www.redhat.com</a></li> <li>• Solaris Data Encryption website: <a href="http://www.sun.com/software/solaris/encryption/">http://www.sun.com/software/solaris/encryption/</a></li> <li>• Sun Microsystems at <a href="http://www.sun.com">http://www.sun.com</a></li> <li>• Sun ONE Software at <a href="http://www.sun.com">http://www.sun.com</a></li> </ul>

## Standards

Cisco SPS supports the following standards and protocols.

Standard or Protocol <sup>1</sup>	Function
Domain Name System (DNS)	Governs translation of names of network nodes into addresses. SIP uses DNS to translate endpoint host names into IP addresses. Translation can include DNS SRV and DNS A records.
Gatekeeper Transaction Message Protocol (GKTMP)	Governs definition of the messages that are used for communication between a Cisco IOS gatekeeper and an external application.
Internet Protocol (IP)	Governs sending of datagram packets between nodes on the Internet. Also provides features for addressing, type-of-service (ToS) specification, fragmentation and reassembly, and security. Cisco SPS supports IP as it is defined in RFC 791 for SIP signaling.
Session Description Protocol (SDP)	Governs multimedia sessions and their related scheduling. The Cisco SIP IP phone uses SDP for session description.
Transport Layer Security (TLS)	Governs provision of secure transactions over TCP by means of encryption.
Transport Control Protocol (TCP)	Governs transport of data packets with guaranteed delivery. Cisco SPS supports TCP as it is defined in RFC 793 for SIP signaling.
User Datagram Protocol (UDP)	Governs transport of data packets without acknowledgments or guaranteed delivery. SIP can use UDP as the underlying transport protocol. With UDP, retransmissions are used to ensure reliability. Cisco SPS supports UDP as it is defined in RFC 768 for SIP signaling.

1. Not all supported standards are listed.

## MIBs

MIBs <sup>1</sup>	MIBs Link
<ul style="list-style-type: none"> <li>• CRITAPP-MIB (critagt)—Start and stop CSPS</li> <li>• LOG-MIB (logagt)—Monitor CSPS error_log and access_log sizes</li> <li>• DISMAN-SCRIPT-MIB (smagt)—Gracefully restart CSPS</li> <li>• DISMAN-EVENT-MIB (eventagt)—Monitor CPU load</li> <li>• HOST-RESOURCES-MIB (hostagt)—Check memory size, disk space</li> <li>• RFC1213-MIB (mib2agt)—Check link up/down status</li> <li>• SYSAPPL-MIB (sappagt)—Check what applications are installed and running on the system</li> </ul>	<p>To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use the Cisco MIB locator at the following URL:</p> <p><a href="http://www.cisco.com/go/mibs">http://www.cisco.com/go/mibs</a></p>

1. Not all supported MIBs are listed.

## RFCs

RFCs <sup>1</sup>	Title
2543	<i>SIP</i>
3261	<i>SIP: Session Initiation Protocol</i>
3263	<i>SIP: Locating SIP Servers</i>

1. Not all supported RFCs are listed.

## Technical Assistance

Description	Link
<p>Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.</p>	<p><a href="http://www.cisco.com/public/support/tac/home.shtml">http://www.cisco.com/public/support/tac/home.shtml</a></p>