

Internet Draft

Document: [draft-gibson-sip-qos-resv-00.txt](#)

Category: Informational

M. Gibson, J. Crowcroft

Nortel Networks, UCL

October 1999

Use of SIP for the Reservation of QoS guaranteed Paths

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at

<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

1. Abstract

This draft defines a modification of the use of the SIP protocol to perform route reservation. These modifications are applied to two of the existing SIP Methods: INVITE and REGISTER. The modifications allow the INVITE method to be used to find the best path across a particular network based on QoS metrics. The route choice is able to be restricted at the network edge. The proposed modifications also allow the REGISTER method to be used as a means of disseminating information necessary to determine these QoS metrics. This information can be used to route INVITES away from known areas of congestion.

2. Conventions used in this document

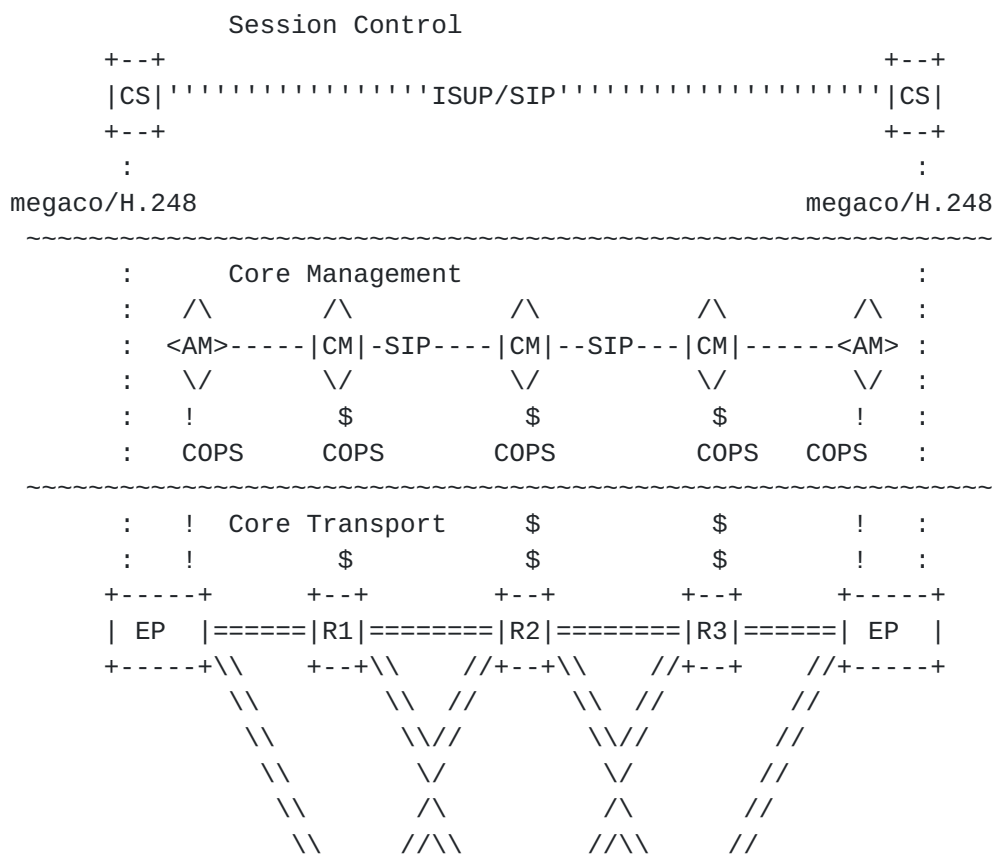
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [2].

3. Protocol Context

The purpose of this draft is to propose that, with suitable extensions, the SIP protocol may be used to perform path determination and reservation across a QoS-capable IP network.

This draft defines a set of extended uses of the basic SIP protocol the motivation for which can be found in [3], however this modified functionality has a wider applicability. The purpose of the network is to enable guaranteed QoS session establishment across an IP trunk network between two separate LANs. The framework document describes solution that can be decomposed into a three-layer model that is repeated here as Figure 1.

The top Session Layer represents a telephony style connection model with user connection stimulus being processed by a Call Server (CS). These CSs may interact using either of ISUP [4] or SIP [5]. The use of this layer is application specific being used by those applications that generate telephony style signaling.






Figure 1 Extended SIP reference architecture

Gibson Category Informational - Expiration April 1999 2

SIP Reservation Extensions October 1999

The Session Layer connects directly to the Transport layer by means of the emerging megaco protocol. This protocol is used by the CS to instruct its endpoint (EP) to establish a new connection for a particular session. In this mode of operation, the EP therefore resembles a Media Gateway.

In those applications where a Session Layer is not present, the EP will receive connection requests directly from the end-user, for example a RSVP Path message, and in this case it resembles an IP gateway router.

The Transport Layer consists of a number of known capacity links capable of carrying IP traffic. The most likely solution being a set of pre-configured LSPs across an MPLS [6] network. In the diagram, where two tunnels meet, a LSR is pictured - however the LSP itself may consist of a number of LSRs.

Whatever stimulus is used, the EP will request a path across the Transport Layer using a COPS Request as specified in [7]. This is sent to an Admission Manager (AM) in the Management Layer. The Admission Managers are responsible for issuing requests for paths across the network, processing these requests and issuing responses. Only AMs can generate or terminate these message types. The COPS interface is also used to carry the response to these requests.

The Management Layer also has a number of Connection Managers (CMs). Note that AMs and CMs are collectively referred to as Management Nodes (MNs). The CMs provide an administrative functionality to one or more LSRs; namely they monitor the bandwidth along each of the tunnels emanating from the LSRs they control and use this information to decide whether a new session may be routed through that tunnel. When a tunnel is established from a LSR, the LSR indicates this to its associated CM using a suitably extended version of COPS. When a session is admitted to or removed from a tunnel, its new available bandwidth is updated. This information is also forwarded to other Management Layer elements using an advertising method on a periodic basis.

The proposed modifications to the SIP protocol contained within this

draft allow it to perform all messaging within the Management Layer. They therefore provide for a route selection and bandwidth reservation function between AMs and also a topology capacity advertising function between all Management Layer elements.

3.1 Management layer elements

From Figure 1 it can be seen that there are three clearly functionally different elements in the Management Layer: Admission Manager, Edge Connection Manager and Core Connection Manager. It is likely, however, in a real implementation that a single Management Node might be required to perform at least two of these roles, depending upon the physical topology of the Transport Network.

Gibson Category Informational - Expiration April 1999 3

SIP Reservation Extensions October 1999

Certainly, were an AM to be attached to the middle CM, it is easy to see that this CM would continue to function as a Core CM to the existing AMs, but as an Edge CM to the newly added AM.

This section will therefore define the functionality that the three MN types will display.

3.1.1 Admission Manager

The Admission Manager is responsible for initiating and terminating session requests. It controls access to the network via a COPS interface to an Endpoint. It receives requests for new sessions over this interface and issues Decision messages based on the outcome of the resultant negotiation and any defined policies (these policies are outside the scope of this document).

The AM will receive and must store the information contained within topology advertisements to determine suitable paths for INVITE messages. It will not generate any advertisements itself. The advertising scheme operates such that the path to each of the Core CMs and the congestion on each of these paths is well known. Each of the reachable Core CMs also advertises the set of domains that can be reached through it. The AM must store all of this information to enable informed path choice.

The AM may also choose to store the complete 4-hop paths used to reach particular destination AMs for re-use in subsequent INVITEs to those AMs. This, however, has the drawback that the congestion information along the length of these paths will be less well known and that large information tables will have to be stored. However, it does allow for a default path across the network to be specified where known.

The AM is therefore able to perform session request blocking at the edge of a network. If there is insufficient capacity on any of the outgoing links the session request is automatically denied.

The AM must also store congestion information on incoming tunnels too. This information is used in the selection of a session path when multiple INVITEs are used.

3.1.2 Edge Connection Manager

An Edge Connection Manager has a peering relationship with an Admission Manager. It is responsible for advertising a set of Core CMs that is reachable by an AM. It must also advertise to its peer Core CMs the set of AMs that it has connection to - this permits the Core CM to build up its set of reachable domains.

