

Internet Draft
Expiration: May 1999
File: [draft-ietf-rap-cops-04.txt](#)

Jim Boyle
Level 3
Ron Cohen
Cisco
David Durham
Intel
Shai Herzog
IPHighway
Raju Rajan
IBM
Arun Sastry
Cisco

The COPS (Common Open Policy Service) Protocol

Last Updated: December 18, 1998

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

To learn the current status of any Internet-Draft, please check the `1id-abstracts.txt` listing contained in the Internet-Drafts Shadow Directories on `ftp.ietf.org`, `nic.nordu.net`, `ftp.isi.edu`, or `munni.oz.au`.

A revised version of this draft document will be submitted to the RFC editor as a Proposed Standard for the Internet Community. Discussion and suggestions for improvement are requested. This document will expire before May 1999. Distribution of this draft is unlimited.

Status of this Memo.....	1
Abstract.....	3
1 . Introduction.....	3
1.1 Basic Model.....	4
2 . The Protocol.....	7
2.1 Common Header.....	7
2.2 COPS Specific Object Formats.....	8
2.2.1 Handle Object (Handle).....	9
2.2.2 Context Object (Context).....	9
2.2.3 In-Interface Object (IN-Int).....	10
2.2.4 Out-Interface Object (OUT-Int).....	11
2.2.5 Reason Object (Reason).....	11
2.2.6 Decision Object (Decision).....	12
2.2.7 LDP Decision Object (LDPDecision).....	14
2.2.8 Error Object (Error).....	14
2.2.9 Client Specific Information Object (ClientSI).....	15
2.2.10 Keep-Alive Timer Object (KATimer).....	15
2.2.11 PEP Identification Object (PEPID).....	15
2.2.12 Report-Type Object (Report-Type).....	16
2.2.13 PDP Redirect Address (PDPRedirAddr).....	16
2.2.14 Last PDP Address (LastPDPAddr).....	17
2.2.15 Accounting Timer Object (AcctTimer).....	17
2.3 Communication.....	17
2.4 Client Handle Usage.....	18
2.5 Synchronization Behavior.....	19
3 . Message Content.....	20
3.1 Request (REQ) PEP -> PDP.....	20
3.2 Decision (DEC) PDP -> PEP.....	21
3.3 Report State (RPT) PEP -> PDP.....	22
3.4 Delete Request State (DRQ) PEP -> PDP.....	22
3.5 Synchronize State Request (SSQ) PDP -> PEP.....	23
3.6 Client-Open (OPN) PEP -> PDP.....	23
3.7 Client-Accept (CAT) PDP -> PEP.....	24
3.8 Keep-Alive (KA) PEP -> PDP, PDP -> PEP.....	24
3.9 Client-Close (CC) PEP -> PDP, PDP -> PEP.....	25
3.10 Synchronize State Complete (SSC) PEP -> PDP.....	25
4 . Common Operation.....	26
4.1 PEP Initialization.....	26
4.2 Outsourcing Operations.....	26
4.3 Configuration Operations.....	27
4.4 Keep-Alive Operations.....	27
4.5 PEP/PDP Close.....	27
5 . Security.....	28
6 . IANA Considerations.....	29
7 . References.....	30
8 . Author Information and Acknowledgments.....	31

Abstract

This document describes a simple client/server model for supporting policy control over QoS Signaling Protocols and provisioned QoS resource management. It is designed to be extensible so that other kinds of policy clients may be supported in the future. The model does not make any assumptions about the methods of the policy server, but is based on the server returning decisions to policy requests.

1. Introduction

This document describes a simple query and response protocol that can be used to exchange policy information between a policy server (Policy Decision Point or PDP) and its clients (Policy Enforcement Points or PEPs). One example of a policy client is RSVP routers that must exercise policy-based admission control over RSVP usage [[RSVP](#)]. We assume that at least one policy server exists in each controlled administrative domain. The basic model of interaction between a policy server and its clients is compatible with the framework document for policy based admission control [[WRK](#)].

A chief objective of policy control protocol is to begin with a simple but extensible design. The main characteristics of the COPS protocol include:

1. The protocol employs a client/server model where the PEP sends requests, updates, and deletes to the remote PDP and the PDP returns decisions back to the PEP.
2. The protocol uses TCP as its transport protocol for reliable exchange of messages between policy clients and a server. Therefore, no additional mechanisms are necessary for reliable communication between a server and its clients.
3. The protocol is extensible in that it is designed to leverage off self-identifying objects and can support diverse client specific information without requiring modifications to the COPS protocol itself. The protocol was created for the general administration, configuration, and enforcement of policies whether signaled or provisioned. The protocol may be extended for the administration of a variety of signaling protocols as well as policy configuration on a device.

4. The protocol relies on existing protocols for security.
Namely IPSEC [[IPSEC](#)] can be used to authenticate and secure the
channel between the PEP and the server.

5. The protocol is stateful in two main aspects:

(1) Request/Decision state is shared between client and server and (2) State from various events (Request/Decision pairs) may be inter-associated. By (1) we mean that requests from the client PEP are installed or remembered by the remote PDP until they are explicitly deleted by the PEP. At the same time, Decisions from the remote PDP can be generated asynchronously at any time for a currently installed request state. By (2) we mean that the server may respond to new queries differently because of previously installed Request/Decision state(s) that are related.

6. Additionally, the protocol is stateful in that it allows the server to push configuration information to the client, and then allows the server to remove such state from the client when it is no longer applicable.

1.1 Basic Model

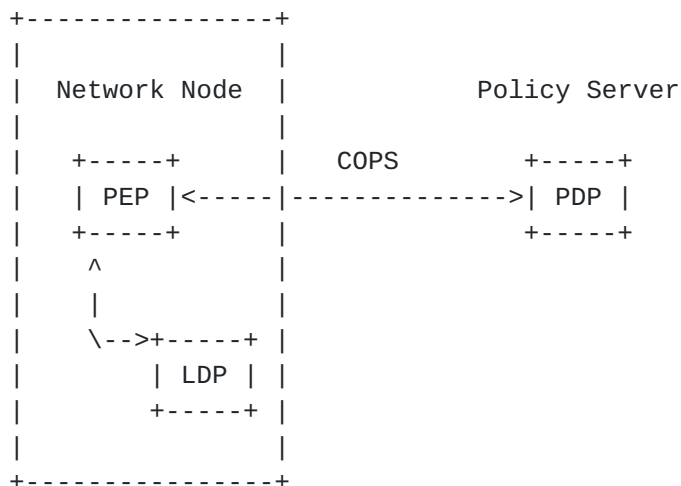


Figure 1: A COPS illustration.