In common with Core CMs, an Edge CM will operate in a similar manner to a SIP proxy by processing and forwarding INVITEs. It will use the

Gibson Category Informational - Expiration April 1999 4

SIP Reservation Extensions October 1999

Record-Route header to ensure its continued presence in the return signalling path.

3.1.3 Core Connection Manager

A Core Connection Manager has connections only to Edge CMs. A Core CM must use the advertisements it receives from these Edge CMs to determine the set of domains that it can reach. It must then advertise this information in a broadcast manner such that it is received by all AMs. (Note: it is the responsibility of the AM to determine whether it has received a reachable domain message from a MN that is one of its Core CMs).

As an AM will only have a clear picture of the network up to a Core CM, when forming a path for an INVITE to traverse it is likely that the next hop after the Core CM address will be wildcarded (See [section 5.2.4](#)). The Core CM will be responsible for correctly forwarding these INVITEs with wildcarded path values.

3.2 Applicability to non-MPLS Networks

Although the remainder of this draft will use the above MPLS-based network as a basis for discussion, the SIP modifications proposed are intended to allow operation over any QoS-capable IP network. Providing a network of IP routers are connected by well-defined Layer Two links whose capacity is well known, our proposal is equally applicable. The term LSR can therefore be read as a

reference to any IP cross-connect between these Layer two links.

4. Basic Protocol Operation

This section shows the basic operation of the modified SIP protocol within the network described above. The messaging used is outlined along with the operation sequence. This is done by way of session messaging examples. The exact content of the messages is dealt with in later sections.

The descriptions of the following message types are based on the four-hop network specified in [3] and pictured in Figure 1. In this case the extended SIP protocol is seen to operate in the Core Management Layer.

4.1 Set-up messaging example

The following is a simple example of the messaging used to establish a session reservation and path within the Management Layer.

4.1.1 INVITE Generation

On receipt of a COPS Request, the AM will form one or more INVITE messages. The process of forming this INVITE follows these steps:

Gibson Category Informational - Expiration April 1999 5

SIP Reservation Extensions October 1999

1) Before any route determination can take place, the AM must determine the resource requirements of the new session. These may come to it in a number of forms, carried by a suitably extended version of COPS, namely:

- SDP [8] session description
- RSVP [9] TSPEC [10]
- RSVP ADSPEC
- RSVP FLOWSPEC
- CR-LDP [11] Traffic Resource TLV
- Simple Bandwidth Object

This information will be coupled with the destination information and used to form a SDP description of the session (see [section 5](#)).

2)(Optional) The originating AM must assign the session a resource class, where this parameter is used when advertising the tunnel characteristics. The mechanism by which a resource class is allocated is to be decided - as yet the MPLS drafts themselves give no clear indication - but will probably be a function of media type, the payment mechanism (where applicable), the DiffServ class or one

of a number of other parameters.

3) Examine available topology information and decide on a route. The AM has a database that contains the information of all its successful session paths plus the information from the advertising REGISTERS that it has received. It can therefore use this information to determine a route for the new session across the network. The proposed extensions to the SIP protocol have been designed such that an incomplete definition of this route can be used. This is implemented by defining part of the path using wildcard values. Note that the chosen path may not have sufficient bandwidth for the session if either another session seizes the resource, or the topology information is out of date, or the path is incompletely specified.

4) Having chosen a path, the AM now assigns it a rank. This rank is based on the AM's view of the congestion of the network, in conjunction with any local policy in force. For example, it is likely that the most favourable path to a particular destination will be that with the lowest congestion. However, a more congested link may have a lower latency and therefore be far more suitable to a real-time session. Also, this allows economic rules to be applied where different service providers own different sections of the Transport Layer. In the case where only one INVITE is sent, the AM may skip this part.

5) The AM must now form a SDP description of the session from the description it is passed by the EP. This is attached to the INVITE in the normal manner for SIP.

6) The AM now forms one INVITE per path selected - the exact number will be dependant on a local policy and may be restricted to reduce the total signaling load on the network. The AM examines the first address in the chosen path and forwards the INVITE to all CMs that match this address. This uses a process like forking in standard SIP.

4.1.1.2 INVITE processing at Connection Manager

When a CM receives an INVITE, it should perform the following processing:

1) Examine the path list and determine if it has a connection to the next hop in the path. If the CM does not have an outgoing tunnel to the next listed CM, it must return an error message to that effect

to the originating AM.

2) The CM now examines the next but one MN in the path too, to ensure this is reachable. Each advertising REGISTER traverses two hops and therefore each CM should have topology information up to two hops away. If none of the available next hop tunnels can be used to reach this two-hop distant MN the CM returns an error message to this effect. Intermediate MNs may read this error message on the path back to the originating AM to update their topology information.

3) The CM now examines the SDP description and determines the resource requirements for the session. Where a choice of next hops is available, the CM may use some policy information to pick how many of the next hop tunnels will be used. Any tunnels that cannot support the session are at this point discarded.

4) The CM can now determine the set of tunnels that can support the session. The CM examines the Call-ID of the INVITE. If a temporary reservation exists for this Call-ID, the tunnel already has sufficient capacity for the session and may be used as a forward path. No extra session bandwidth reservation need be made.

5) Otherwise the CM makes a temporary resource reservation for the new session on each of the tunnels capable of supporting the session. This reservation should match the requirements of the session and will be tagged with the Call-ID of the INVITE that prompted it. No other session may now be offered this part of the tunnel resource.

6) The CM adds its SIP-URL into the Record Route header of the INVITE and then forwards the INVITE to the MNs at the other end of those tunnels.

This process is repeated at each CM along the route to the destination AM. The INVITES will thus fan out towards the centre of the network and converge at the destination AM.

4.1.3 INVITE processing at destination Admission Manager

The destination AM will probably receive a number of INVITES, depending on how tightly defined each original INVITE's path was, how many INVITES were sent and how many of the paths could be used. These will arrive asynchronously, however the AM may process each individually before issuing a response. The AM must perform the

following on each of the INVITEs:

- 1) Examine the Record-Route header and compare it to the topology information the AM has stored. Using this topology information it too assigns a rank to the path in the INVITE. Note this may result in multiple ranks assigned to a single original INVITE due to forking - identical Record Routes should have identical ranks.
- 2) By convolution with the original rank assigned by the originating AM, the preferred path is chosen. This should happen after a time period long enough to allow all INVITEs to arrive but short enough to avoid too much setup latency.
- 3) The INVITE that contained the chosen path is now replied to with a 200 OK. The Record-Route header is copied into a Route header to direct the 200 OK through the MNs that forwarded the INVITE. The same Call-ID must be used in this response, again copied from the INVITE.

```

      +--+      +--+      +--+
    ..1.>|CM|....1*....>|CM|.    .>|CM|.3..
      |11| .      |24| .      .##|31|<###.
      +--+ .      +--+ 1* .# +--+ #.
      1          . .200OK      #.
/\ .          .          .#      #.>/\
|AM|          .          .#      #|AM|
\/.          .          3#      .>\/.
# .          .          .#      .
200OK.      +--+      +--+.#      +--+
#...2.>|CM|....2.....>|CM|...1.....>|CM|.1/1*
#####|13|<##200OK###|29|<#      |36|
      +--+      +--+      +--+

...n... INVITE n Path
##### 200OK Path

```

Figure 2. 200 OK operation

[4.1.4](#) Action of 200 OK Response

The 200 OK response is directed back over the same path as the INVITE to which it is responding. At each CM in the path it has the action of confirming the temporary resource reservation made by the INVITE. This is achieved simply by correlating the Call-ID of the response with the Call-ID used to tag of the temporary reservation. The tag should not be discarded, as it will aid in the removal of the reservation using a BYE message. The principle of this operation can be seen in Figure 2. The left hand AM sends 2 INVITEs that get routed over diverse paths to the destination AM. This AM chooses INVITE 2 as the preferred route and issues a 200 OK back over the same path. At each CM in this path the 200 OK confirms the temporary reservation, made by the INVITE, for the duration of the session. That is, until it is explicitly removed.