Figure 1 Illustrates the layout of various policy components in a typical COPS example (taken from [WRK]). Here, COPS is used to communicate policy information between a Policy Enforcement Point (PEP) and a remote Policy Decision Point (PDP) within the context of a particular type of client.

It is assumed that each participating policy client is functionally consistent with a PEP [WRK]. The PEP may communicate with a policy server (herein referred to as a remote PDP [WRK]) to obtain policy decisions or directives.

The PEP is responsible for initiating a persistent TCP connection to

a PDP. The PEP uses this TCP connection to send requests to and receive decisions from the remote PDP. Communication between the PEP and remote PDP is mainly in the form of a stateful request/decision exchange, though the remote PDP may occasionally send unsolicited

decisions to the PEP to force changes in previously approved request states. The PEP also has the capacity to report to the remote PDP that it has committed to an accepted request state for purposes of accounting and monitoring. The PEP is responsible for notifying the PDP when a request state has changed on the PEP. Finally, the PEP is responsible for the deletion of any state that is no longer applicable due to events at the client or decisions issued by the server.

When the PEP sends a configuration request, it expects the PDP to continuously send named units of configuration data to the PEP via decision messages as applicable for the configuration request. When a unit of named configuration data is successfully installed on the PEP, the PEP should send a report message to the PDP confirming the installation. The server may then update or remove the named configuration information via a new decision message. When the PDP sends a decision to remove named configuration data from the PEP, the PEP will delete the specified configuration and send a report message to the PDP as confirmation.

The policy protocol is designed to communicate self-identifying objects which contain the data necessary for identifying request states, establishing the context for a request, identifying the type of request, referencing previously installed requests, relaying policy decisions, reporting errors, and transferring client specific/name space information.

To distinguish between different kinds of clients, the type of client is identified in each message. Different types of clients may have different client specific data and may require different kinds of policy decisions. It is expected that each new client-type will have a corresponding usage draft specifying the specifics of its interaction with this policy protocol.

The context of each request corresponds to the type of event that triggered it. COPS identifies three types of outsourcing events: (1) the arrival of an incoming message (2) allocation of local resources, and (3) the forwarding of an outgoing message. Each of these events may require different decisions to be made. Context sub types are also available to describe the type of message that triggered the policy event. The content of a COPS request/decision message depends on the context. A fourth type of request is useful for types of clients that wish to receive configuration information from the PDP. This allows a PEP to issue a configuration request for a specific named device or module that requires configuration information to be installed.

The PEP may also have the capability to make a local policy decision

via its Local Decision Point (LDP) [\[WRK\]](#), however, the PDP remains the authoritative decision point at all times. This means that the relevant local decision information must be relayed to the PDP. That is, the PDP must be granted access to all relevant information to make a final policy decision. To facilitate this functionality, the

PEP must send its local decision information to the remote PDP via a LDP decision object. The PEP must then abide by the PDP's decision as it is absolute.

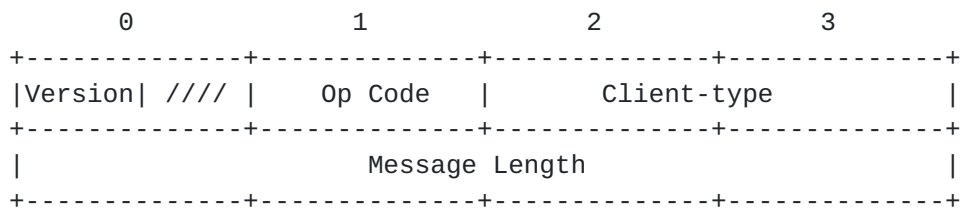
Finally, fault tolerance is a required capability for this protocol, particularly due to the fact it is associated with the security and service management of distributed network devices. Fault tolerance can be achieved by having both the PEP and remote PDP constantly verify their connection to each other via keep-alive messages. When a failure is detected, the PEP must try to reconnect to the remote PDP or attempt to connect to a new/alternative PDP. While disconnected, the PEP should revert to making local decisions. Once a connection is reestablished, the PEP is expected to notify the PDP of any deleted state or new events that passed local admission control after the connection was lost. Additionally, the remote PDP may request that all the PEP's internal state be resynchronized (all previously installed requests are to be reissued). After failure and before the new connection is fully functional, disruption of service can be minimized if the PEP caches previously communicated decisions and continues to use them for some limited amount of time.

2. The Protocol

This section describes the message formats and objects exchanged between the PEP and remote PDP.

2.1 Common Header

Each COPS message consists of the COPS header followed by a number of typed objects.



Global note: //// implies field is reserved, set to 0.

The fields in the header are:

Version: 4 bits

COPS version number. Current version is 1.

Op Code: 8 bits

The COPS operations:

1 = Request	(REQ)
2 = Decision	(DEC)
4 = Report State	(RPT)
5 = Delete Request State	(DRQ)
6 = Synchronize State Req	(SSQ)
7 = Client-Open	(OPN)
8 = Client-Accept	(CAT)
9 = Keep-Alive	(KA)
10= Client-Close	(CC)
11= Synchronize Complete	(SSC)

Client-type: 16 bits

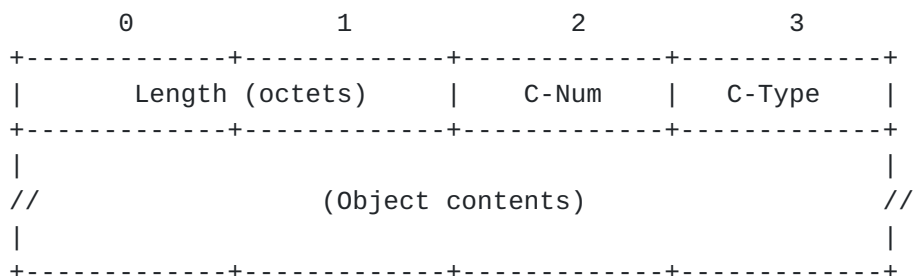
The Client-type identifies the policy client. Interpretation of all encapsulated objects is relative to the client-type. Client-types that set the most significant bit in the client-type field are enterprise specific (these are client-types 0x8000 - 0xFFFF). (See the specific client usage documents for particular client-type IDs). For KA Messages, the client-type in the header should always be set to 0 as the KA is used for connection verification (not per client session verification).