4.1.5 200 OK processing at originating Admission Manager

When the 200 OK is received at the originating AM it does the following:

- 1) The AM makes the temporary reservation on the LSP to its Edge CM permanent. The path across the network for the new session request from the EP is now fully specified at the Connection Layer.
- 2) Replies to the EP using a COPS Decision, indicating that the session reservation was successful. This message must contain the list of LSRs that will support the session (and whose CMs have accepted the session). The AM must therefore convert the SIP-URLs in the Route header of the 200 OK into IP addresses. This will be achieved through DNS lookup, either from the set of locally stored values from previous sessions or from a central server.
- 3) If no 200 OK responses are received, the AM may either respond with a COPS Decision rejecting the session, or try another set of routes across the network by issuing further INVITEs.
- 4) To complete the INVITE process, an ACK is returned to the destination AM, providing a 200 OK was received. This message is purely informational.

Once the EP receives a COPS Decision, the session path can be considered fully provisioned and the CR-LDP process can be initiated using the returned LSR address list.

4.2 Session Tear-down example

This section gives an example of how session reservations are removed as the session itself is torn-down.

4.2.1 Tear-down stimulus

Either participant in the session can initiate the tear down of a session. The stimulus for the removal will be a COPS Delete Request State (DRQ) with the same Client Handle as a previously received Request. The AM matches this to the original session Request and re-uses the session information in the teardown process.

4.2.2 Originating AM action

The originating AM will correlate the Client Handle of the DRQ with the Call-ID used to originally establish the reservation. A BYE will now be issued with the corresponding Call-ID and will also include the QoS characteristics used to establish the session, described using SDP.

A Route header must be used in the BYE message that corresponds to the route that the session has reserved across the network. This should again be identified using the Client Handle as an index to a table.

4.2.3 BYE operation at each CM

When a CM receives a BYE, it should remove from its set of confirmed reservations, the one that corresponds to the Call-ID of the BYE. The SDP description can be used as a fail-safe check of the resource to be freed.

If it is unable to perform this then it must resolve the size of reservation described by the SDP description. It must also examine the next SIP-URL in the Route header to determine the tunnel along which the original reservation was made. It then frees this amount of resource from the tunnel. This is accomplished either by removing a specific reservation identified by the Call-ID and re-calculating the total bandwidth.

The CM then forwards the BYE to the next CM in the Route header.

4.2.4 Session Clearing notes

The Removal of reservations in the Management Layer can occur concurrently with the removal of label mappings in the Transport Layer.

The action of a BYE is unidirectional therefore a similar BYE must be issued in the opposite direction to remove the other half of a two-way session. The trigger for this will be provided by the

session application and signaled using a COPS DRQ at the other Endpoint.

4.3 Resource advertisement example

This section offers an overview of how SIP can be extended to advertise topology information. This information includes the size and destination of tunnels and the set of reachable domains via particular CMs.

This advertisement mechanism will be implemented using the REGISTER method that exists within SIP, modified such that the status of the tunnel initially registered is regularly updated using a new message body type that uses a similar syntax to that of SDP.

4.3.1 Initial tunnel advertisement

When a tunnel is established in the Transport Layer, its presence must be indicated to the Management Layer in order for sessions to be routed over it. The mechanism and reasons by which a tunnel is initiated is outside of the scope of the document - it will likely be driven by the owner of the network and may be caused by network initialisation or re-provisioning.

Whatever the reason, the tunnel must always be established using a known set of traffic parameters which are tightly controlled. Once the tunnel exists, the LSR issues a COPS Request to its MN to register the tunnel. The issuing of a COPS Decision is more of a formality, as the MN is unlikely to refuse to allow the tunnel. The Request must include at a minimum the address of the destination LSR and the bandwidth of the tunnel. More likely it will include the full description of its QoS parameters e.g. the Traffic TLV used by CR-LDP. This information will be included in the message body.

When the MN has received this information, it must advertise it for it to be of any use within the Management Layer. The MN that has received the information will be at the upstream end of the tunnel and will be controlling admission to the tunnel for sessions towards the destination LSR. The MN must therefore perform a DNS lookup on the IP address of the destination LSR to determine the SIP-URL of the MN controlling it. It is this SIP-URL that will be used in the advert.

The MN is now able to form a REGISTER message that it will send to all of its peer MNs plus the MN that is the destination of the tunnel. The contents of the REGISTER will be the resource

characteristics of the tunnel and the SIP-URLs of the MNs it runs between. The purpose of this REGISTER message is twofold:

- 1) It establishes two-hop topology information in those upstream (peer) MNs that receive the REGISTER. That is, each of these MNs now has a route to the tunnel destination via the REGISTER sending MN.
- 2) It establishes a peering relationship with the MN at the destination of the tunnel. On receipt of the REGISTER, the

Gibson Category Informational - Expiration April 1999 11


SIP Reservation Extensions October 1999

destination MN now knows that there is a new inbound tunnel to it and therefore a new upstream MN that will want to receive its update REGISTERS.

The above process can be seen in Figure 3. CM_X receives a COPS Request for the new tunnel from LSR_X to LSR_Z. CM_X therefore forms an REGISTER which it sends to its existing peer CM_Y telling it of the new route. CM_X also sends the same REGISTER to CM_Z. The action this time is to inform CM_Z that a new peering relationship should be formed.

In general terms, if a MN receives a REGISTER that has a From: header containing the URI of a MN it does not recognise, this should be regarded as a request for peering from that MN.

Peering REGISTER messages should be sent to the new peer directly, not multicast. A MN without a peer relationship will ignore multicast REGISTERS from a MN it is not yet the peer of, as it has not been told to listen for them.



Existing tunnel	New Tunnel
-----------------	------------

Figure 3. Action of Tunnel Advert

4.3.2 Tunnel update advertisement

Each MN must update the information about the capacity of each of its tunnels on a regular basis. This information should be sent to each of its peers based on the following metrics:

Time - as a baseline, each MN must send an advertising REGISTER on a regular basis at the expiry of a timer - every time an advertising REGISTER is sent, the timer should be restarted;

Gibson	Category Informational - Expiration April 1999	12
--------	--	----

SIP Reservation Extensions	October 1999
----------------------------	--------------

Activity - a MN may choose to send a new advert based on a number of INVITE/BYE messages;

Resource - every time there is a significant change in tunnel resource a MN may choose to send a new ADVERT.

Once a peering relationship is established, the same Call-ID for each subsequent REGISTER to the same peer should be retained to allow easy tunnel identification. The CSeq header should be used to indicate an updated value. The peer should always use the advertising REGISTER information with the highest CSeq value and discard all lower value REGISTERS with the same Call-ID. Multiple tunnel descriptions can be included in a single REGISTER as each description is self-contained (See [section 5.3.2](#)) however care should be taken that the total message length does not get too long.

As the tunnel advert information is contained in a message body, the update information could be piggybacked on other messages as they traverse the network to reduce the total messaging.

4.3.3 Reachable domain advertisement

Over time, a MN will build up a view of the domains it is two hops distant from, based on the REGISTERS it receives from its peers. This information may be advertised too to aid route selection and ranking. This information should be cascaded back to the edges of the Management Layer to the AMs. The To header of each domain REGISTER will be formed from the end descriptor of the tunnel

advertises the Core CM receives. The SIP-URL used in the To header may only ever have a MN-type of AM (see [section 5.3.2](#) for details).

When an AM is choosing a route to a destination AM, it will only have detailed topology information for the first two hops of the journey to a set of reachable Core CMs. It therefore is essential that the AM knows the set of nodes reachable from each of these Core CMs so that the final two MNs in the path to the destination AM can be determined.

The domain information can be gleaned by examining the domain segment of the destination SIP-URL from a tunnel advertisement. A Core CM is capable of choosing an onward route for an INVITE based on its topology information, so by sending an INVITE via a certain Core CM and allowing (through wildcarding) the CM to choose the route for the final two hops of the path, the AM can reduce the total amount of information it needs to store.

5. Protocol details

Gibson Category Informational - Expiration April 1999 13

SIP Reservation Extensions October 1999

The aim of this section is to describe in some detail the format and use of the messages described in brief above. It will concentrate on the areas where there is significant divergence from the standard operation of SIP, either in operation or interpretation of fields, and on the key message elements used where operation is similar.

5.1 Standard SIP

The message format of the standard SIP protocol, represented using Augmented Backus-Naur format (ABNF, as described in [Appendix C](#) of [5]), is as follows:

```
generic-message           = start-line
                           *message-header
                           CRLF
                           [ message body ]

start-line                 = Request-Line |
                           Status-Line

message-header             = ( general-header
                              | request-header
                              | response-header
                              | entity header )
```

The first element is always a start-line, which can be either Request line or Status line. Generally a Request line describes the SIP method that the message is instigating and the Status line shows the nature of a response to such a message. The Request line has the following format:

Request-Line = Method SP Request-URI SP SIP-Version CRLF
e.g. INVITE sip:mark@nortelnetworks.com SIP/2.0

This would be for an INVITE from mark@nortelnetworks.com using version 2.0 of the SIP software. For a response to this INVITE, the following might be used as the Status-Line of the message:

Status-Line = SIP-Version SP Status-Code SP Reason-Phrase CRLF
e.g. SIP/2.0 200 OK

in this case an OK response that has a Status-Code of 200.

After the start-line there are multiple message-headers (denoted by the * in ABNF) that can be any of the shown types except that a request-header is specific to a request and similarly a response-header is specific to a response message. This draft proposes a number of changes to the standard set of these headers as detailed in [section 5.3](#).

Gibson Category Informational - Expiration April 1999 14

SIP Reservation Extensions October 1999

Finally, in the generic message, comes any message body attached to the message. The format and length of the body type is described by the entity-headers at the end of the message-header list. A message body is attached in native format, hence a SDP body is included using its standard text-description. In order to accommodate capacity advertisements we will describe an additional advert body type.

[5.1.1](#) Change of INVITE meaning

The meaning of the INVITE message is changed when compared to the standard SIP protocol. Under normal operation, this message is used to invite the called party to join a session and indicates the type of media that the caller can receive. The indication of what is to be sent is an option.

The INVITE is now used to make a forward reservation for the session from the caller. Therefore the session description contained within the INVITE is that for the forward direction and should be used to

reserve bandwidth in tunnels from the caller to the callee.

5.1.2 Change of REGISTER use

The REGISTER method is currently used to register the address listed in the To header with a SIP server. This draft proposes two alterations to the use of the message.

In the link advertisement we REGISTER the existence of a link to the location described by the To header and advertise its available capacity. The description of this path is then included in a new message body type, as described in [section 5.3.2](#).

Secondly, the set of domains reachable via a particular CM is registered with other MNs. In this instance the To header has no real meaning, with the information stored in the reachable domains body type attached to the REGISTER message, as described in [section 5.3.3](#).

To accommodate this functionality, all MNs must operate as Registration servers for both path and reachable domain REGISTER messages. These REGISTER messages should be sent using a different multicast address than the well-known "all SIP servers" address "sip.mcast.net" (224.0.1.75). It is suggested that two new well-known addresses be used: "sip-path.mcast.net" for path REGISTER messages; "sip-domain.mcast.net" for reachable domain REGISTER messages.

To control the level of signalling information present in the network, the REGISTER messages must be controlled using the TTL field. It is recommended that in the 4-hop network shown in Figure 1, the value TTL = 1 be used for path REGISTER messages and TTL = 2 for reachable domain messages.

Gibson Category Informational - Expiration April 1999 15

SIP Reservation Extensions October 1999

The information in both new REGISTER message types must be periodically updated without the need for outside stimulus. Updates will be indicated by an increased CSeq. Previously stored information must be discarded and the new information used instead.

5.2 Message formats

We now address the specific SIP headers that need to be modified to allow path reservation described within this draft to be performed. Also included in this section is the details of the single new header type required.

5.2.1 From header

The From header takes the standard SIP form. It will only ever have the SIP-URL of the Admission Manager that is originating the INVITE or of the Connection Manager originating the REGISTER method. To aid the determination the MN-type originating the message, it is recommended that the tag is used to identify the type of MN node originating the message.

The From header also has a vital role in the establishment of a peering relationship between adjacent Management Nodes. If a path REGISTER is received with a SIP-URL in its From header of a Management Node that the receiving Management Node does not already have a peering relationship with, the SIP-URL must be added to its peer list.

The from header takes the standard SIP form as described below:

```
From = ( "From" | "f" ) ":" ( name-addr | addr-spec )
      *( ";" addr-params )
name-addr   = [display name] "<" addr-spec ">"
addr-spec   = SIP-URL|URI
display-name = *token|quoted-string
addr-params = tag-param
tag-param   = "tag=" UUID
UUID        = 1*( hex | "-" )
```

e.g. From: "M. Gibson" <sip:mrg@nortelnetworks.com>; tag=AM

The above example shows that the message originated from user mrg whose SIP application has AM functionality.

5.2.2 To header

The main part of the To header is used in the normal SIP manner. It will only ever contain the address of the Admission Manager that is the destination of the INVITE.

Gibson Category Informational - Expiration April 1999 16

SIP Reservation Extensions October 1999

The use of the tag part, however, is to be changed to accommodate the formation of multiple INVITEs for a particular session. It will be used to indicate which of how many total INVITEs for the session this INVITE is.

The To address takes the standard SIP form as described below:

```

To    = ( "To" | "t" ) ":" ( name-addr | addr-spec)
      *( ";" addr-params )
addr-params = tag-param|alternate-param
alternate-param = instance "(" total ")"      ;instance <= total
instance      = 1*digit
total         = 1*digit

```

e.g. To: sip:j.bloggs@sip.net; tag=2(4)

The above example is for an INVITE to j.bloggs and indicates that this is the second of four different INVITEs.

5.2.3 Record-Route header

The Record-Route header is formed in the exact same manner as normal SIP. It must always be used in an INVITE to allow determination of the path over which reservation is being made.

The Record-Route header will never be used in an ADVERT message type. The Record-Route header takes the standard SIP form as described below:

```
Record-Route = "Record-Route" ":" 1# name-addr
```

e.g Record-Route: <sip:a.g.bell@bell-telephone.com>,
<sip:a.bell@ieee.org>

This indicates that the request has traversed the host ieee.org then bell-telephone.com.

5.2.4 Route header

The Route header is used in an INVITE to define the path it should take across the network. In contrast to normal SIP, the list of SIP-URLs may contain non-explicit information. E.g. *@domain.name.com is an acceptable value. Where this is used, the INVITE may be forwarded to all MNs that match the non-asterisked parts of the SIP-URL. The determination of which of these MNs to use is performed using the topology information stored by the MN in conjunction with the next element in the Route header. Only those matches that provide a route to the two-hop distant SIP-URL may be used. The chosen value is written into the Record-Route header as above. The Route header uses the standard SIP form as described below:

```
Route          = "Route" ":" 1# name-addr|wildcard-addr
```

Gibson Category Informational - Expiration April 1999 17

SIP Reservation Extensions October 1999

```
Wildcard-addr  = "*"@" *(" | domainlabel ".") (" | toplabel
```

["."])

e.g Route: <sip:a.g.bell@bell-telephone.com>,
<sip:*@ieee.org>

The Route header of an INVITE should be discarded by the AM that receives it. Only the Record-Route header will contain useful route information.

It is recommended that the Route header is not used in REGISTER messages, thereby reducing the packet size and therefore signaling overhead within the Management Layer.

5.2.4.1 Wildcard Usage

The SIP-URL contained in a Route header may be replaced in total or just part by a wildcard value, as illustrated in the ABNF rule above. This allows the originator to direct an INVITE through a particular domain but to allow local SIP servers to determine the path across that domain.

Where the whole of a Route element is wildcarded the INVITE may be forwarded over any next hop, provided that there is a route to the next but one Route element via the chosen next hop. This forwarding decision should be made using the reachable domain information REGISTER'd for this next hop. See Annex B for further usage examples.

The Route header must always contain the SIP_URL of the destination AM its last element. It must never be wildcarded. This ensures that looping is avoided and allows the looping detection afforded by the use of Via headers to be suppressed. Where there is any doubt about the address of the destination AM, the originating AM must set the TTL header of an INVITE such that it will time-out shortly after it should have reached its destination. In the four-hop network example, this implies a TTL=5.