Message Length: 32 bits

Size of message in octets, which includes the standard COPS header and all encapsulated objects. Messages must be aligned on 4 octet intervals.

2.2 COPS Specific Object Formats

All the objects follow the same object format; each object consists of one or more 32-bit words with a four-octet header, using the following format:



The length is a two-octet value that describes the number of octets (including the header) that compose the object. If the length in octets does not fall on a 32-bit word boundary, padding must be added to the end of the object so that it is aligned to the next 32-bit boundary before the object can be sent on the wire. On the receiving side, a subsequent object boundary can be found by simply rounding up the previous stated object length to the next 32-bit boundary.

Typically, C-Num identifies the class of information contained in the object, and the C-Type identifies the subtype or version of the information contained in the object.

C-num: 8 bits

- 1 = Handle
- 3 = Context
- 4 = In Interface
- 5 = Out Interface
- 6 = Reason code
- 7 = Decision
- 8 = LDP Decision
- 9 = Error
- 10 = Client Specific Info
- 11 = Keep-Alive Timer
- 12 = PEP Identification
- 13 = Report Type
- 14 = PDP Redirect Address
- 15 = Last PDP Address
- 16 = Accounting Timer

C-type: 8 bits

Values defined per C-num.

2.2.1 Handle Object (Handle)

The Handle Object encapsulates a unique value that identifies an installed state. This identification is used by most COPS operations. A state corresponding to a handle must be explicitly deleted when it is no longer applicable.

C-Num = 1

C-Type = 1, Client Handle.

Variable-length field, no implied format other than it is unique from other client handles. It is always initially chosen by the PEP and then deleted by the PEP when no longer applicable. The client handle is used to refer to a request state initiated by the PEP and installed at the PDP. A PEP will specify a client handle in its Request messages, Report messages and Delete messages sent to the PDP. In all cases, the client handle is used to uniquely identify the PEP request.

The client handle value is set by the PEP and is opaque to the PDP. The PDP simply performs a byte-wise comparison on the value in this object with respect to the handle object values of other currently installed requests.

2.2.2 Context Object (Context)

Specifies the type of event(s) that triggered the query. Required for request messages. Admission control, resource allocation, and forwarding requests are all amenable to client-types that outsource their decision making facility to the PDP. For applicable client-types a PEP can also make a request to receive named configuration information from the PDP. This named configuration data may be in a form useful for setting system attributes on a PEP, or it may be in the form of policy rules that are to be directly verified by the PEP.

Multiple flags can be set for the same request. This is only allowed, however, if the set of client specific information in the combined request is identical to the client specific information that would be specified if individual requests were made for each specified flag.

C-num = 3, C-Type = 1

0	1	2	3
+-----+-----+-----+-----+			
	R-Type		M-Type
+-----+-----+-----+-----+			

R-Type (Request Type Flag)

0x01 = Incoming-Message/Admission Control request

Boyle et al.

Expires May 1999

[Page 9]

0x02 = Resource-Allocation request
 0x04 = Outgoing-Message request
 0x08 = Configuration request

M-Type (Message Type)

Client Specific 16 bit values of protocol message types

2.2.3 In-Interface Object (IN-Int)


The In-Interface Object is used to identify the incoming interface on which a particular request applies and the address where the received message originated. For flows or messages generated from the PEP's local host, the loop back address and ifindex are used.

This Interface object is also used to identify the incoming (receiving) interface via its ifindex. The ifindex may be used to differentiate between sub-interfaces and unnumbered interfaces (see RSVP's LIH for an example). When SNMP is supported by the PEP, this ifindex integer must correspond to the same integer value for the interface in the SNMP MIB-II interface index table.

Note: The ifindex specified in the In-Interface is typically relative to the flow of the underlying protocol messages. The ifindex is the interface on which the protocol message was received.

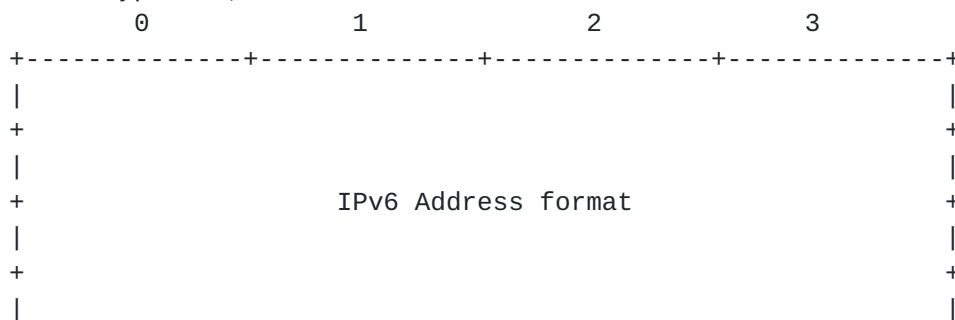
C-Num = 4

C-Type = 1, IPv4 Address + Interface



For this type of the interface object, the IPv4 address should specify the IP address that the incoming message came from.

C-Type = 2, IPv6 Address + Interface



+	-----	+	-----	+	-----	+	-----	+
+	-----	+	-----	+	-----	+	-----	+

For this type of the interface object, the IPv6 address should specify the IP address that the incoming message came from. The ifindex is used to refer to the MIB-II defined local incoming interface on the PEP as described above.

2.2.4 Out-Interface Object (OUT-Int)

The Out-Interface is used to identify the outgoing interface to which a specific request applies and the address for where the forwarded message is to be sent. For flows or messages destined to the PEP's local host, the loop back address and ifindex are used. The Out-Interface has the same formats as the In-Interface Object.

This Interface object is also used to identify the outgoing (forwarding) interface via its ifindex. The ifindex may be used to differentiate between sub-interfaces and unnumbered interfaces (see RSVP's LIH for an example). When SNMP is supported by the PEP, this ifindex integer must correspond to the same integer value for the interface in the SNMP MIB-II interface index table.

Note: The ifindex specified in the Out-Interface is typically relative to the flow of the underlying protocol messages. The ifindex is the one on which a protocol message is about to be forwarded.