5.2.5 No-Loop header

The use of wildcards within the route header allows INVITEs to diverge and then converge at the same point as they travel across a network. Under normal operation, SIP proxies will add Via headers to allow for loop detection. However, this would suppress the multiple route finding function that the wildcarded Route header introduces. A No_Loop header is therefore introduced to suppress this action:

No-Loop = "noloop"

Care must obviously be taken when using this header that looping is not introduced. By use of the reachable domains message body and the Route header it is possible to ensure that looping is avoided by

SIP Reservation Extensions

October 1999

rejecting INVITEs whose next-hop Route SIP-URL cannot be reached through the current MN.

5.3 Body formats

This section sets out the body formats needed to perform session reservations. Firstly the extensions needed to SDP are detailed. Then secondly the two new message body types needed to implement advertisement are detailed.

5.3.1 SDP extensions

To exploit SDP so that it properly most profitably interworks with the modified SIP protocol proposed within this draft, both the bandwidth (b) and session information (i) fields need an extended useage and new fields of route favourability (f) and resource class (r) need to be defined.

5.3.1.1 Bandwidth (b)

The current bandwidth definition used by the SDP protocol is used to indicate only the required bandwidth in kilobits per second of the session. While this may sufficient for simple implementations, there is a need to support more complex description of the session to make better use of an MPLS or similar tunnel technology. This also allows for a better mapping of RSVP TSpec and FlowSpec and CR-LDP Traffic TLV parameters.

The bandwidth description will be in terms of the policed session flow from the AM. As the AM will tell the EP what committed rates the network will offer, the AM need only use these rates in INVITE. The process by which the committed rates are derived from the requested rates is outside the scope of this document.

The 3-tuple that defines the peak and maximum date rates is therefore defined to be used in the bandwidth description containing:

Data rate - the mean data rata of the application. (This correlates roughly to the Token Bucket Rate of RSVP and PDR parameter of CR-LDP Traffic TLV.)

Peak rate - the peak data rate of the application. (This correlates roughly to the Token Bucket Size of RSVP and PBR parameter of CR-LDP Traffic TLV.)

Burst rate - either the maximum burst rate that the application might produce, or the maximum burst size that a tunnel can support. (This correlates roughly to the Peak Data Rate of RSVP and EBS

parameter of CR-LDP Traffic TLV.)

The syntax for this SDP description is as shown:
b= <Data rate> <Peak rate> <Burst rate>

Gibson Category Informational - Expiration April 1999

19

SIP Reservation Extensions

October 1999

5.3.1.1.1 Peak vs Burst rate

The use of the Peak and Burst Rate parameters will depend on the context of the bandwidth parameter. In an INVITE, the bandwidth parameter will be used to describe the session requirements. Therefore the Peak rate and Burst rate should be a good match to each other as the AM will provide a policing function for the session.

In an ADVERT the Peak rate may be used to indicate a preferred maximum burst rate for new sessions and the Burst rate can indicate the maximum rate that the tunnel can support. A tunnel may therefore accept a session with a higher Peak rate than the advertised burst rate - this will probably have the action of reducing the advertised Burst rate next time around.

5.3.1.2 Information (i)

The information description of SDP can be used to convey the same information as the tag used in the To header. Namely which of multiple INVITEs this message is. This implies that the INVITEs are forked at the originating AM using Route headers, which may not sit comfortably with the SIP standard. However, it does allow a single 200 OK to be legitimately issued for all the different paths and keeps the differentiation at the session level.

The syntax for this SDP description is as shown:

i=<m> of <n>

Where $m < n$ and n is the maximum number of different routes that an AM may attempt reservation over.

5.3.1.3 Favourability (f)

This type is used to indicate the rank or favourability of the route in the INVITE for a particular session, as perceived by the originating AM. The higher the value assigned, the more favourable the route is. It is recommended that this be an integer value with a pre-agreed maximum value. The method for deriving the favourability value and how it is interpreted by the receiving AM is beyond the

scope of this document.

The syntax for this description is as shown below:

f=<value>

5.3.1.4 Resource Class (r)

The resource class will be decided by the AM for a particular session and used in MPLS networks where multiple tunnels exist

Gibson Category Informational - Expiration April 1999 20

SIP Reservation Extensions October 1999

between LSRs. It must be included in all INVITEs to enable suitable route choice.

The syntax for this description is as shown below:

r=<value>

When used in an INVITE, every effort must be made to match the resource class of the INVITE sdp with that of the tunnel the session is to be forwarded along.

5.3.2 Advert body type

The Advert body type is used to advertise the existence of a tunnel between two LSRs and will be carried by a REGISTER message. It is sent initially to announce the establishment of a new tunnel - part of the peering establishment operation - and then subsequently to update the spare capacity it has left.

The Advert body type takes the same basic syntax form as the SDP protocol and has six descriptors:

Start [s] - the starting LSR of the tunnel
End [e] - the end or terminating LSR of the tunnel
Total Capacity [c*] - the total size of the tunnel
Free Capacity [f] - the available capacity of the tunnel
Latency [l*] - the time taken to traverse the tunnel
Resource Class [r*] - the resource class of the traffic carried by the tunnel.

The descriptors marked with an asterisk are all optional in every Advert and their use will depend upon the complexity of the network

and the parameters used to make routing decisions. This set of descriptors is not seen as necessarily definitive and there is scope to add further descriptors as required.

Each of the above descriptors will now be dealt with in more depth in the following sub-sections.

5.3.2.1 Start [s] descriptor

This is used to indicate the LSR (or other Layer Two element) at the beginning of the path being advertised. The Start descriptor must always be present in an advert message body. Note: the tunnel is seen as running from the start to the end such that the start performs the label mapping of incoming session packets to send them to the end.

Gibson Category Informational - Expiration April 1999 21

SIP Reservation Extensions October 1999

Although it is the LSRs that form the tunnel's endpoints, it is the SIP-URL of the MN controlling the LSR at the start of the tunnel that is advertised. This simplifies the peering process and allows multiple different LSRs controlled by the same MN to appear as one super-LSR. The SIP-URL used to describe the start of the tunnel will be that of the MN controlling the LSR. The syntax for the Start descriptor is shown below:

s=<SIP-URL>
e.g. s=CM_london@SIP.co.uk

There is no assumed policy for the allocation of SIP-URLs to describe the endpoints of a tunnel and is left flexible to allow tailoring to particular situations and architectures.

5.3.2.2 End [e] descriptor

The End descriptor is used to indicate the destination of the path being advertised. It must always be present in an advertisement.

As each advertisement is sent by a MN to each of its peers, the End descriptor not only describes the path destination but also provides reachable domain information as the destination MN of the path is two hops distant from each of these peers. A CM that is acting as a Core CM should therefore log all advertisements that emanate from AM-type SIP-URLs.

The End descriptor has the following syntax:

e=<SIP-URL>
e.g. e=AM_Paris@SIP.fr

5.3.2.3 Total Capacity [c] descriptor

The total capacity descriptor indicates the size of the path being advertised in terms of its bandwidth. The same 3-tuple used in the extended bandwidth descriptor for SDP is re-used here. The tunnel is therefore described in terms of:

Total data rate (kbps) - essentially the capacity of the path

Preferred peak rate (kbps) - the preferred maximum session burst rate

Maximum allowable data burst (kbps) - the largest burst the path can handle.

This descriptor is optional in all Adverts as in the most part, the useful information when forwarding an INVITE is the free tunnel capacity. If advertising the total capacity is deemed necessary, it is recommended that the total capacity be sent only once in an

Gibson Category Informational - Expiration April 1999 22

SIP Reservation Extensions October 1999

initial REGISTER and only re-sent if it changes during the lifetime of the tunnel.

The syntax for the total capacity descriptor is:

c=<total data rate> <peak rate> <max burst>
e.g. c=10000 50 200

5.3.2.4 Free Capacity [f] descriptor

The free capacity descriptor is mandatory in all tunnel REGISTER messages. It uses the same format as the tunnel capacity descriptor but advertises the free or available capacity of the tunnel.