C-Num = 5

C-Type = 1, IPv4 Address + Interface

Same C-Type format as the In-Interface object. The IPv4 address should specify the IP address to which the outgoing message is going. The ifindex is used to refer to the MIB-II defined local outgoing interface on the PEP.

C-Type = 2, IPv6 Address + Interface

Same C-Type format as the In-Interface object. For this type of the interface object, the IPv6 address should specify the IP address to which the outgoing message is going. The ifindex is used to refer to the MIB-II defined local outgoing interface on the PEP.

2.2.5 Reason Object (Reason)

This object specifies the reason why the request state was deleted. It should appear in the delete request (DRQ) message. The Reason Sub-code field is reserved for more detailed client-specific reason codes defined in the corresponding documents.

C-Num = 6, C-Type = 1

0	1	2	3
+-----+-----+-----+-----+			
	Reason-Code		Reason Sub-code
+-----+-----+-----+-----+			

Reason Code:

- 1 = Unspecified
- 2 = Management
- 3 = Preempted (Another request state takes precedence)
- 4 = Tear (Used to communicate a signaled state removal)
- 5 = Timeout (Local state has timed-out)
- 6 = Route Change (Change invalidates request state)
- 7 = Insufficient Resources (No local resource available)
- 8 = PDP's Directive (PDP decision caused the delete)
- 9 = Unsupported decision (PDP decision not supported)
- 10= Synchronize Handle Unknown
- 11= Transient Handle (stateless event)
- 12= Malformed Decision (could not recover)

2.2.6 Decision Object (Decision)

Decision made by the PDP. Must appear in replies. The specific non-mandatory decision objects required in a decision to a particular request depend on the type of client.

C-Num = 7

C-Type = 1, Decision Flags (Mandatory)

0	1	2	3
+-----+-----+-----+-----+			
	Flags		
+-----+-----+-----+-----+			

Flags:

- 0x01 = Admit Signaled (Reject if set)
- 0x08 = Trigger Error (Trigger error message if set)
- Note: For the above flags, a flag bit set to 1 implies a negative decision for that flag. Not setting a flag implies a positive decision.
- 0x10 = NULL Configuration (No configuration data if set)
- 0x20 = Install Configuration (Install named data if set)
- 0x40 = Remove Configuration (Remove named data if set)
- Note: If NULL Configuration ignore Install/Remove.
- Note: Only one of Install OR Remove may be set in one decision flags object.
- 0x200= Solicited Decision
- (Initial decision after a new/updated request if set)

Flag values not applicable to a given context's R-Type MUST be ignored by the PEP. See table below:

Boyle et al.

Expires May 1999

[Page 12]

Decision Flag	: Valid Context R-Type

Admit Signaled	: Incoming-Message/Admission Control
Admit Signaled	: Resource-Allocation
Admit Signaled	: Outgoing-Message
Trigger Error	: Any of the above R-Types
NULL Configuration	: Configuration
Install Config.	: Configuration
Remove Config.	: Configuration
Solicited Decision	: Any R-Type

C-Type = 2, Resource Allocation Data

This type of decision object carries additional stateless information that can be applied by the PEP locally. It is a variable length object and its internal format should be specified in the relevant COPS extension document for the given client-type. This object is optional in Decision messages and is interpreted relative to a given context.

It is expected that even outsourcing PEPs will be able to make some simple stateless policy decisions locally in their LDP. As this set is well known and implemented ubiquitously, PDPs are aware of it as well (either universally, through configuration, or using the Client-Open message). The PDP may also include this information in its decision, and the PEP should apply it to the resource allocation event that generated the request.

As an example, reservations may be admitted by a PDP contingent on some type of per-session preemption priority. A RSVP PEP could have a set of stateless policy rules for when to preempt other reservations in favor of a new one (e.g. higher-priority pre-empts any of lower priority). The PDP would need to include appropriate priority information for each reservation in its decisions that the PEP can use to apply its rules.

C-Type = 3, Replacement Data

This type of decision object carries replacement data that is to replace existing data in a signaled message. It is a variable length object and its internal format should be specified in the relevant COPS extension document for the given client-type. It is optional in Decision messages and is interpreted relative to a given context.

C-Type = 4, Client Specific Decision Data

Additional decision types can be introduced using the Client Specific Decision Data Object. It is a variable length object and its internal format should be specified in the relevant COPS extension document for the given client-type. It is optional in

Decision messages and is interpreted relative to a given context.

C-Type = 5, Named Decision Data

Named configuration information should be encapsulated in this version of the decision object in response to configuration requests. It is a variable length object and its internal format should be specified in the relevant COPS extension document for the given client-type. It is optional in Decision messages and is interpreted relative to both a given context and decision flags.

2.2.7 LDP Decision Object (LDPDecision)

Decision made by the PEP's local decision point (LDP). May appear in requests. These objects correspond to and are formatted the same as the client specific decision objects defined above.

C-Num = 8

C-Type = (same C-Type as for Decision objects)

2.2.8 Error Object (Error)

This object is used to identify a particular COPS protocol error. The error sub-code field contains additional detailed client specific error codes. The appropriate Error Sub-codes for a particular client-type should be specified in the relevant COPS extensions document.

C-Num 9, C-Type = 1

0	1	2	3
Error-Code		Error Sub-code	

Error-Code:

- 1 = Bad handle
- 2 = Invalid handle reference
- 3 = Bad message format (Malformed Message)
- 4 = Unable to process (server gives up on query)
- 5 = Mandatory client-specific info missing
- 6 = Unsupported client-type
- 7 = Mandatory COPS object missing

8 = Client Failure
9 = Communication Failure
10= Unspecified
11= Shutting down

Boyle et al.

Expires May 1999

[Page 14]

2.2.9 Client Specific Information Object (ClientSI)

The various types of this object are required for requests, and used in reports and opens when required. It contains client-type specific information.

C-Num = 10,

C-Type = 1, Signaled ClientSI.

Variable-length field. All objects/attributes specific to a client's signaling protocol or internal state must be encapsulated within one or more signaled Client Specific Information Objects. The format of the data encapsulated in the ClientSI object is determined by the client-type.

C-Type = 2, Named ClientSI.

Variable-length field. Contains named configuration information useful for relaying specific information about the PEP, a request, or configured state to the PDP server.


2.2.10 Keep-Alive Timer Object (KATimer)

Times are encoded as 2 octet integer values and are in units of seconds. The timer value is treated as a delta.

C-Num = 11,

C-Type = 1, Keep-alive timer value

Timer object used to specify the maximum time interval over which a COPS message must be sent or received. The value of zero implies infinity.



2.2.11 PEP Identification Object (PEPID)

The PEP Identification Object is used to identify the PEP client to the remote PDP. It is required for Client-Open messages.

C-Num = 12, C-Type = 1

Variable-length field. It is a NULL terminated ASCII string that is also zero padded to a 32-bit word boundary (so the object length is a multiple of 4 octets). The PEPID must contain an ASCII string that

uniquely identifies the PEP within the policy domain in a manner that is persistent across PEP reboots. For example, it may be the PEP's statically assigned IP address or DNS name. This identifier may safely be used by a PDP as a handle for identifying the PEP in its policy rules.

2.2.12 Report-Type Object (Report-Type)

The Type of Report on the request state associated with a handle:

C-Num = 13, C-Type = 1

0	1	2	3
+-----+-----+-----+-----+			
Report-Type		///////////////	
+-----+-----+-----+-----+			

Report-Type:

- 1 = Commit : PEP's local resources now allocated
- 2 = Accounting: Accounting update for an installed state
- 3 = No Commit : PEP's resource allocation failure
- 4 = Installed : Named configuration installed
- 5 = Removed : Named configuration removed
- 6 = NotInstall: Named configuration was not installed
- 7 = NotRemoved: Named configuration was not removed

2.2.13 PDP Redirect Address (PDPRedirAddr)

A PDP when closing a PEP session for a particular client-type may optionally use this object to redirect the PEP to another PDP server:

C-Num = 14,

C-Type = 1, IPv4 Address (4 octets)

0	1	2	3
+-----+-----+-----+-----+			
IPv4 Address format			
+-----+-----+-----+-----+			

C-Type = 2, IPv6 Address (16 octets)

0	1	2	3
+-----+-----+-----+-----+			
+			+
+	IPv6 Address format		+

+ +
| |
+-----+-----+-----+-----+

2.2.14 Last PDP Address (LastPDPAddr)

When a PEP sends a Client-Open message for a particular client-type the PEP should specify the last PDP it has successfully opened (meaning it received a Client-Accept) since the PEP last rebooted. If no PDP was used since the last reboot, the PEP will simply not include this object in the Client-Open message.

C-Num = 15,

C-Type = 1, IPv4 Address (Same format as PDPRedirAddr)

C-Type = 2, IPv6 Address (Same format as PDPRedirAddr)


2.2.15 Accounting Timer Object (AcctTimer)

Times are encoded as 2 octet integer values and are in units of seconds. The timer value is treated as a delta.

C-Num = 16,

C-Type = 1, Accounting timer value

Optional timer value used to determine the minimum interval between periodic accounting type reports. It is used by the PDP to describe to the PEP an acceptable interval between accounting updates via Report messages where applicable. The value of zero implies there are no timing constraints on accounting updates.



2.3 Communication

The COPS protocol uses a single persistent TCP connection between the PEP and a remote PDP. The remote PDP listens on a well-known port number (COPS=3288 [[IANA](#)]), and the PEP is responsible for initiating the connection. The location of the remote PDP can either be configured, or obtained via a service location mechanism [SRVLOC]. Service discovery is outside the scope of this protocol, however.

If a single PEP can support multiple client-types, it may send multiple Client-Open messages, each specifying a particular client-

type to a PDP over one or more TCP connections. Likewise, a PDP residing at a given address may support one or more client-types. Given the client-types it supports, a PDP has the ability to either accept or reject each client-type independently. If a client-type is

rejected, the PDP can redirect the PEP to an alternative PDP for a given client-type via COPS. Additional provisions for supporting multiple client-types (perhaps from independent PDP vendors) on a single remote PDP server are not provided by the COPS protocol, but, rather, are left to the software architecture of the given server platform.

It is possible a single PEP may have open connections to multiple PDPs. This is the case when there are physically different PDPs supporting different client-types as shown in figure 2.

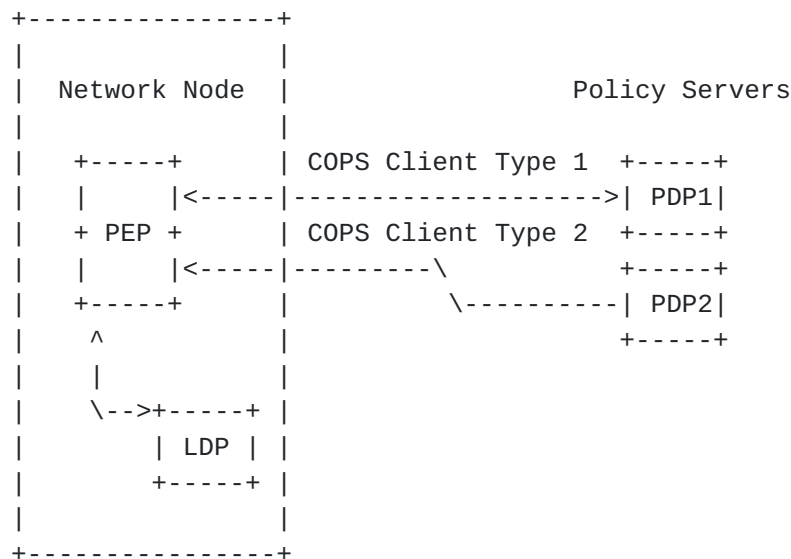


Figure 2: Multiple PDPs illustration.

When a TCP connection is torn down or is lost, the PDP is expected to eventually clean up any outstanding request state related to request/decision exchanges with the PEP. When the PEP detects a lost connection due to a timeout condition it should explicitly send a Client-Close message for each opened client-type containing an <Error> object indicating the "Communication Failure" Error-Code. Additionally, the PEP should continuously attempt to contact the primary PDP or, if unsuccessful, any known backup PDPs. Specifically the PEP should keep trying all relevant PDPs with which it has been configured until it can establish a connection. If a PEP is in communication with a backup PDP and the primary PDP becomes available, the backup PDP is responsible for redirecting the PEP back to the primary PDP (via a <Client-Close> message containing a <PDPRedirAddr> object indicating the primary PDP to use for each affected client-type).