The syntax for the free capacity descriptor is:

f=<available data rate> <peak rate> <max burst>
e.g. f=300 40 200

While the maximum allowable data burst size will likely remain constant, if a large number of bursty sessions have already been routed along the path, the preferred peak/burst rate may be reduced to inhibit the admission of further bursty sessions. The available

data rate will indicate the amount of free bandwidth along the path.

5.3.2.5 Latency [l] descriptor

The latency descriptor describes the time taken to traverse the path and is optional. As with the total path capacity, this figure is unlikely to change over time and so where used it need only be sent once, unless the path changes - note resizing of the path will likely have no effect.

The latency descriptor will have the following syntax:

```
l=<time (ms)>
e.g. l=10
```

The means of determining the latency is outside the scope of this document, however it will always be in milliseconds.

5.3.2.6 Resource Class [r] descriptor

The resource class descriptor describes the MPLS resource class assigned to the path and is optional. The resource class is used to group similar types of traffic together, though its complete use is still poorly defined in MPLS drafts. Where used, it will help in the forwarding of INVITEs along the correct paths for the session type.

The syntax for the resource class descriptor is:

```
r=<resource class>
```

Gibson Category Informational - Expiration April 1999 23

SIP Reservation Extensions October 1999

```
e.g. r=45
```

5.3.2.7 SIP-URL usage in REGISTER messages and their Responses

SIP-URLs have a basic form of user@host.domain. As noted in [section 5.2.2](#), it is proposed that the tag be used to indicate the type of MN node that generated the REGISTER message.

When sending the first in a sequence of REGISTER messages about a particular path, the originating MN must include its MN type in the To header tag.

However, the use of the SIP-URL within the Advert message body will be changed such that the information it contains is unambiguous. The user part will be altered according to the following rules:

1) The user part must take the form <MN-type>_<identifier>

2) The MN-type must be either AM or CM

The recipient of an Advert body type can therefore determine the destination of the path it is advertising without examination of the SIP headers of the message it arrived in.

When a Core CM is logging domain information it should discard all the user part and just use the host.domain information.

5.3.3 Domains body type

The domains body type is used to list the set of reachable domains from a particular Core CM. It will be included in a REGISTER message and will be sent to the AMs the Core CM is serving i.e. those to which it has a REGISTER'd path.

This body type will again use an SDP type notation and will consist of a list of domains. The syntax is as shown:

```
n=<number of domains in list>
d=<domain.name>
```

```
e.g. n=3
      d=london.sip.co.uk
      d=harlow.sip.co.uk
      d=cambridge.sip.co.uk
```

5.4 New Status Codes

To indicate the status of the amended INVITE method proposed within this draft requires the introduction of a number of new Status Codes, to deal with the QoS-based nature of the protocol extensions. A new set of numbers, using the format 8xx will be used to differentiate them from normal SIP codes. The following subsections list these codes and provide guidelines for their usage.

Gibson	Category Informational - Expiration April 1999	24
--------	--	----

SIP Reservation Extensions	October 1999
----------------------------	--------------

5.4.1 Informational Status Codes

The first set are informational error codes and are used between MNs to indicate the progress of a forked INVITE. They may optionally be sent back to the INVITE originator to update topology knowledge. These are denoted by 88x codes.

5.4.1.1 881 No capacity in tunnel

This error message will be issued when the bandwidth of the new

session is greater than the free capacity of the tunnel.

5.4.1.2 882 Not available

This error message indicates that the LSP is not currently available for reasons other than it has reached its maximum capacity. This may be due to some fault in the Connection Layer or because the destination CM/AM has developed a fault. The process by which either of these events is detected is beyond the scope of this document.

5.4.1.3 883 No such Tunnel

This error message is sent back in response to an INVITE that cannot be forwarded to its next-hop AM/CM because there is not now, nor has there been in the recent past, a tunnel to the AM/CM. The definition of recent past will depend on the rate of reconfiguration of tunnel within the network. For very static networks, it will probably be minutes.

5.4.2 Other Status Codes

A Second set of status codes are needed to indicate that an INVITE has failed to make a reservation across a network. These are denoted by 8xx codes.

5.4.2.1 801 No Path

This response will be sent by a forking proxy that has been unable to find a single forward path for the INVITE's resource requirements. This should be sent back to the INVITE originator.

5.4.2.2 802 Unable to change

This error message is sent when there is insufficient spare capacity in the tunnel for a session to increase its reserved bandwidth. This will occur when an INVITE is sent for an existing reservation whose extra bandwidth requirements are in excess of the spare capacity of the tunnel.

6. Messaging Examples

Gibson Category Informational - Expiration April 1999 25

SIP Reservation Extensions October 1999

To show the full use of the extended SIP headers, there now follows a series of example messages. These will also highlight the use of the new body types.

This section gives details of the amended message formats from the

standard SIP protocol and shows the use of the various header and message bodies described in this document.

6.1 INVITE

Having received the COPS Request and decoded its contents, the originating AM is now ready to form an INVITE.

The first action that must be performed is the computation of the number of paths that the AM wants to use to attempt reservation over. The SIP-URL derived from the To information of the COPS Address object will be used as the final element of the Route header for each of these paths. The other SIP-URLs of the Route header will be decided upon using the REGISTER information the AM has. Having decided the set of paths, the AM also assigns a rank to each of them, which will be encoded as a favourability SDP string.

The From and To headers can be directly filled in from the converted COPS Address object. The Record-Route header is added to the INVITE and the SIP-URL of the originating AM is filled in, with AM added as a tag. The Route header is added after this and the Alternative header added to indicate which of the INVITES for the session this is.

A Call-ID will be assigned for the new SIP exchange and a CSeq, so that changes to the session characteristics can be logged.

The SDP description of the session is now appended to the end of the INVITES. The bandwidth string can be formed directly from the COPS Bandwidth object and the favourability string assigned to the particular Route will be added after this. A completed INVITE should therefore be formed as illustrated below. Note that the user strings here contain the MN-type to aid readability. They are not mandatory but may improve message parsing.

```
INVITE sip:AM_Croydon@London.SIP.co.uk SIP/2.1
From: sip:AM_Croydon@London.SIP.co.uk; tag=AM
To: sip:AM_Fort_Worth@Dallas.SIP.com; tag=1(1)
Route: <sip:CM_London_Gateway@London.SIP.co.uk>,
<sip:*@New_York.SIP.com>,
<sip:*@*>, <sip:AM_Fort_Worth@Dallas.SIP.com>
Record-Route: <sip:AM_Croydon@London.SIP.co.uk>
Call-ID: 2496513480@London.SIP.co.uk
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
```

```
v=1
o=croydon 56348948 20159786 IN IP4 47.25.3.106
s=new session
b=8 32 64
f=6
```

The above INVITE is for a session reservation from an AM in Croydon to an AM in Fort Worth, Dallas. It has selected a path through the London CM Gateway, via New York with an unspecified penultimate node to its destination. The CSeq count is set to 1 for the start of a new session - any changes to the session will re-use the unique Call-ID but increment the CSeq by 1. The tag in the To header indicates that this is the INVITE generated for this session. Notice though, that by underspecifying the Route, it is going to end up with a number of paths beyond the London Gateway.

The AM has also used SDP to describe the specifics of the session. It indicates it is the owner of the description - if the original description had been passed up over the COPS interface it would have been modified and used in preference. As the AM has issued the description, only the bandwidth and favourability strings are used - it knows nothing of media types.

Notice the favourability has a rank of 6. Despite the fact that only one Path is specified that might lead you to expect a rank of 10, this clearly is not an optimal solution for this new session.

6.2 200 OK

A 200 OK response to an INVITE will be formed directly from one of the incoming INVITES. The receiving AM will examine the Record-Route header and use its ADVERT information to assign a rank to the path it describes. It does this for each forked and alternative INVITE it receives. Then by convolution with the original rank, it chooses the highest overall ranked path.

So the information used in the response is virtually that used in the chosen INVITE.

```
SIP/2.1 200 OK
From: sip:AM_Croydon@London.SIP.co.uk; tag=AM
To: sip:AM_Fort_Worth@Dallas.SIP.com; tag=1(1)
Route: <sip:CM_Chicago_Transit@Chicago.SIP.com>,
<sip:CM_Manhattan_Gateway@New_York.SIP.com>,
<sip:CM_London_Gateway@London.SIP.co.uk>,
<sip:AM_Croydon@London.SIP.co.uk>
Call-ID: 2496513480@London.SIP.co.uk
CSeq: 1 INVITE
```

Notice that the sdp is no longer needed as the Call-ID identifies


the session and the favourability does not need to be returned as it

has been processed. The Alternative, however, is needed so that the originating EP can determine which Route was successful.

The contents of the Record-Route header of the successful INVITE is used as the Route header for the 200 OK. This is vital to update the temporary reservations at each traversed MN.

6.3 Forking

In this section we will consider the process of forking an INVITE based on the Route header contained in an INVITE. We will re-use the diagram shown in Figure 2, only this time we will specify the Route for each of the two INVITES. For the purposes of this section we will assume that the session characteristics can be supported across the network by each hop in the end-to-end route. We will therefore leave the SDP message body out of the INVITES in this section.



...n... INVITE n Path
200OK Path

INVITE 1: Route {CM11, *, CM36, AM}
INVITE 2: Route {CM13, CM29, CM31, AM}

Figure 4. Routed INVITES

The originating AM issues two INVITES. The first of the two takes the following form:

AM->CM11

```
INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=1(2)
Route: <sip:CM11@SIP.net>, <sip:*@SIP.net>,
<sip:CM36@SIP.net>, <sip:AM_T@SIP.net>
```

Gibson Category Informational - Expiration April 1999

28

SIP Reservation Extensions

October 1999

```
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...
```

with the second INVITE as:

AM->CM13

```
INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=2(2)
Route: <sip:CM13@SIP.net>, <sip:CM29@SIP.net>,
<sip:CM31@SIP.net>, <sip:AM_T@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...
```

Notice that the difference between the two is merely the Route header and the tag on the To header. Both INVITEs are transmitted at the same time towards their differing first CMs. At CM11, the Route header of INVITE 1 is examined to determine where to forward it. The next hop is wildcarded and CM11 has paths to both CM24 and CM29. It therefore logs the session characteristics in a temporary store and adjusts the available capacity of both forwarding tunnels accordingly. It then forwards the following INVITE over both tunnels:

CM11->CM24 and CM11->CM29

```
INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=1(2)
Route: <sip:CM11@SIP.net>, <sip:*@SIP.net>,
```


<sip:CM36@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM11@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...

Meanwhile CM13 only has CM29 as the forwarded path, so it makes a temporary reservation along that path and sends the following:

CM13->CM29

Gibson Category Informational - Expiration April 1999 29

SIP Reservation Extensions October 1999

INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=2(2)
Route: <sip:CM13@SIP.net>, <sip:CM29@SIP.net>,
<sip:CM31@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM13@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...

At CM24, INVITE 1 is received. It determines that it must forward to CM36, so makes a temporary reservation for this forwarding path and sends the following:

CM24->CM36

INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=1(2)
Route: <sip:CM11@SIP.net>, <sip:*@SIP.net>,
<sip:CM36@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM24@SIP.net>, <sip:CM11@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...

At CM29, we will make the assumption that INVITE 2 arrives first. CM29 examines the Route header and determines that CM31 is the next CM in the path. It makes a reservation along the path to that node and sends:

CM29->CM31

```
INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=2(2)
Route: <sip:CM13@SIP.net>, <sip:CM29@SIP.net>,
<sip:CM31@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM29@SIP.net>, <sip:CM13@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...
```

Gibson Category Informational - Expiration April 1999

30

SIP Reservation Extensions

October 1999

Meanwhile it also receives INVITE 1. It examines its Route header and determines that the next CM is CM36. It makes a separate temporary reservation along the path to this CM and forwards:

CM29->CM36

```
INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=1(2)
Route: <sip:CM11@SIP.net>, <sip:*@SIP.net>,
<sip:CM36@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM29@SIP.net>, <sip:CM11@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...
```

At CM31, INVITE 2 is received and CM31 determines that it has a route to AM_T, so it makes a reservation along that path and forwards the INVITE:

CM31->AM_T

```
INVITE sip:AM_0@SIP.net SIP/2.1
```

From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=2(2)
Route: <sip:CM13@SIP.net>, <sip:CM29@SIP.net>,
<sip:CM31@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM31@SIP.net>, <sip:CM29@SIP.net>,
<sip:CM13@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...

At CM36, two versions of INVITE 1 are received. The CM makes a reservation for the first to AM_T but not for the second. As the reservation is indexed using the Call-ID a 200OK response for either INVITE would have the same Call-ID and there is therefore no need for two reservations along the same path. The CM logs itself in the Record-Route headers of both INVITES and forwards them both:

CM36->AM_T

INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=1(2)

Gibson Category Informational - Expiration April 1999

31

SIP Reservation Extensions

October 1999

Route: <sip:CM11@SIP.net>, <sip:*@SIP.net>,
<sip:CM36@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM36@SIP.net>, <sip:CM29@SIP.net>,
<sip:CM11@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...

INVITE sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=1(2)
Route: <sip:CM11@SIP.net>, <sip:*@SIP.net>,
<sip:CM36@SIP.net>, <sip:AM_T@SIP.net>
Record-Route: <sip:CM36@SIP.net>, <sip:CM24@SIP.net>,
<sip:CM11@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE

No-Loop: noloop
Content-Type: application/sdp
Content-Length: 74
...

AM_T will receive all 3 INVITEs and uses their favourability rank along with its own congestion information to choose one of them. In this case, INVITE 2 is chosen and its Record-Route header is used as the Route for the 200 OK response message:

```
SIP/2.1 200 OK
From: sip:AM_0@SIP.net; tag=AM
To: sip:AM_T@SIP.net; tag=2(2)
Route: <sip:CM31@SIP.net>,<sip:CM29@SIP.net>,<sip:CM13@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 INVITE
```

At each of the CMs it traverses, the soft reservation is now made permanent before the response is forwarded.

Finally, the process is completed by sending an ACK:

```
ACK sip:AM_0@SIP.net SIP/2.1
From: sip:AM_0@SIP.net
To: sip:AM_T@SIP.net
Route: <sip:CM13@SIP.net>,<sip:CM29@SIP.net>,<sip:CM31@SIP.net>
Call-ID: 2496513480@SIP.net
CSeq: 1 ACK
```

7. Future Work

- 1) Investigation is needed into the frequency of transmission of advertisement traffic. The information within the messages must remain relevant and timely, however there must not be so many that the network becomes overloaded.
- 2) Control of the number of possible routes that can be examined using a wildcarded Route in an INVITE.
- 3) Control of the number of simultaneous INVITEs to the same destination AM. Where a Call Server model is used, this may be solved by using call-blocking at the destination AM.

- 4) Investigation of the apparent synergy between the extensions proposed in the draft and those proposed by the DCS Group.
- 5) Investigation of the OA&M mechanisms needed to ensure safe delivery of messages and recovery from Layer Two failure.

8. Security Considerations

This draft makes no explicit considerations about security over and above those already made for the basic SIP protocol.

9. Notice Regarding Intellectual Property Rights

Nortel Networks may seek patent or other intellectual property protection for some or all of the technologies disclosed in the document. If any standards arising from this disclosure are or become protected by one or more patents assigned to Nortel Networks, Nortel Networks intends to disclose those patents and license them on reasonable and non-discriminatory terms. Future revisions of this draft may contain additional information regarding specific intellectual property protection sought or received.

10. References

- 1 S. Bradner, "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- 2 S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997
- 3 M. Gibson, "The management of MPLS LSPs for scalable QoS Service Provision", <[draft-gibson-manage-mpls-qos-00.txt](#)>, October 1999.
- 4 "Integrated Services Digital Network User Part (ISDN UP)", CCITT Recommendations Q.761 - Q. 766, (ITU, Geneva).

Gibson Category Informational - Expiration April 1999 33

SIP Reservation Extensions October 1999

- 5 M. Handley et al., "SIP: Session Initiation Protocol", [RFC 2543](#), March 1999.
- 6 R. Callon et al., "A Framework for Multiprotocol Label Switching", <[draft-ietf-mpls-framework-05.txt](#)>, September 1999.
- 7 J. Boyle et al., "The COPS (Common Open Policy Service) Protocol",

<[draft-ietf-rap-cops-07.txt](#)>, August 1999.