2.4 Client Handle Usage

The client handle is used to identify a unique request state. Client handles are chosen by the PEP and are opaque to the PDP. The PDP simply uses the request handle to uniquely identify the request state and generically tie its decisions to a corresponding request.

Client handles are initiated in request messages and are then used by subsequent request, decision, and report messages to reference the same request state. When the PEP is ready to remove a local request state, it will issue a delete message to the PDP for the corresponding client handle. A handle **MUST** be explicitly deleted by the PEP before it can be used to identify a new request state. Handles referring to different request states must be unique.

2.5 Synchronization Behavior

When disconnected from a PDP, the PEP should revert to making local decisions. Once a connection is reestablished, the PEP is expected to notify the PDP of any events that have passed local admission control. Additionally, the remote PDP may request that all the PEP's internal state be resynchronized (all previously installed requests are to be reissued) by sending a Synchronize State message.

After a failure and before a new connection is fully functional, disruption of service can be minimized if the PEP caches previously communicated decisions and continues to use them for some appropriate length of time. Specific rules for such behavior are to be defined in the appropriate COPS client-type extension specifications.

A PEP that caches state from a previous exchange with a disconnected PDP must communicate this fact to any PDP with which it is able to later reconnect. This is accomplished by including the address of the last PDP for which the PEP is still caching state in the Client-Open message. The <LastPDPAddr> object will only be included for the last PDP with which the PEP was completely in sync. If the service interruption was temporary and the PDP still contains the complete state for the PEP, the PDP may choose not to synchronize all states. It is still the responsibility of the PEP to update the PDP of all state changes that occurred during the disruption of service including any states communicated to the previous PDP that had been deleted after the connection was lost. These must be explicitly deleted after a connection is reestablished. If the PDP issues a synchronize request the PEP must pass all current states to the PDP followed by a Synchronize State Complete message (thus completing the synchronization process).

3. Message Content

This section describes the basic messages exchanged between a PEP and a remote PDP as well as their contents. As a convention, object ordering is expected as shown in the BNF for each COPS message unless otherwise noted. Malformed messages are to be silently dropped unless otherwise specified.

3.1 Request (REQ) PEP -> PDP

The PEP establishes a request state client handle for which the remote PDP may maintain a state. The remote PDP then uses this handle to refer to the exchanged information and decisions.

Once a stateful handle is established for a new request, any subsequent modifications of the request can be made using the REQ message specifying the previously installed handle. The PEP is responsible for notifying the PDP whenever its local state changes so the PDP's state will be able to accurately mirror the PEP's state.

The format of the Request message is as follows:

```

<Request Message> ::= <Common Header>
                        <Client Handle>
                        <Context>
                        [<IN-Int>]
                        [<OUT-Int>]
                        [<ClientSI(s)>]
                        [<LDPDecision(s)>]

<ClientSI(s)> ::= <ClientSI> | <ClientSI(s)> <ClientSI>

<LDPDecision(s)> ::= <LDPDecision> |
                    <LDPDecision(s)> <LDPDecision>

```

The context object is used to determine the context within which all the other objects are to be interpreted. It also is used to determine the kind of decision to be returned from the policy server. This decision might be related to admission control, resource allocation, object forwarding and substitution, or configuration.

The interface objects are used to determine the corresponding interface on which a signaling protocol message was received or is about to be sent. They are typically used if the client is participating along the path of a signaling protocol or if the

client is requesting configuration data for a particular interface.

ClientSI, the client specific information object, holds the client-type specific data for which a policy decision needs to be made. In

the case of configuration, the Named ClientSI may include named information about the module, interface, or functionality to be configured. The ordering of multiple ClientSIs is not important.

Finally, LDPDecision object holds information regarding the local decision made by the LDP.

3.2 Decision (DEC) PDP -> PEP

The PDP responds to the REQ with a DEC message that includes the associated client handle and one or more decision objects grouped relative to a Context object and Decision Flags object type pair. If there was a protocol error an error object is returned instead.

It is assumed that the first decision for a new/updated request will set the solicited decision flag in all included Decision Flag object types. This avoids the issue of keeping track of which updated request (that is, a request reissued for the same handle) a particular decision corresponds. It is important that, for a given handle, there be at most one outstanding solicited decision per request. This essentially means that the PEP should not issue more than one REQ (for a given handle) before it receives a corresponding DEC with the solicited decision flag set.

To avoid deadlock, the client can always timeout after issuing a request. It must then delete the timed-out handle, and possibly try again using a different (new) handle.

The format of the Decision message is as follows:

```
<Decision Message> ::= <Common Header>
                        <Client Handle>
                        <Decision(s)> | <Error>

<Decision(s)> ::= <Decision> | <Decision(s)> <Decision>

<Decision> ::= <Context>
               <Decision: Flags>
               [<Decision: Resource Alloc Data>]
               [<Decision: Replacement Data>]
               [<Decision: ClientSI Data>]
               [<Decision: Named Data>]
```

The Decision message may include either an Error object or one or more context plus associated decision objects. COPS protocol problems are reported in the Error object (e.g. an error with the format of the original request, including malformed request messages). The applicable Decision object(s) depend on the context

and the type of client. The only ordering requirement for decision objects is that the required Decision Flags object type must proceed the other Decision object types per context binding.

3.3 Report State (RPT) PEP -> PDP

This message is used by the PEP to communicate a change in the status of a previously installed state to the PDP. A commit or no-commit report-type indicates to the PDP that a particular policy directive has or has not been acted upon as is relevant for accounting purposes. (In RSVP this would mean that a reservation passed or failed local capacity admission control [[RSVP](#)]. For a configuration decision, it would mean the configuration identified in the ClientSI either could or could not be installed by the PEP).

The Report State may also be used to provide periodic updates of client specific information for accounting and state monitoring purposes depending on the type of the client. In such cases the accounting report type should be specified utilizing the appropriate client specific information object.

```
<Report State> ::= <Common Header>
                  <Client Handle>
                  <Report-Type>
                  [<ClientSI>]
```

3.4 Delete Request State (DRQ) PEP -> PDP

When sent from the PEP this message indicates to the remote PDP that the state identified by the client handle is no longer available/relevant. This information will then be used by the remote PDP to initiate the appropriate housekeeping actions. The reason code object is interpreted with respect to the client-type and signifies the reason for the removal.