8 M. Handley et al., "SDP: Session Description Protocol", [RFC 2327](#), April 1998.

9 R. Braden et al., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.

10 J. Wroclawski, "The use of RSVP with IETF Integrated Services", [RFC 2210](#), September 1999.

11 B. Jamoussi et al., "Constraint-Based LSP Setup using LDP", <[draft-ietf-mpls-cr-ldp-03.txt](#)>, September 1999.

11. Acknowledgments

Thanks go to Tom Taylor, Roy Mauger, Simon Brueckheimer and Dave Stacey of Nortel Networks for their comments on the design of the protocol extensions.

12. Author's Addresses

Mark Gibson
Nortel Networks
London Road, Harlow, Essex, UK
Phone: +44 1279 402621
Email: mrg@nortelnetworks.com

Jon Crowcroft
Department of Computer Science
University College London
Gower Street, London, UK
Email: J.Crowcroft@cs.ucl.ac.uk

Gibson Category Informational - Expiration April 1999 34

SIP Reservation Extensions October 1999

The proposed extensions to SIP have use outside of the management of MPLS network resource. There is much scope for its use in terminal to terminal setup procedures, though this operation requires a slight change in the operational philosophy.

- o A terminal is now seen to act like an AM.
- o The reservation being made by an INVITE is now for a bi-directional path and resource is allocated based on the description in the calling party's INVITE. Where only one media description is listed, this is assumed to reflect the requirements in both directions.
- o Each CM is capable of issuing a CANCEL for a particular INVITE, if the session description in a 200 OK exceeds the resource allocated for the new session. A new Response Code is needed to indicate to the caller, that this event has occurred.

The use of wildcards in the Route header of an INVITE can be used to allow a user to specify the set of service providers they wish to use to place a call, with the best match to the user's needs being chosen.

Imagine the scenario depicted in Figure 5. Terminals A and B are hosted on the networks of two different ISPs. Terminal A wishes to call Terminal B using its local ISP_A but is prepared to let the network decide its best path to Terminal B. ISP_A has a connection to 3 different carriers, namely Carrier_X, Carrier_Y and Carrier_Z.

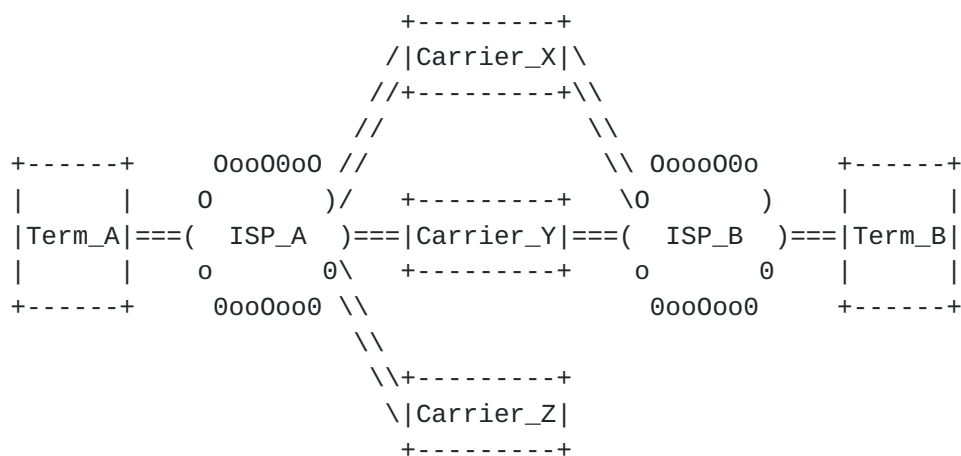


Figure 5. Multi-domain route selection

Terminal B is hosted on ISP_B. ISP_B has connections to two of the above Carriers, Carrier_X and Carrier_Y. For the purposes of this example, we will ignore the session description except to say that it is used to derive a service level requirement for the new session.

Terminal A therefore issues the following INVITE:

```
INVITE sip:Term_A@ISP_A.com SIP/2.1
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_A.com>, <*@*>, <sip:*@ISP_B.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: ..
...
```

Notice that Term_A acts like an AM in this case. When this INVITE is received at the CM of ISP_A, it examines the Route header to determine the next hop. It sees that it has a choice of next hops, as long as they provide a path to ISP_B.

Of the set of possible Carriers, the CM sees that Carrier_Z has no forwarding path and therefore it is immediately discarded. The CM therefore forwards the following INVITE to Carrier_X and Carrier_Y:

```
INVITE sip:Term_A@ISP_A.com SIP/2.1
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_A.com>, <*@*>, <sip:*@ISP_B.com>
Record-Route: <sip:CM@ISP_A.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: ..
...
```

The INVITE is received at both the Carrier networks. At Carrier_X, the SDP requirements are examined and the Carrier decides it is unwilling to support the session. An error response is therefore sent back to the calling terminal, via CM at ISP_A as follows:

```
SIP/2.1 802 Not Available
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
```


Route: <sip:CM@ISP_A.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE

SIP Reservation Extensions

October 1999

The CM at ISP_A should log this response for use in future forwarding decisions.

At Carrier_Y, the SDP requirements are also examined and the Carrier decides it can support the session. It therefore forwards the INVITE to the CM at ISP_B, which has REGISTER'd its connection to Term_B, as below:

```
INVITE sip:Term_A@ISP_A.com SIP/2.1
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_A.com>, <*@*>, <sip:*@ISP_B.com>
Record-Route: <CM@Carrier_Y.com>, <sip:CM@ISP_A.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: ..
...
```

The CM at ISP_B receives the INVITE and determines that it has a connection to Term_B. It is willing to support the new session and therefore forwards the INVITE as below:

```
INVITE sip:Term_A@ISP_A.com SIP/2.1
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_A.com>, <*@*>, <sip:*@ISP_B.com>
Record-Route: <sip:CM@ISP_B.com>, <sip:CM@Carrier_Y.com>,
<sip:CM@ISP_A.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE
No-Loop: noloop
Content-Type: application/sdp
Content-Length: ..
...
```

When Term_B receives the INVITE it chooses to accept the session and therefore the following response is sent back, using the Record-Route header as the Route:

SIP/2.1 200 OK

From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_B.com>, <sip:CM@Carrier_Y.com>,
<sip:CM@ISP_A.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE

When the 200 OK response is received by Term_A, the new reservation is confirmed using an ACK:

Gibson Category Informational - Expiration April 1999

37

SIP Reservation Extensions

October 1999

ACK sip:Term_A@ISP_A.com SIP/2.1
From: sip:Term_A@ISP_A.com
To: sip:Term_B@ISP_B.com
Route: <sip:CM@ISP_A.com>, <sip:CM@Carrier_Y.com>,
<sip:CM@ISP_B.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 ACK

A.1 Return Path Failure

We now consider the above example where the session indicated by the 200 OK to be propagated in the return direction is significantly different to that described in the INVITE. The INVITE processing occurs as above and Term_B forms its 200 OK response, except that the SDP it contains now has a resource requirement significantly greater than that indicated in the INVITE.

The 200 OK is propagated to CM@ISP_B, which is able to cope with the increased load. It now forwards the 200 OK to CM@Carrier_Y.com, which is not able to carry the session under the new resource requirements. It therefore issues a CANCEL for the INVITE to Term_B and sends an error response back to Term_A.

CANCEL sip:CM@Carrier_Y.com SIP/2.1
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_B.com>
Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE

It also issues an 801 response back upstream.

SIP/2.1 801 No Path
From: sip:Term_A@ISP_A.com; tag=AM
To: sip:Term_B@ISP_B.com; tag=1(1)
Route: <sip:CM@ISP_A.com>

Call-ID: 563845256-26@ISP_A.com
CSeq: 1 INVITE

Gibson Category Informational - Expiration April 1999 38

SIP Reservation Extensions October 1999

Full Copyright Statement

"Copyright (C) The Internet Society (date). All Rights Reserved.
This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into

Gibson	Category Informational - Expiration April 1999	39
	SIP Reservation Extensions	October 1999