The format of the Delete Request State message is as follows:

```
<Delete Request> ::= <Common Header>
                    <Client Handle>
                    <Reason>
```

Given the stateful nature of COPS, it is important that when a request state is finally removed from the PEP, a DRQ message for this request state is sent to the PDP so the corresponding state may likewise be removed on the PDP. Request states not explicitly deleted by the PEP will be maintained by the PDP until either the client session is closed or the connection is terminated.

Malformed Decision messages should trigger a DRQ specifying the appropriate erroneous reason code (Bad Message Format) and any associated state on the PEP should either be removed or re-requested.

3.5 Synchronize State Request (SSQ) PDP -> PEP

The format of the Synchronize State Query message is as follows:

```
<Synchronize State> ::= <Common Header>
                        [<Client Handle>]
```

This message indicates that the remote PDP wishes the client (which appears in the common header) to re-send its state. If the optional Client Handle is present, only the state associated with this handle is synchronized. If the PEP does not recognize the requested handle, it should immediately send a DRQ message to the PDP for the handle that was specified in the SSQ message. If no handle is specified in the SSQ message, all the active client state should be synchronized with the PDP.

The client performs state synchronization by re-issuing request queries of the specified client-type for the existing state in the PEP. When synchronization is complete, the PEP must issue a synchronize state complete message to the PDP.

3.6 Client-Open (OPN) PEP -> PDP

The Client-Open message can be used by the PEP to specify to the PDP the client-types the PEP can support, the last PDP to which the PEP connected for the given client-type, and/or client specific feature negotiation. A Client-Open message can be sent to the PDP at any time and multiple Client-Open messages for the same client-type are allowed (in case of global state changes).

```
<Client-Open> ::= <Common Header>
                  <PEPID>
                  [<ClientSI>]
                  [<LastPDPAddr>]
```

The PEPID is a symbolic, variable length name that uniquely identifies the specific client to the PDP.

A named ClientSI object can be included for relaying additional global information about the PEP to the PDP when required (as specified in the appropriate extensions document for the client-type).

Finally, the PEP may provide a Last PDP Address object in its Client-Open message specifying the last PDP (for the given client-type) for which it is still caching decisions since its last reboot. A PDP can use this information to determine the appropriate

synchronization behavior (See [section 2.5](#)).

3.7 Client-Accept (CAT) PDP -> PEP

The Client-Accept message is used to positively respond to the Client-Open message. This message will return to the PEP a timer object indicating the maximum time interval between keep-alive messages. Optionally, a timer specifying the minimum allowed interval between accounting report messages may be included when applicable.

```
<Client-Accept> ::= <Common Header>
                    <KA Timer>
                    [<ACCT Timer>]
```

If the PDP refuses the client, it will instead issue a Client-Close message.

The KA Timer corresponds to maximum acceptable intermediate time between the generation of messages by the PDP and PEP. The timer value is determined by the PDP and is specified in seconds. A timer value of 0 implies no secondary connection verification is necessary.

The optional accounting timer allows the PDP to indicate to the PEP that periodic accounting reports should not exceed the specified timer interval. This allows the PDP to control the rate at which accounting reports are sent by the PEP (when applicable). In general, accounting type Report messages are sent to the PDP when determined appropriate by the PEP. The accounting timer merely is used by the PDP to keep the rate of such updates in check (i.e. Preventing the PEP from blasting the PDP with accounting reports).

3.8 Keep-Alive (KA) PEP -> PDP, PDP -> PEP

The keep-alive message must be transmitted by the PEP within the period defined by the minimum of all KA Timer values specified in all received CAT messages for the connection. A KA message must be generated randomly between 1/4 and 3/4 of this minimum TA timer interval. When the PDP receives a keep-alive message from a PEP, it must echo a keep-alive back to the PEP. This message provides validation for each side that the connection is still functioning even when there is no other messaging.

Note: The client-type in the header should always be set to 0 as the KA is used for connection verification (not per client session verification).

```
<Keep-Alive> ::= <Common Header>
```

Both client and server may assume the TCP connection is insufficient for the client-type with the minimum time value (specified in the CAT message) if no communication activity is detected for a period exceeding the timer period. For the PEP, such detection implies the

remote PDP or connection is down and the PEP should now attempt to use an alternative/backup PDP.

3.9 Client-Close (CC) PEP -> PDP, PDP -> PEP

The Client-Close message can be issued by either the PDP or PEP to notify the other that a particular type of client is no longer being supported.

```
<Client-Close> ::= <Common Header>
                  <Error>
                  [<PDPRedirAddr>]
```

The Error object is included to describe the reason for the close (e.g. the requested client-type is not supported by the remote PDP or client failure).

A PDP may optionally include a PDP Redirect Address object in order to inform the PEP of the alternate PDP it should use for the client-type specified in the common header.

3.10 Synchronize State Complete (SSC) PEP -> PDP

The Synchronize State Complete is sent by the PEP to the PDP after the PDP sends a synchronize state request to the PEP and the PEP has finished synchronization. It is useful so that the PDP will know when all the old client state has been successfully re-requested and, thus, the PEP and PDP are completely synchronized.

```
<Synchronize State Complete> ::= <Common Header>
```


4. Common Operation

This section describes the typical exchanges between remote PDP servers and PEP clients.

4.1 PEP Initialization

Sometime after a connection is established between the PEP and a remote PDP, the PEP will send one or more Client-Open messages to the remote PDP, one for each client-type supported by the PEP. The Client-Open message must contain the address of the last PDP with which the PEP is still caching a complete set of decisions. If no decisions are being cached from the previous PDP the LastPDPAddr object must not be included in the Client-Open message (see [Section 2.5](#)). Each Client-Open message should at least contain the common header noting one client-type supported by the PEP. The remote PDP will then respond with separate Client-Accept messages for each of the client-types requested by the PEP that the PDP can also support.

If a specific client-type is not supported by the PDP, the PDP will instead respond with a Client-Close specifying the client-type is not supported and will possibly suggest an alternate PDP address. Otherwise, the PDP will send a Client-Accept specifying the timer interval between keep-alive messages and the PEP may begin issuing requests to the PDP.

4.2 Outsourcing Operations

In the outsourcing scenario, when the PEP receives an event that requires a new policy decision it sends a request message to the remote PDP. What specifically qualifies as an event for a particular client-type should be specified in the specific document for that client-type. The remote PDP then makes a decision and sends a decision message back to the PEP. Since the request is stateful, the request will be remembered, or installed, on the remote PDP. The unique handle, specified in both the request and its corresponding decision identifies this request state. The PEP is responsible for deleting this request state once the request is no longer applicable.

The PEP can update a previously installed request state by reissuing a request for the previously installed handle. The remote PDP is then expected to make new decisions and send a decision message back to the PEP. Likewise, the server may change a previously issued decision on any currently installed request state at any time by issuing another decision message. At all times the PEP module is

expected to abide by the PDP's decisions and notify the PDP of any state changes.

4.3 Configuration Operations

In the configuration scenario, as in the outsourcing scenario, the PEP will make a configuration request to the PDP for a particular interface, module, or functionality that may be specified in the named client specific information object. The PDP will then send potentially several decisions containing named units of configuration data to the PEP. The PEP is expected to install and use the configuration locally. A particular named configuration can be updated by simply sending additional decision messages for the same named configuration. When the PDP no longer wishes the PEP to use a piece of configuration information, it will send a decision message specifying the named configuration and a decision flags object with the remove configuration flag set. The PEP should then proceed to remove the corresponding configuration and send a report message to the PDP that specifies it has been deleted.

In all cases, the PEP may notify the remote PDP of the local status of an installed state using the report message where appropriate. The report message is to be used to signify when billing should begin, what actions were taken, or to produce periodic updates for monitoring and accounting purposes depending on the client. This message can carry client specific information when needed.

4.4 Keep-Alive Operations

The Keep-Alive message is used to validate the connection between the client and server is still functioning even when there is no other messaging from the PEP to PDP. The PEP must generate a COPS KA message randomly within one-fourth to three-fourths the negotiated minimum KA Timer interval. On receiving a Keep-Alive message from the PEP, the PDP must then respond to this Keep-Alive message by echoing a Keep-Alive message back to the PEP. If either side does not receive a Keep-Alive or any other COPS message within the minimum KA Timer interval from the other, the connection should be considered lost.

4.5 PEP/PDP Close

Finally, Client-Close messages are used to negate the effects of the corresponding Client-Open messages, notifying the other side that the specified client-type is no longer supported/active. When the PEP detects a lost connection due to a keep-alive timeout condition it should explicitly send a Client-Close message for each opened client-type specifying a communications failure error code. Then the PEP may proceed to terminate the connection to the PDP and attempt to reconnect again or try a backup/alternative PDP. When the PDP is shutting down, it should also explicitly send a Client-Close to all

connected PEPs for each client-type, perhaps specifying an alternative PDP to use instead.

5. Security

The security of RSVP messages is provided by inter-router MD5 authentication [[MD5](#)]. This assumes a chain-of-trust model for inter PEP authentication. Security between the client (PEP) and server (PDP) is provided by IPSEC [[IPSEC](#)].

To ensure the client (PEP) is communicating with the correct policy server (PDP) involves two issues: authentication of the policy client and server using a shared secret, and consistent proof that the connection remains valid. The shared secret requires manual configuration of keys, which is a maintenance issue. IPSEC AH may be used for the validation of the connection; IPSEC ESP may be used to provide both validation and secrecy.

6. IANA Considerations

The Client-type identifies the policy client application to which a message refers. Client-type values within the range 0x0000-0x3FFF are reserved Specification Required status as defined in [IANA-CONSIDERATIONS]. These values must be registered with IANA and their behavior and applicability must be described in a COPS extension document.

Client-type values in the range 0x4000 - 0x7FFF are reserved for Private Use as defined in [[IANA-CONSIDERATIONS](#)]. These Client-types are not tracked by IANA and are not to be used in standards or general-release products, as their uniqueness cannot be assured.

Client-type values in the range 0x8000 - 0xFFFF are First Come First Served as defined in [[IANA-CONSIDERATIONS](#)]. These Client-types are tracked by IANA but do not require published documents describing their use. IANA merely assures their uniqueness.

7. References

- [RSVP] Braden, R. ed. et al., "Resource ReSerVation Protocol (RSVP) Version 1 - Functional Specification", [RFC 2205](#), September 1997.
- [WRK] Yavatkar, R. et al., "A Framework for Policy-Based Admission Control", Internet-Draft, [draft-ietf-rap-framework-01.txt](#), November 1998.
- [SRVLOC] Guttman, E. et al., "Service Location Protocol", Internet-Draft, [draft-ietf-srvloc-protocol-v2-01.txt](#), October 1997.
- [INSCH] Shenker, S., Wroclawski, J., "General Characterization Parameters for Integrated Service Network Elements", [RFC 2215](#), September 1997.
- [IPSEC] Atkinson, R., "Security Architecture for the Internet Protocol", [RFC1825](#), August 1995.
- [MD5] Baker, F., "RSVP Cryptographic Authentication", Internet-Draft, [draft-ietf-rsvp-md5-05.txt](#), August 1997.
- [RSVPPR] Braden, R., Zhang, L., "Resource ReSerVation Protocol (RSVP) - Version 1 Message Processing Rules", [RFC 2209](#), September 1997.
- [IANA] <http://www.isi.edu/in-notes/iana/assignments/port-numbers>
- [IANA-CONSIDERATIONS] Alvestrand, H. and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.

8. Author Information and Acknowledgments

Special thanks to Andrew Smith and Timothy O'Malley our WG Chairs, Raj Yavatkar, Russell Fenger, Fred Baker, Laura Cunningham, Roch Guerin, Ping Pan, and Dimitrios Pendarakis for their valuable contributions.

Jim Boyle
Level 3 Communications
1450 Infinite Drive13
Louisville, CO 80027
303.926.3100
email: jboyle@l3.net

Ron Cohen
Cisco Systems
Hasadna St.
Ra'anana 43650 Israel
972.9.7462020
ronc@classdata.com

David Durham
Intel
2111 NE 25th Avenue
Hillsboro, OR 97124
503.264.6232
David_Durham@mail.intel.com

Raju Rajan
IBM T.J. Watson Research Cntr
P.O. Box 704
Yorktown Heights, NY 10598
914.784.7260
raju@watson.ibm.com

Shai Herzog
IPHighway
2055 Gateway Pl., Suite 400
San Jose, CA 95110
408.390.3045
herzog@iphighway.com

Arun Sastry
Cisco Systems
506210 W Tasman Drive
San Jose, CA 95134
408.526.7685
asastry@cisco.com

